The Ronald O. Perelman Center for Political Science and Economics (PCPSE)
133 South 36th Street
Philadelphia, PA 19104-6297

pier@econ.upenn.edu
http://economics.sas.upenn.edu/pier

PIER Working Paper
20-021

Subjective Information Choice Processes

DAVID DILLENBERGER
University of Pennsylvania

R. VIJAY KRISHNA
Florida State University

PHILIPP SADOWSKI
Duke University

March 6, 2020

Subjective Information Choice Processes*

David Dillenberger† R. Vijay Krishna‡ Philipp Sadowski§

March 6, 2020 at 18:39 hrs

Abstract

We propose a class of dynamic models that capture subjective (and hence unobservable) constraints on the amount of information a decision maker can acquire, pay attention to, or absorb, via an Information Choice Process (ICP). An ICP specifies the information that can be acquired about the payoff-relevant state in the current period, and how this choice affects what can be learned in the future. In spite of their generality, wherein ICPs can accommodate *any* dependence of the information constraint on the history of information choices and state realizations, we show that the constraints imposed by them are identified up to a dynamic extension of Blackwell dominance. All the other parameters of the model are also uniquely identified. Behaviorally, the model is characterized by a novel recursive application of static properties.

Key Words: Dynamic Preferences, Information Choice Process, Dynamic Blackwell Dominance, Rational Inattention, Subjective Markov Decision Process

JEL Classification: D80, D81, D90

1. Introduction

In a typical dynamic choice problem, a decision maker (henceforth DM) must choose an action that, contingent on the state of the world, determines a payoff for the

(*) We thank Attila Ambrus, Alexander Bloedel, Andrew Ellis, Itzhak Gilboa, Ed Green, Botond Kőszegi, Jay Lu, George Mailath, Robert Miller, Pietro Ortoleva, Wolfgang Pesendorfer, Todd Sarver, and Ran Spiegler for helpful comments, and the National Science Foundation for support under awards SES 1461469 and SES 1461502. This work began while Dillenberger and Krishna were visiting the Economics departments at NYU and Duke University, respectively. They are most grateful to these institutions for their hospitality.

(†) University of Pennsylvania <ddill@sas.upenn.edu>
(‡) Florida State University <rvk3570@gmail.com>
(§) Duke University <p.sadowski@duke.edu>

current period as well as the collection of actions available in the next period. A standard example is a consumption-investment problem, wherein DM simultaneously chooses what to consume and how to invest his residual wealth, thereby determining the consumption-investment choices available to him in the future, contingent on the evolution of the stock market and retail prices.

Faced with a dynamic choice problem, DM wants to acquire information about the state of the world, but often is constrained by the amount of information he can acquire, pay attention to, or simply absorb. For example, consumers cannot at all times be aware of relevant prices at all possible retailers and firms have limited resources they can expend on market analysis. While accounting for such information constraints can significantly change theoretical predictions (see, for instance, Geanakoplos and Milgrom (1991), Stigler (1961), Persico (2000), and the literature on rational inattention pioneered by Sims (1998, 2003)), an inherent difficulty in modeling them, as well as the actual choice of information, is that they are often private and unobservable.

In this paper we provide — and fully identify — a class of dynamic models that incorporate *intertemporal* information constraints. Just as with intertemporal budget constraints, these constraints have the property that information choice in one period can affect the set of feasible information choices in the future. However, unlike budget constraints, they need not be linear and can accommodate many patterns, such as developing expertise in processing information or feeling fatigued after paying a lot of attention. Indeed, our information constraints can encode arbitrary history dependence.[1] Our framework unifies behavioral phenomena that arise in the presence of such constraints, regardless of their nature; that is, it applies whether the constraints are cognitive, so that individuals have limited ability to take into account available information, or physical, where they reflect the scarcity of information.

To fix ideas, suppose in each period DM can manage his portfolio by choosing from a set of possible investments. Depending on the state of the economy, each choice of investment results in an instantaneous payoff (eg, a dividend) and a new realization of the monetary value of the portfolio, which determines the continuation investment problem for the next period. Further suppose that to improve his portfolio choice, DM can acquire information about the true state in a way that may also affect the feasible set of information strategies in the next period. For instance, it may be that DM is

(1) This is in contrast, for instance, to extant dynamic models of rational inattention, which typically assume that attention constraints (or direct information costs) are time-separable. For a sampling of such models, see Sims (2011).

subject to fatigue, and so can acquire information only if he did not do so in the last period. Alternatively, he may gain expertise, so that acquiring a particular piece of information in one period makes it easier to acquire that same information in subsequent periods. These information constraints may become increasingly complicated as the length of DM's history of past choices grows. The difficulty for the analyst is that while the actual portfolio choice is in principle observable, DM's information choice, and it's impact on the feasibility of subsequent information strategies, is typically not. The following two questions are natural:

(a) Can (unobservable) information constraints be identified from DM's preferences and what type of data is needed to achieve this identification?

(b) What observable choice behavior can be rationalized by unobservable information choices?

Our first main result (Theorem 1) shows that the class of dynamic choice problems we consider is sufficiently rich to identify the entire set of parameters governing DM's preferences over those problems; the parameters being (i) state dependent utilities, (ii) (time varying) beliefs about the state, (iii) the discount factor, and (iv) the information constraint up to a dynamic extension of Blackwell informativeness.[2] This answers question (a) above. Identifying the subjectively controlled process (which we call Information Choice Process, or simply ICP) from behavior is our main conceptual contribution. We should mention that since our focus is on the novel dynamic aspect of the information constraint, we allow *any* history dependence (see examples in Section 2), but make the model as tractable as possible otherwise. We will discuss the behavioral content of the main restrictions we impose in Section 3.3, right after laying out the formal model.

Apart from establishing that our model does not have free parameters, we see three main benefits of this result. First, it establishes that one snapshot of preferences over continuation problems is sufficient to forecast future choice frequencies between continuation problems for any menu of acts (where each act results in a state-contingent lottery that yields current consumption and a new menu of acts for the next period) and after any history of state realizations. Second, in Theorem 2 we use the identification result to show that comparing individuals with respect to their affinity for dynamic choice amounts to comparing who has, in a well defined

(2) Earlier work has demonstrated how to identify information constraints from observed choice data in static settings (see de Oliveira et al. (2017) and Ellis (2018)). We discuss the relation with these papers and others in Section 6.

sense, more informative information plans available. Third, identifying information constraints can be important for policy decisions. Mullainathan and Shafir (2013), for example, suggest that poorer people make suboptimal investment choices because they face too many demands on their time or cognitive resources to fully inform themselves. In order to ameliorate the effects of such constraints, a policy maker must first understand them.

Our second main result (Theorem 3) answers (b), and provides behavioral foundations for the model. The main challenge is that we cannot condition continuation preferences directly on the status of the subjectively controlled process that governs DM's flow of information. We overcome this by showing that it suffices to condition on the history of appropriately constrained choices between acts. In effect, a decision-theoretic version of the dynamic programming principle then lets us impose our axioms recursively, even after histories that have subjective (and hence unobservable) components. The construction of the representation is based on an idea that is similar to self-generation as introduced by Abreu, Pearce, and Stacchetti (1990), but applied to preferences.

The paper is organized as follows. Section 2 presents examples of ICPs and some of the behavioral patterns they can generate. Section 3 introduces the analytical framework, states our utility representation, and describes our notion of comparative informativeness for ICPs. Section 4 establishes our identification result. Section 5 discusses our axiomatic approach, and provides the representation theorem. Section 6 surveys the most related decision-theoretic literature. All proofs are in the appendices.

2. Examples of ICPs and Patterns of Behavior

An ICP is an entirely subjective control problem, which specifies how future information constraints depend on past choices of information. Formally, an ICP is parametrized by a control state θ, a function $\Gamma(\theta)$ that determines the set of feasible partitions of the space of payoff relevant states S, and an operator τ that governs the transition of θ in response to the choice of partition and the realization of $s \in S$. A trivial ICP, where $\Gamma(\theta)$ is always a singleton, corresponds to DM facing an exogenous information process, and so our framework subsumes the standard model of dynamic decision making without information choice. Another special case is where $\Gamma(\theta)$ is independent of θ and is non-trivial, so that the set of available signals is constant over time.

The novelty in our theory is that ICPs can accommodate arbitrary history dependence. We first describe an ICP that is based on models in the literature on managerial decision making and constrains the accumulation of human capital, or expertise, from experience.[3]

Example 2.1. (Geanakoplos and Milgrom (1991) meets Bolton and Dewatripont (1994)). Following Geanakoplos and Milgrom (1991), suppose a manager needs to process information from N different sources and has total time κ to allocate among them each period. Allocating time κ_i to source i yields the partition $P^i_{\theta_i\kappa_i}$, where θ_i is a parameter measuring the manager's efficiency in processing information from source i, and P^i_σ becomes (weakly) finer as σ increases. The manager's information in a given period is the coarsest refinement of all the partitions he chooses. With $\theta = (\theta_1, \ldots, \theta_N) \in \mathbb{R}^N_+$ as the information control state, the information constraint is then

$$\Gamma(\theta) = \left\{ P : P = P^1_{\theta_1\kappa_1} \vee \cdots \vee P^N_{\theta_N\kappa_N} \text{ with } \sum_i \kappa_i \leq \kappa \right\}$$

Bolton and Dewatripont (1994) consider a dynamic setting where processing information of a particular kind in a given period makes it easier to repeat the same task in the future. We can adapt the formulation of Geanakoplos and Milgrom (1991) to this dynamic setting as follows.[4] Interpret κ as the time available between any two periods to allocate among the different information sources. To specify a simple rule for the evolution of the manager's efficiency in processing information, fix $\lambda > 0$ and $\beta \in (0, 1)$, and let $\tau(\theta, (\kappa_i), s) = \theta' \in \mathbb{R}^N_+$ be given by

$$\theta'_i := (1 - \beta)\theta_i\lambda\kappa_i + \beta\theta_i \qquad \text{for } i = 1, \ldots, N$$

The parameter β measures the persistence of expertise over time, so a greater β means that expertise gained in one period carries over to the next to a greater extent. The parameter λ measures the marginal impact on expertise of spending more time on an information source. It is easy to see that θ_i increases from one period to the next if, and only if, $\kappa_i \geq 1/\lambda$ in that period. In the interesting case where $\kappa \leq N/\lambda$, this

(3) We refer exclusively to expertise that improves an individual's ability to make the right decision, rather than the ability to execute that decision. See Currie and MacLeod (2017) for a discussion of the two types of expertise in the context of medical decision making.

(4) Bolton and Dewatripont (1994) do not explicitly model constraints on information processing. Instead, they directly assume that agents learn the most relevant information each period, and that repetition makes the agent better at processing information. Instead, our example allows the agent to choose directly which source of information he repeatedly (and hence more and more efficiently) processes.

amounts to requiring that $\kappa_i/\kappa \geq 1/N$, which, in turns, implies that to gain expertise in processing one source of information, the manager must lose expertise in another.

Our next example builds on the entropy based constraints found in the literature on rational inattention, for example in Maćkowiak and Wiederholdt (2009), and on an analogy between optimal choice of information and a standard consumption-investment problem.

Example 2.2. In each period DM receives an attention income $\kappa \geq 0$. Any stock of attention not used in the current period can be carried over to the next one at a decay rate of β. Let K denote the attention stock in the beginning of a period. Learning the partition P costs $c(P)$, for some cost function c (measured in units of attention and not utils).[5] Formally, with attention stock K, any partition $P \in \Gamma(K) = \{P : c(P) \leq K + \kappa\}$ can be chosen, whereupon the stock transitions to $K' = \tau(K, P) = \beta[K + \kappa - c(P)]$ to determine the continuation constraint. An ICP of this sort is parametrized by the 4-tuple (K_0, κ, c, β) where K_0 is the initial stock of attention. The case with $\beta = 0$ corresponds to a typical per period constraint in the literature on rational inattention.

In order to incorporate expertise in this example, suppose the cost of learning a partition depends on past choices. In particular, if partition Q was chosen in the previous period, then the cost of learning P now is $c(P \mid Q) = (1-b)H_\mu(P) + bH_\mu(P \mid Q)$, where, given a probability μ over S, $H_\mu(P)$ is the entropy of P and $H_\mu(P \mid Q)$ is the relative entropy of P with respect to Q. Note that $H_\mu(P \mid P) = 0$ and hence $c(P, P) = (1-b)H_\mu(P)$. That is, while learning P initially costs $H_\mu(P)$, learning P again in the subsequent period costs only a fraction $(1-b)$ thereof. The parameter b measures the degree to which DM can gain expertise.

Examples 2.1 and 2.2 capture, in alternative ways, the notion of *expertise* in processing information, that is, a complementarity between information processed in different periods. This can lead to a 'locked-in' phenomenon, where DM is reluctant to switch away from familiar choice problems, even in favor of options that are deemed superior in the absence of familiarity.[6]

(5) For example, as is common in the rational inattention literature, $c(P)$ can be the entropy of P calculated using some probability distribution over S.

(6) It has been argued that home bias in portfolio choice among investors who manage their own portfolio (rather than use index funds) is driven by informational advantages — see Coeurdacier and Rey (2013). Evidence that this bias persists in favor of the old home even after a move to a new location — see Massa and Simonov (2006) — nicely illustrates this phenomenon.

ICPs can also accommodate *fatigue* in acquiring information, for instance when acquiring information now diminishes the ability to do so in the future, as in Example 2.2 for $\beta > 0$. A similar mechanism that gives rise to fatigue is resource exhaustion, where DM is endowed with an initial stock of attention which he draws down whenever he chooses to learn.

We note that every ICP has a finite horizon truncation that mimics the original ICP for t periods, after which it admits only the coarsest partition. One of our technical contributions lies in metrizing the space of ICPs and showing that every ICP can be approximated by its finite horizon truncations.

3. Representation with Information Choice Processes

3.1. Domain

Let S be a finite set of objective or observable states. For any compact metric space Y, we denote by $\Delta(Y)$ the space of probability measures over Y, by $\mathscr{F}(Y)$ the set of acts that map each $s \in S$ to an element of Y, and by $\mathscr{K}(Y)$ the space of closed and non-empty subsets of Y.

Let C be a compact metric space representing consumption. A one-period consumption problem is $x_1 \in X_1 := \mathscr{K}\big(\mathscr{F}(\Delta(C))\big)$. It consists of a menu of acts, each of which results in a state-dependent lottery over instantaneous consumption prizes. Then, the space of two-period consumption problems is $X_2 := \mathscr{K}\big(\mathscr{F}(\Delta(C \times X_1))\big)$, so that each two-period problem consists of a menu of acts, each of which results in a lottery over consumption and a one-period problem for the next period. Proceeding inductively, we may define t-period problems as $X_t := \mathscr{K}\big(\mathscr{F}(\Delta(C \times X_{t-1}))\big)$.

Our domain consists of infinite horizon *dynamic choice problems* (henceforth, choice problems, or simply menus) and is denoted by X which is itself homeomorphic to $\mathscr{K}\big(\mathscr{F}(\Delta(C \times X))\big)$.[7] Note that both current and the continuation menus are in X. For any $x, y \in X$ and $t \in [0,1]$, we let $tx + (1-t)y := \{tf + (1-t)g : f \in x,\, g \in y\} \in X$.

A consumption stream is a degenerate choice problem that does not involve choice at any point in time. As we show in Appendix D, the space L of all consumption streams can be written recursively as $L \simeq \mathscr{F}(\Delta(C \times L))$. Thus, each $\ell \in L$ is an act that yields a state-dependent lottery over instantaneous consumption and continuation consumption streams (an $\ell' \in L$). There is a natural embedding of L in X. We analyze DM's preference relation $\succsim$ over X, and denote its restriction to L by $\succsim|_L$.

(7) The homeomorphism is written as $X \simeq \mathscr{K}\big(\mathscr{F}(\Delta(C \times X))\big)$. See Appendix A.2 for details.

The space X of menus subsumes some domains previously studied in the literature. For instance, if S is a singleton, X reduces to the domain considered by Gul and Pesendorfer (2004). Furthermore, if the horizon is also finite, it reduces to the domain in Kreps and Porteus (1978). The subspace L of consumption streams is also a subspace of the domain in Krishna and Sadowski (2014).

3.2. ICP-Representation

Given a menu x, DM chooses a partition in every period. Let $\mathscr{P}$ be the space of all partitions of S. DM's choice of partition is constrained by an *Information Choice Process* (ICP). Formally, an ICP is a tuple $\mathscr{M} = (\Theta, \Gamma, \tau, \theta_0)$, where Θ is a set of control states; the mapping $\Gamma : \Theta \to 2^{\mathscr{P}} \setminus \varnothing$ specifies the set of feasible partitions in a given control state θ; the transition operator $\tau : \mathscr{P} \times \Theta \times S \to \Theta$ determines the transition of the control state θ, given a particular choice of partition and the realization of an objective state; and θ_0 is the initial control state. Let $\mathbf{M}$ be the space of ICPs.

In addition, let $(u_s)_{s \in S}$ be a collection of real-valued continuous functions on C such that at least one u_s is non-constant and let $\delta \in (0, 1)$ be a discount factor. Let Π be a fully connected transition operator for a Markov process on S, where $\Pi(s, s') =: \pi_s(s')$ is the probability of transitioning from state s to state s' and $\Pi(s, s') > 0$ for all $s, s' \in S$. Let $0 \notin S$ be an auxiliary state, and denote by π_0 the unique invariant measure of Π.

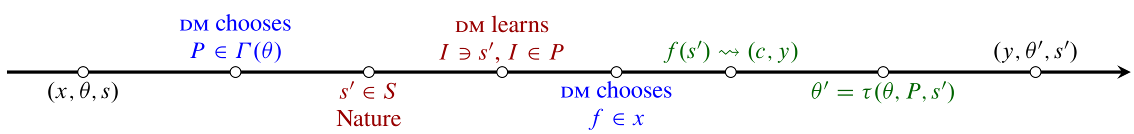

Figure 1: *Timeline*

Our model suggests the following timing of events and decisions, as illustrated in Figure 1. DM enters a period facing a menu x, while being equipped with a prior belief π_s over S and an information control state θ. He first chooses a partition $P \in \Gamma(\theta)$. For any realization of a cell $I \in P$, DM updates his beliefs using Bayes' rule to obtain $\pi_s(\cdot \mid I)$ and then chooses an act $f \in x$. At the end of the period, the true state s' is revealed and DM receives the lottery $f(s')$, which determines current consumption c and continuation menu y for the next period. At the same time, a new control state $\theta' = \tau(\theta, P, s')$ and a new belief $\pi_{s'}$ are determined for next period.

DM's objective is to maximize expected utility, which consists of state-dependent consumption utilities and the discounted continuation value, as we now define.

Definition 3.1. A preference $\succsim$ on X has an ICP *representation* $\big((u_s)_{s\in S}, \delta, \Pi, \mathscr{M}\big)$ if it is represented by a function $V(\cdot, \theta_0, 0) : X \to \mathbb{R}$, where $V : X \times \Theta \times (S \cup \{0\}) \to \mathbb{R}$ satisfies

$$V(x,\theta,s) = \tag{Val}$$

$$\max_{P\in\Gamma(\theta)} \sum_{I\in P} \left[\max_{f\in x} \sum_{s'\in I} \mathsf{E}^{f(s')} \Big[u_{s'}(c) + \delta V(y, \tau(P,\theta,s'), s') \Big] \pi_s(s' \mid I) \right] \pi_s(I)$$

In the representation above, for each $s' \in S$, $f(s') \in \Delta(C \times X)$ is a probability measure over $C \times X$, so that $\mathsf{E}^{f(s')}$ is the expectation with respect to it.[8]

A dynamic information plan prescribes a choice of $P \in \Gamma(\theta)$ for each tuple (x, θ, s). Thus, an ICP describes the set of feasible information plans available to DM. The next proposition ensures the existence of the value function and an optimal dynamic information plan.

Proposition 3.2. Each tuple $\big((u_s)_{s\in S}, \delta, \Pi, \mathscr{M}\big)$ induces a unique function $V : X \times \Theta \times S \cup \{0\} \to \mathbb{R}$ that is continuous on X and satisfies [**Val**]. Moreover, an optimal dynamic information plan exists.

The proofs of all the results from this section are in Appendix B.

3.3. Restrictions

We now discuss four restrictions that the ICP representation in Definition 3.1 imposes, and comment on the behavioral patterns to which these restrictions correspond.

First, the payoff relevant state is observed at the end of each period (and hence s appears as a state variable in the function [**Val**]). This feature is implied by two standard assumptions that we make: (i) DM fully understands the continuation problem he receives; and (ii) preferences are continuous. To see this, note that if DM did not

(8) One of the central properties of dynamic choice is dynamic consistency, which requires DM's ex post preferences to agree with his ex ante preferences over plans involving the contingency in question. Because our primitive is ex ante choice between menus, we cannot investigate dynamic consistency directly in terms of behavior. However, our representation [**Val**] describes behavior as the solution to a dynamic programming problem with state variables (x, θ, s), so that implied behavior is dynamically consistent contingent on those state variables. The novel aspect is that the state θ is controlled by DM and is not observed by the analyst.

freely learn the state, then he would be willing to pay a premium for the information content of fully revealing acts (that is, acts that pay different continuation problems in different states), and these acts are dense in the set of all acts.[9] Assumption (i) is in contrast to some models of bounded attention that, rather than focusing on intertemporal constraints, study the implication of choosing without full awareness of the choice set; see, for instance, Gabaix (2014) for analysis of static problems where budget sets are misperceived, and Lian (2019) for a dynamic model where DM has an imperfect perception of his total wealth.

Second, observed preferences over menus are according to the stationary (ie, ergodic) distribution, π_0, of the Markov process that governs the evolution of states. This property is implied if preferences are stationary on the subdomain of consumption streams, since then the same beliefs are used for the evaluation of future consumption acts (those acts that have no continuation values beyond their instantaneous payoffs in a certain period), independently of the date of consumption. The interpretation is that DM does not learn the state in the period prior to the observed choice, and aggregates state-dependent preferences accordingly, using π_0 as his prior beliefs. One could instead assume that DM does learn the realization of the state in the period prior to his initial choice. Formally, this would not cause any complication; it would simply mean replacing our primitive, $\succsim$, with a state-dependent family of initial preferences. Of course, induced preferences in future periods are already state dependent, so that aggregated ex ante preferences can be thought of as an expositionally convenient summary of state dependent preferences.

Third, learning in our model is via partitions of the space of payoff relevant states, that is, signals are deterministic contingent on the true state. In general, signals could be noisy, and since the state space is given, it is not without loss of generality to restrict the class of permissible information structures. Deterministic signals are not essential for our results, but for technical reasons we rely on DM choosing from finite sets of finite valued information structures; while this finiteness can be imposed in a variety of ways, partitional learning is a parsimonious way to achieve it.

(9) One way DM may be able to learn s for free at the end of a period is to place a side bet on a fully revealing act with arbitrarily small stakes. Once the state is learned, there is no additional explorative motive for consumption choice; the choice of act has no information value, just as information choice has no direct consumption value; the two interact only through the instrumental value of information. Steiner, Stewart, and Matějka (2017) solve a dynamic rational inattention problem that also allows for the free incorporation of information at the end of each period. In their model the information cost or constraint is history-independent and invariant.

Finally, ICPs can generate opportunity costs of information acquisition by tightening future information constraints. Our model allows us to focus entirely on the behavioral implications of this new type of dynamic cost. While information constraints are a common modelling choice, an alternative way to model limitations on information acquisition is via direct information costs, measured in consumption 'utils'.[10] In the presence of information costs, whether or not an information plan is ever optimal (so that its cost could potentially be identified), would depend on the instrumental value it can generate over time. With a compact consumption space (and hence bounded utilities) as in our model, this would in turn depend on all preference parameters, namely state-dependent consumption utilities, beliefs, and the discount factor (as, for instance, in Ergin and Sarver (2010)). While a remedy for this problem in static models may be found with unbounded consumption and utilities (see de Oliveira et al. (2017)), in a dynamic setting such as ours the existence of a recursive value function cannot be guaranteed with unbounded utility; additional unappealing assumptions on the domain and/or preferences are required. Our focus on information constraints allows us to steer clear of these complications. In particular, the dominance order on $\mathbf{M}$, which we introduce in the next subsection and characterize behaviorally in Theorem 2, is independent of the other preference parameters.

3.4. Comparative Informativeness of ICPs

As noted in Section 3.2, an ICP can be viewed as circumscribing the set of available dynamic information plans. We now show that the space of ICPs has a natural order.

Partitions can be compared in terms of fineness, which coincides with Blackwell's comparison of informativeness. To extend this idea to ICPs, first consider only how two ICPs $\mathscr{M}$ and $\mathscr{M}'$ differ in the first period. Notice that as far as dynamic information plans are concerned, all that matter are the partitions each ICP permits. This suggests the following one-period order: $\mathscr{M}$ *one-period Blackwell dominates* $\mathscr{M}'$ if for every $P' \in \Gamma'(\theta_0')$, there exists $P \in \Gamma(\theta_0)$ such that P is finer than P'.

In turn, this suggests a natural extension to two-periods. $\mathscr{M}$ two-period Blackwell dominates $\mathscr{M}'$ if for every $P' \in \Gamma'(\theta_0')$, there exists $P \in \Gamma(\theta_0)$ such that (i) P is finer than P', and (ii) for all $s \in S$ and for every $Q' \in \Gamma'(\tau'(P', \theta_0', s))$, there exists $Q \in \Gamma(\tau(P, \theta_0, s))$ such that Q is finer than Q'. Thus, for any information plan in $\mathscr{M}'$,

(10) See, for example, Ergin and Sarver (2010), Woodford (2012), Caplin and Dean (2015), de Oliveira et al. (2017), and Hébert and Woodford (2016).

there is another plan in $\mathscr{M}$ that is more informative in every period and state.

To extend our construction to more periods, we note that requirement (ii) above amounts to the continuation ICP $(\Theta, \tau(P, \theta_0, s), \Gamma, \tau)$ one-period Blackwell dominating $(\Theta', \tau'(P, \theta_0', s), \Gamma', \tau')$. In a similar fashion, we inductively define an order extending Blackwell dominance to t periods, whereby one ICP t-period Blackwell dominates another if for each information plan from the latter, there is another plan from the former that is more informative in the first period and, for any s, leads to a more informative $(t-1)$-period plan starting in the second period.

To illustrate, consider two ICPs, $\mathscr{M}$ and $\mathscr{M}'$, for which the left and right panels of Figure 2 display the respective first two periods. Both ICPs allow DM to commit at the outset to learn either partition P or Q for two successive periods, where P and Q are not ordered in terms of fineness. In the left panel (depicting $\mathscr{M}$) DM can alternatively postpone the choice of partition until the second period — at the cost of not learning anything (ie, learning $\{S\}$) in the first period. It follows that $\mathscr{M}$ two-period Blackwell dominates $\mathscr{M}'$, but not vice versa. To see this, note that every two period information plan available on the right is also feasible on the left, while only the constraint on the left allows the following plan: Pick $\{S\}$ in the first period, wait for the second-period consumption problem to realize, and then choose one of the partitions P or Q.

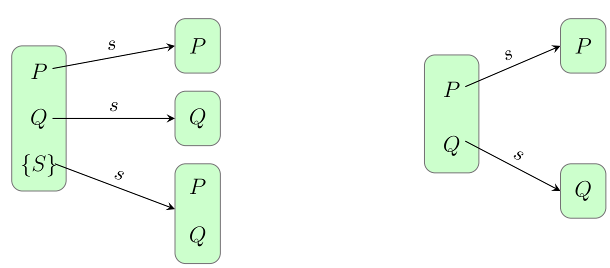

Figure 2: *Two-period* ICPs*, $\mathscr{M}$ and $\mathscr{M}'$*

As another example, consider the ICPs introduced in Example 2.2 where attention stock is drawn down, decays, and is renewed with attention income. Such an ICP is parametrized by the 4-tuple (K_0, κ, c, β). Consider now two ICPs $\mathscr{M}^i$, for $i = a, b$, parametrized by $(K_0, \kappa, c^i, \beta)$ that only differ in their costs of acquiring information. It is easy to see that $\mathscr{M}^a$ one period Blackwell dominates $\mathscr{M}^b$ if, and only if, $c^a \leq c^b$ (ie, $c^a(P) \leq c^b(P)$ for all $P \in \mathscr{P}$). Similarly, $\mathscr{M}^a$ t-period Blackwell dominates $\mathscr{M}^b$ if, and only if, $c^a \leq c^b$.

Ordering ICPs over the infinite horizon is more delicate, because arbitrary ICPs may not have a final period of non-trivial information choice, and hence may not permit backwards induction. Instead, we exploit the recursive structure of ICPs, so that our ordering of informativeness for arbitrary ICPs is also recursive.

Proposition 3.3. There exists a unique largest order that satisfies the following: For all $\mathscr{M}$, $\mathscr{M}^\dagger \in \mathbf{M}$, $\mathscr{M}$ dominates $\mathscr{M}^\dagger$ if for all $P^\dagger \in \Gamma^\dagger\big(\theta_0^\dagger\big)$ there is $P \in \Gamma(\theta_0)$ such that (i) P is (weakly) finer than $P^\dagger$, and (ii) $(\Theta, \Gamma, \tau, \tau(P, \theta_0, s))$ dominates $\big(\Theta^\dagger, \Gamma^\dagger, \tau^\dagger, \tau^\dagger(P^\dagger, \theta_0^\dagger, s)\big)$ for all $s \in S$.

We refer to this largest order as the *dynamic Blackwell order*. It is reflexive and transitive. We say that $\mathscr{M}$ strictly dynamically Blackwell dominates $\mathscr{M}'$ if $\mathscr{M}$ dynamically Blackwell dominates $\mathscr{M}'$, but not vice versa: and that they are dynamically Blackwell equivalent if each dominates the other.

We note that there are other ways to define dynamic extensions of the static Blackwell order; see, for instance, Greenshtein (1996) and de Oliveira (2018), which we further discuss in Section 4.1 when explaining our identification result. Our approach differs from these in that instead of comparing signal processes, we compare controlled signal processes that allow the decision maker to choose his signal (in our case, partition) as a function of the (potentially private) past. That our approach is particularly well suited to our problem is demonstrated by our main identification result, Theorem 1, in the next section.

4. Unique Identification

Theorem 1. *Let $\big((u_s), \delta, \Pi, \mathscr{M}\big)$ be an* ICP *representation of $\succsim$. Then, the functions $(u_s)_{s\in S}$ are unique up to the addition of constants and a common scaling, δ and Π are unique, and $\mathscr{M}$ is unique up to dynamic Blackwell equivalence.*[11]

The proof of all the results in this section is in Appendix C. On the subdomain L, V satisfies Independence as it is independent of $\mathscr{M}$, and is thus completely characterized by the parameters $((u_s), \delta, \Pi)$. Krishna and Sadowski (2014, Corollary 5) show that such a representation on L is unique up to the addition of constants and a common scaling of (u_s). Our challenge then is to identify the ICP $\mathscr{M}$. In Section 4.1 we discuss the main ideas behind the identification strategy for finitely many periods; the extension

(11) In other words, for any additional representation of $\succsim$ with parameters $\big((u_s^\dagger), \delta^\dagger, \Pi^\dagger, \mathscr{M}^\dagger\big)$, it is the case that $\delta^\dagger = \delta$, $\Pi^\dagger = \Pi$, $u_s^\dagger = au_s + b_s$, for some $a > 0$ and $b_s \in \mathbb{R}$ for each $s \in S$, and $\mathscr{M}$ and $\mathscr{M}^\dagger$ dynamically Blackwell dominate each other.

to infinite horizon is similar, but involves technical issues that we discuss in detail in Appendix C.

Notice that our model is a Markov decision process for DM with state (x, θ, s), where x and s (and their transitions) can be verified by the analyst, while θ cannot. Thus, the Markov decision process is subjective with partially unknown transitions and partially unobservable actions and states. Theorem 1 achieves full identification of this subjective Markov decision process. To the best of our knowledge, this is the first result of this sort in the literature.

An immediate benefit of identifying all the parameters is that it allows a meaningful comparison of decision makers. The next result demonstrates that dynamic Blackwell dominance plays the same role in our environment as does Blackwell dominance in a static setting.

Consider two decision makers with preferences $\succsim$ and $\succsim^\dagger$, respectively. We say that $\succsim$ has a *greater affinity for dynamic choice* than $\succsim^\dagger$ if for all $x \in X$ and $\ell \in L$, $x \succsim^\dagger \ell$ implies $x \succsim \ell$.[12] The comparison in the definition implies that $\succsim$ and $\succsim^\dagger$ have the same ranking over consumption streams in L.[13] While consumption streams require no choice of information, a typical choice problem x may allow DM to wait for information to arrive over multiple periods before making a choice. This option should be more valuable the more information plans DM's ICP renders feasible. The uniqueness established in Theorem 1 allows us to formalize this intuition.

Theorem 2. *Let $((u_s), \delta, \Pi, \mathscr{M})$ and $\left((u_s^\dagger), \delta^\dagger, \Pi^\dagger, \mathscr{M}^\dagger\right)$ be* ICP *representations of $\succsim$ and $\succsim^\dagger$ respectively. The preference $\succsim$ has a greater affinity for dynamic choice than $\succsim^\dagger$ if, and only if, $\Pi = \Pi^\dagger$, $\delta = \delta^\dagger$, $(u_s)_{s \in S}$ and $(u_s^\dagger)_{s \in S}$ are identical up to the addition of constants and a common scaling, and $\mathscr{M}$ dynamically Blackwell dominates $\mathscr{M}^\dagger$.*

Theorem 2 connects a purely behavioral comparison of preferences to dynamic Blackwell dominance of ICPs, which is independent of preferences, and hence of utilities and beliefs. This indicates a duality between our domain of choice and the information constraints that can be generated by ICPs, a theme we will return to in Section 4.1. A useful corollary of Theorem 2 is the following characterization of the dynamic Blackwell order: $\mathscr{M}$ dynamically Blackwell dominates $\mathscr{M}^\dagger$ if, and only if,

(12) This definition is the analogue of notions of 'greater preference for flexibility' in the dynamic settings of Higashi, Hyogo, and Takeoka (2009) and Krishna and Sadowski (2014).

(13) That is, $\ell \succsim \ell'$ if, and only if, $\ell \succsim^\dagger \ell'$ for all $\ell, \ell' \in L$. This is Lemma 34 in Appendix F of Krishna and Sadowski (2014), and uses the fact that both $\succsim$ and $\succsim^\dagger$ satisfy Independence on L.

every discounted expected utility maximizer prefers to have the ICP $\mathscr{M}$ instead of $\mathscr{M}^\dagger$ regardless of the menu in question.[14]

4.1. Identifying the ICP

We now illustrate the main idea behind the identification of the ICP and also discuss an alternative notion of Blackwell dominance in dynamic settings. To simplify matters, suppose for the rest of this section that (i) consumption is in the set $C = [0, 1]$, (ii) utilities are state independent (so $u_s = u$ for all $s \in S$), and (iii) u is strictly increasing, with $u(0) = 0 = 1 - u(1)$. Rather than providing a general, more abstract intuition, we will base our discussion first on a static example, and then on an ICP that allows non-trivial information acquisition only in the first two periods; the same ideas extend to any finite horizon truncation, and a limit argument extends them to arbitrary ICPs.

We start with identification in the static setting. For each $J \subset S$, define the simple act $f_{1,J}$ by

$$f_{1,J}(s) := \begin{cases} (1, \mathbf{1}) & \text{if } s \in J \\ (0, \mathbf{0}) & \text{if } s \notin J \end{cases}$$

where $\mathbf{c}$ means receiving $c \in C$ in perpetuity independently of the state. To test if DM is able to learn some partition that is weakly finer than P, consider the menu $x_1(P) := \{f_{1,J} : J \in P\}$. Contingent on event $J \in P$, the act $f_{1,J}$ delivers permanent consumption 1 with certainty. Therefore, upon learning an event in Q that is weakly finer than P, the optimal consumption strategy will guarantee the same consumption, so that $V(x_1(P); Q) = 1 = V(x_1(P); P)$. Conversely, if Q is not finer than P, then with positive probability DM will learn a set of states I which is not a subset of any $J \in P$. In that case, any choice of act from $x_1(P)$ will generate 0 with positive probability, and hence $V(x_1(P); Q) < 1 = V(x_1(P); P)$.

In order to extend this intuition to two periods, recall the two ICPs, $\mathscr{M}$ and $\mathscr{M}'$, respectively, in the left and right panels of Figure 2 (Section 3.4). Suppose the analyst believes DM's information constraint is either $\mathscr{M}$ or $\mathscr{M}'$. How can she verify it is $\mathscr{M}$, and not $\mathscr{M}'$? To answer this, consider the act $f_{2,\{S\}}(s) := \big(1, \mathrm{Unif}\{x_1(P), x_1(Q)\}\big)$ where $\mathrm{Unif}\{y, z\}$ is the uniform lottery over y and z, and the menu $x_2(\{S\}) = \{f_{2,\{S\}}\}$. Note that $x_2(\{S\})$ requires no choice in the first period, but instead offers the bet

(14) This result thus generalizes the seminal characterization of the standard Blackwell order for partitions, according to which P is finer than Q if, and only if, every decision maker prefers P to Q regardless of the (static) choice problem.

Unif$\{x_1(P), x_1(Q)\}$ that provides choice from either $x_1(P)$ or $x_1(Q)$ in the second period. To guarantee consumption of 1, DM must therefore have the option to choose in the second period whether to learn (at least) P or Q, which is not feasible under $\mathscr{M}'$. Therefore, $V(x_2(\{S\}); \mathscr{M}) = 1 > V(x_2(\{S\}); \mathscr{M}')$.

In order to fully identify $\mathscr{M}$, we also need to distinguish it from ICPs other than $\mathscr{M}'$. For instance, in order to test whether DM can learn the partition P twice in a row, let $f_{2,J}$ be the act that pays $f_{2,J}(s) = (1, x_1(P))$ if $s \in J$ and $(0, \mathbf{0})$ otherwise, and define the two-period choice problem $x_2(P) := \{f_{2,J} : J \in P\}$. An analogous construction for learning Q twice in a row yields $x_2(Q)$. It is easy to see that for any $y \in \{x_2(P), x_2(Q), x_2(\{S\})\}$, $V(y, \mathscr{M}) = 1$. Moreover, for any ICP $\mathscr{M}''$, we have that $\mathscr{M}''$ dynamically Blackwell dominates $\mathscr{M}$ if, and only if, $V(y, \mathscr{M}'') = 1$ for all $y \in \{x_2(P), x_2(Q), x_2(\{S\})\}$. In essence, each such y amounts to a betting game, where in each period DM is told a random and history dependent partition and is asked to bet on the correct event in it to receive payoff 1 and stay in the game, rather than exiting the game and receiving 0 indefinitely. Such a betting game generates the same value as a constant stream of 1 if, and only if, DM can learn the relevant partition in each period. We say that the collection of menus $\{x_2(P), x_2(Q), x_2(\{S\})\}$ is *strongly aligned* with the ICP $\mathscr{M}$.

An analogous construction of a strongly aligned set of menus is possible for any ICP. Consider, then, two ICPs that do not dynamically Blackwell dominate each other, and the two corresponding sets of betting games. At least one of the two sets contains a game in which it will be possible to stay forever with certainty under the ICP the set is strongly aligned with (generating the same value as consuming 1 forever) but not under the other ICP (generating a lesser value). In other words, the ICP in our model is identified up to dynamic Blackwell dominance.

4.2. Remarks

4.2.1. *An Alternative Dynamic Extension of Blackwell Dominance*

We now compare our notion of dynamic Blackwell dominance over ICPs with an alternative definition, which was introduced in Greenshtein (1996) and further used by de Oliveira (2018). To ease exposition, we again confine our attention in the definition below to only two periods.

Definition 4.1. Let $\mathscr{M}$ be an ICP. A (random) two-period sequence of partitions $(P_0, (P_{1,s}))$ is feasible in $\mathscr{M}$ if $P_0 \in \Gamma(\theta_0)$ and $P_{1,s} \in \Gamma(\tau(\theta_0, P_0, s))$. Let $\mathscr{M}'$ be another

ICP. Then, $\mathscr{M}$ *two-period sequentially Blackwell dominates* $\mathscr{M}'$ if for any sequence of partitions $(Q_0, (Q_{1,s}))$ feasible in $\mathscr{M}'$, there is another sequence $(P_0, (P_{1,s}))$ feasible in $\mathscr{M}$ such that (i) P_0 is finer than Q_0, and (ii) $P_{1,s}$ is finer than $Q_{1,s}$ for all $s \in S$.

To understand the discrepancy between the comparative notion in Definition 4.1 and our notion of dynamic Blackwell dominance, let $\mathscr{M}$ and $\mathscr{M}'$ again be two ICPs that correspond to the left and right panels in Figure 2, respectively. Observe that both $\mathscr{M}$ and $\mathscr{M}'$ two-period sequentially Blackwell dominate each other: $\mathscr{M}$ contains more sequences of partitions than $\mathscr{M}'$, but any of the additional sequences in $\mathscr{M}$ starts with $\{S\}$ in the first period and is dominated by some sequence feasible in $\mathscr{M}'$.

Notice that in the menu $x_2(\{S\})$ from Section 4.1, DM learns whether he faces $x_1(P)$ or $x_1(Q)$ only in the second period. Our notion allows for plans that condition on this information, thereby allowing DM to postpone the choice of information in $\mathscr{M}$ for the second period until this uncertainty has resolved, a plan that is unavailable under $\mathscr{M}'$. The sequential Blackwell order does not allow information plans to condition on the resolution of uncertainty regarding the menu of choices available, and therefore treats $\mathscr{M}$ and $\mathscr{M}'$ as being equivalent.

4.2.2. *Non-Identification of General Markov Decision Processes*

There are negative results in the econometric literature about the identifiability of subjective Markov decision processes. For instance, Rust (1994) and Magnac and Thesmar (2002) show that in a general Markov decision processes, where (i) utilities depend on the Markov state and (ii) choice affects the stochastic evolution of that state, observing (stochastic) choice is insufficient to identify the evolution of shocks or other parameters of the model. Recall that the Markov state in our model is (x, θ, s). Our identification problem is even more ambitious than the one mentioned above, because we have truly unobservable control variables: θ cannot be observed by the analyst, and the set of available information choices given θ as well as the transition of θ as a function of the information choice are both unknown. Crucially for us, the set of available observable actions in Markov state (x, θ, s) depends only on x, and the analyst can effectively observe choice from all possible continuation problems for the same combination of s and θ. This contrasts with the aforementioned econometric literature where the analyst does not have access to a rich set of tradeoffs at a given state. For instance, even though in Magnac and Thesmar (2002) the evolution of the Markov state is observable by the analyst, the distribution of (iid) taste shocks cannot be identified because DM always has the same set of actions to choose from. Similarly, in Rust (1994), the set of feasible actions is completely determined by the Markov state, and therefore precludes observations from a rich set of continuation problems in a given Markov state.

It is important to note that the question here is not one of observing ex ante menu choice versus observing ex post choice frequency from a menu. For example, in a stochastic choice framework, Gul and Pesendorfer (2006) show that with random choice data from a rich set of menus, distributions over taste shocks can be identified. It is the ability to observe choices from a rich set of alternatives at any state that is essential for identification.

4.2.3. *Inference from Limited Data*

Our identification strategy suggests that inference about the ICP can be made from a small number of observations. In particular, inference benefits from three of its features. First, identification of the ICP is (almost) independent of the other preference parameters, as it only uses the best and worst outcomes in each state (1 and 0 in the example discussed in Section 4.1). Second, while identification of the ICP relies on randomization over continuation problems, the exact probabilities used in this randomization are not important; for example, we could replace the uniform distribution in the proof of Theorem 1 with any distribution with full support. Finally, to verify whether DM can follow a particular information plan, the analyst only needs to observe *one* appropriate binary choice — between the best consumption stream and a choice problem that is strongly aligned with the plan in question. For instance, the question whether DM is able to follow the information plan that chooses $\{S\}$ in the first period and from $\{P, Q\}$ in the second period, as in the ICP $\mathscr{M}$ in Figure 2, is answered immediately by offering DM a choice between the betting game $x_2(\{S\})$ and the consumption stream $\mathbf{1}$. Indifference is observed if, and only if, DM can follow the information plan in question (or one that is more informative than it).[15]

5. Behavioral Characterization

In this section we provide an axiomatic foundation for our model. Before describing our axioms in section 5.1, we first explain our approach.

(15) The last observation, in particular, can be useful in many applications. For instance, a policy maker may not need to know exactly what information constraint individuals face, but may only wish to know whether they are able to process the specific information needed to behave optimally vis a vis a new policy or not.

Consider static preferences over dynamic choice problems that are stationary, that is, contingent on the decision relevant state, induced preferences over continuation problems are the same as the original preferences. In that case, a dynamic value function is commonly derived from a representation of the static preferences via a recursive operator that ensures that the induced value function over continuation problems has the same form as the original static representation, and so forth.

As is evident from [**Val**], preferences over menus in our model are stationary conditional on the state (x, θ, s). To replicate the dynamic extension of a static value function in terms of preferences, we would thus like to (i) take as given a tight axiomatic description of static behavior, (ii) define a recursive operator that ensures that induced preferences over continuation problems also satisfy those static axioms, and (iii) combine the restrictions derived this way with additional dynamic axioms as needed. In order to falsify the model, it would then suffice to find a violation of the static axioms for some history of state realizations. Our main conceptual challenge in following this approach is that the decision relevant state is now (x, θ, s), and its second component is not observable by the analyst. We will show how to use the history of observable choices to control for θ.

Since the work of extending static axioms to the dynamic environment will be done by the recursively applied operator, it is possible to replace a more primitive set of static axioms by a static representation that has already been shown to be equivalent to them. The static representation that will be our starting point (Axiom 0 below) is derived in Dillenberger, Krishna, and Sadowski (2020).

5.1. Axioms

We are interested in the structure of a binary relation (a preference) $\succsim \subset X \times X$. Our first axiom assumes a representation that has all of the properties of the ICP-representation, except that it is static, that is, the evaluation of continuation problems does not take into account the recursive structure of our domain. This representation is axiomatized in Dillenberger, Krishna, and Sadowski (2020) based on appropriate relaxations of separability, strategic rationality and independence.

Axiom 0. The relation $\succsim$ is represented on X by $V_0 \in \mathbb{R}^X$, which has the form

$$V_0(x) = \max_{P \in \mathbb{Q}} \sum_{I \in P} \left[\max_{f \in x} \sum_{s \in I} \mathsf{E}^{f(s)} \left[u_s(c) + v_s(y, P) \right] \pi(s \mid I) \right] \pi(I) \qquad [5.1]$$

where $\mathbb{Q} \subset \mathscr{P}$ is a set of partitions of S, $u_s \in \mathbf{C}(C)$ and $v_s(\cdot, P) \in \mathbf{C}(X)$[16] for each $s \in S$ and $P \in \mathbb{Q}$, with the property that for all $P, P' \in \mathbb{Q}$ and $s \in S$, we have $v_s(\cdot, P)|_L = v_s(\cdot, P')|_L$.

Recall from Section 3.1 the homeomorphism $X \simeq \mathscr{K}\big(\mathscr{F}(\Delta(C \times X))\big)$. Axiom 0 is static because it ignores that X shows up on both sides of this homeomorphism. To emphasize this, let us for a moment write $X = \mathscr{K}\big(\mathscr{F}(\Delta(C \times W))\big)$ and require only that W is a compact metric space. The representation in [**5.1**] can then be interpreted as follows: DM chooses a partition to learn about the payoff relevant state, taking into account that the outcome has two components. The value of the first component, $u_s(c)$, is independent of the chosen partition P, while the value of the second component, $v_s(w, P)$, depends directly on P. This is plausible since we interpret each w as a menu that may benefit from future information, and the set of information choices available in the future may depend on past information choices.

Note that in display [**5.1**] we can discard partitions that are never optimal for any $x \in X$. Thus, we may restrict attention to *minimal* representations.[17]

Before recursively applying Axiom 0, we pose one additional axiom that captures the special role played by consumption streams, which are dynamic but leave no choice to be made in the future and therefore require no information (that is, all information alternatives perform equally well). The axiom requires $\succsim$ to satisfy additional standard assumptions when consumption streams are involved.

In what follows, for any $c \in C$ and $\ell \in L$, let (c, ℓ) be the constant act that yields consumption c and continuation stream ℓ with probability one in every state $s \in S$. Let $\ell^*, \ell_* \in L$ be such that $\ell^* \succsim \ell \succsim \ell_*$ for all $\ell \in L$.

Fix $c \in C$ and let

$$\ell_s(s') = \begin{cases} \ell_* & s' \neq s \\ (c, \ell) & s' = s \end{cases}$$

Define the induced relation $\succsim_s$ on L by $\ell \succsim_s \hat{\ell}$ if $\ell_s \succsim \hat{\ell}_s$. It will follow from Axiom 1 below that $\succsim_s$ is independent of the particular choice of c.

Axiom 1 (Consumption Stream Properties)**.**

(a) L-Nontriviality: $\ell^* \succ \ell_*$.

(b) L-Independence: For all $x, y \in X$, $t \in (0, 1]$, and $\ell \in L$, $x \succ y$ implies $tx + (1-t)\ell \succ ty + (1-t)\ell$.

(16) We denote by $\mathbf{C}(Z)$ the space of continuous functions on a space Z.

(17) The collection $\mathbb{Q}$ is minimal if any subset of it gives a lower utility for some menu x.

(c) L-History Independence: For all $\ell, \hat{\ell} \in L$, $c \in C$, and $s, s', s'' \in S$, $(c, \ell_s) \succsim_{s'} (c, \hat{\ell}_s)$ implies $(c, \ell_s) \succsim_{s''} (c, \hat{\ell}_s)$.

(d) L-Stationarity: For all $\ell, \hat{\ell} \in L$ and $c \in C$, $\ell \succsim \hat{\ell}$ if, and only if, $(c, \ell) \succsim (c, \hat{\ell})$.

(e) L-Indifference to Timing: $\frac{1}{2}(c, \ell) + \frac{1}{2}(c, \ell') \sim (c, \frac{1}{2}\ell + \frac{1}{2}\ell')$.

Axiom 1(b) resembles the C-Independence axiom in Gilboa and Schmeidler (1989), and is similarly motivated: Because consumption streams require no information choice, mixing two menus with the same consumption stream should not alter the ranking between these menus. For a discussion of properties (c) through (e) see Krishna and Sadowski (2014).

As noted in Section 3.1, the domain L is recursive. Therefore, using Axiom 0, we can write $\ell^*(s) = (c_s^+, \ell_s^+)$ and $\ell_*(s) = (c_s^-, \ell_s^-)$, where $c_s^+, c_s^- \in C$ are (respectively) the u_s-maximal and -minimal elements, and $\ell_s^+, \ell_s^- \in L$ are (respectively) the v_s-maximal and -minimal elements in L for each $s \in S$.

5.2. A Recursively Formulated Representation Theorem

As noted above, the history of state variables that are relevant for DM consists of the observed history of the state s, and the history generated by DM's unobserved choices of information. The challenge in imposing axioms on continuation preferences is that continuation preferences can (and typically will) depend on those unobserved information choices which naturally vary with the menu in question. This conceptual difficulty arises already in two-period problems, which we now focus on to simplify the discussion. The solution for longer horizons will then be immediate.

The key is to note that, due to the finiteness of S, if P is DM's unique optimal information choice at a menu x in state s, then P will also be optimal under small perturbations of x. In that case, the menu- and state-dependent preferences over continuation problems can be inferred from preferences over perturbations of the continuation problems available in x.

Formally, we provide a definition, in terms of $\succsim$, for a finite x to have a unique optimal partition. (In the static utility representation [**5.1**] above, this corresponds to one partition generating a strictly higher value than any other available partition.) For such x that has a unique optimal partition, if (i) DM prefers to simultaneously perturb all acts $f \in x$ by mixing the continuation problems they specify in s with y rather than mixing them with y', and (ii) x and the perturbed menus have the same optimal partitions (we refer to such menus as concordant and also define concordance based on $\succsim$ below), then we write $y \succsim_{(x,s)} y'$. Because x has a unique optimal partition, sufficiently small perturbations do not upset concordance with x, and so $\succsim_{(x,s)}$ is well defined on X (ie, is complete and transitive), and therefore amenable to the imposition of our axioms.

For each $f, g \in \mathscr{F}(\Delta(C \times X))$, $s \in S$, and $\varepsilon \in [0,1]$, define

$$\big(f \oplus_{\varepsilon,s} g\big)(s') := \begin{cases} (1-\varepsilon)f(s') + \varepsilon g(s') & s' = s \\ f(s') & s' \neq s \end{cases}$$

That is, $f \oplus_{\varepsilon,s} g$ perturbs the lottery $f(s)$ in state s in the direction of $g(s)$, while leaving it untouched in all other states. Fix a $c \in C$. For $x, y \in X$, define $x \oplus_{\varepsilon,s} y := \{f \oplus_{\varepsilon,s} (c, y) : f \in x\}$ which represents the menu where each act in x is perturbed in state s in the same way.

Our goal is to find a behavioral notion that allows us to say that DM has partition P available. Generalizing our construction in Section 4.1, we now let $x_1(P)$ denote the menu that requires DM to bet on the cells of the partition P, where betting correctly now pays (c_s^+, ℓ^*) (instead of $(1, \mathbf{1})$), and a wrong bet pays (c_s^-, ℓ_*) (instead of $(0, \mathbf{0})$). Then, a partition at least as fine as P is available in the first period, if, and only if, $x_1(P) \sim \ell^*$. For a choice problem x we then have $\frac{1}{2}x + \frac{1}{2}x_1(P) \sim \frac{1}{2}x + \frac{1}{2}\ell^*$ if, and only if, some partition that is at least as fine as P is optimal for x. Thus, the same collection of partitions is optimal for two menus x and y, if for all $P \in \mathscr{P}$ we have $\frac{1}{2}x + \frac{1}{2}x_1(P) \sim \frac{1}{2}x + \frac{1}{2}\ell^*$ if, and only if, $\frac{1}{2}y + \frac{1}{2}x_1(P) \sim \frac{1}{2}y + \frac{1}{2}\ell^*$.

Definition 5.1. Choice problems x and y are *concordant*, if $\frac{1}{2}x + \frac{1}{2}x_1(P) \sim \frac{1}{2}x + \frac{1}{2}\ell^*$ if, and only if, $\frac{1}{2}y + \frac{1}{2}x_1(P) \sim \frac{1}{2}y + \frac{1}{2}\ell^*$ for all $P \in \mathscr{P}$.

Definition 5.2. If $x = \frac{1}{2}x' + \frac{1}{2}x_1(P) \succsim \frac{1}{2}x' + \frac{1}{2}x_1(Q)$ for all Q, with strict preference if P is not finer than Q, then we say that P is the uniquely optimal information choice for x. If such a P exists, we say that x has a unique information choice.[18]

To understand this definition, notice that $x = \frac{1}{2}x' + \frac{1}{2}x_1(P) \succsim \frac{1}{2}x' + \frac{1}{2}x_1(Q)$ for all Q implies that there is an information choice that is optimal for x' and is (weakly) finer than P.[19] Because $\frac{1}{2}x' + \frac{1}{2}x_1(P) \succ \frac{1}{2}x' + \frac{1}{2}x_1(Q)$ if Q is not finer than P, learning

(18) Since all information choices in this paper are partitions, we also refer to optimal information choices as optimal partitions.

(19) If there is no feasible partition at least as fine as P, then the representation implies $\frac{1}{2}x' + \frac{1}{2}x_1(\{S\}) \succ \frac{1}{2}x' + \frac{1}{2}x_1(P)$; ie, there is some partition finer than $\{S\}$, the coarsest partition, but none finer than P.

exactly P must be optimal for x', and any other optimal information choice for x' must consist of a partition that is coarser than P. But in that case, P must be the unique optimal partition for $x = \frac{1}{2}x' + \frac{1}{2}x_1(P)$.

For a fixed P, denote the representation in [5.1] by $V_0\big(x; P\big)$, and define X'_P as follows:

$$X'_P := \big\{x : V_0(x) = V_0\big(x; P\big) \text{ for some } P \in \mathbb{Q} \text{ and}$$
$$V_0(x) > V_0\big(x; Q\big) \text{ for all } Q \in \mathbb{Q} \text{ such that } P \neq Q\big\}$$

Thus, X'_P is the set of all menus for which P is the unique optimal information choice according to V_0. If $x \in X'_P$ then it follows from Axiom 0 that X'_P is the set of menus that are concordant with x and that also have a unique optimal partition.

Definition 5.3. Let x have a unique optimal partition, $y, y' \in X$, $s \in S$, and $\varepsilon \in (0, 1)$ such that $x \oplus_{\varepsilon,s} y$, $x \oplus_{\varepsilon,s} y'$, and x are pairwise concordant. Then, $y \succsim_{(x,s)} y'$ if $[x \oplus_{\varepsilon,s} y] \succsim [x \oplus_{\varepsilon,s} y']$.

Lemma E.6 in Appendix E.1 shows that for any $x \in X'_P$ and $y, y' \in X$, there exists an $\varepsilon \in (0, 1)$ such that $x \oplus_{\varepsilon,s} y, x \oplus_{\varepsilon,s} y' \in X'_P$, and that the choice of $c \in C$ in the definition of $x \oplus_{\varepsilon,s} y$ is irrelevant, so that $\succsim_{(x,s)}$ is indeed well defined.

Recall that our objective is to impose restrictions on the functions $v_s(\cdot, P)$ in [5.1], but that these functions are not directly observable. Nevertheless, Lemma E.6 further establishes that for every $v_s(\cdot, P)$ in a minimal representation, there exists some x (indeed, all $x \in X'_P$) such that $v_s(\cdot, P)$ represents $\succsim_{(x,s)}$, and so behavioral restrictions on $\succsim_{(x,s)}$ correspond to structural restrictions on $v_s(\cdot, P)$. With $\succsim_{(x,s)}$ defined above, we can now recursively apply Axiom 0.

Let Ψ_1 be the space of all binary relations on X that satisfy Axiom 0. That is,

$$\Psi_1 := \{\ \succsim \text{ on } X : \succsim \text{ satisfies Axiom 0 }\}$$

Define the operator $\mathscr{B} : 2^{\Psi_1} \to 2^{\Psi_1}$ by:

$$\text{[5.2]} \qquad \mathscr{B}(\Psi) := \left\{ \succsim \in \Psi_1 : \begin{array}{l} \succsim_{(x,s)} \in \Psi \text{ for all } s \in S \text{ and finite } x \in X \\ \text{with a unique optimal partition} \end{array} \right\}$$

This allows us to define the set $\Psi_2 := \mathscr{B}(\Psi_1)$, which consists of all preferences that satisfy our static axiom with the additional property that continuation preferences (ie, $\succsim_{(x,s)}$ for all finite x with a unique optimal partition) also satisfy the static Axiom 0, by virtue of being in Ψ_1. We can extend this construction to arbitrary finite horizons by inductively defining $\Psi_{n+1} := \mathscr{B}(\Psi_n)$ for all $n \geq 1$. The construction of Ψ_n for $n \geq 2$ parallels the recursive construction of our domain of menus described in Section 3.1.

For the infinite horizon, we proceed to the limit, and define the set Ψ^* as the largest fixed point of the operator $\mathscr{B}$ in [**5.2**], which is easily seen to be monotone. It can be shown that $\Psi^* = \lim_{n\to\infty} \Psi_n = \bigcap_{n\geq 1} \Psi_n$. We say that a preference relation on X *satisfies Axiom 0 recursively* if it is in Ψ^*, or equivalently, lies in Ψ_n for all $n \geq 1$.

Our representation theorem below characterizes the set of preferences in Ψ^* via a well defined recursive value function, and establishes that it is non-empty.

Theorem 3. *Let $\succsim$ be a binary relation on X. Then, the following are equivalent:*

(a) *The relation $\succsim$ satisfies Axiom 0 recursively and satisfies Axiom 1.*

(b) *There exists an* ICP *representation of $\succsim$.*

As just argued, the recursive application of Axiom 0 requires not only preferences, but also induced preferences over the next period's continuation problems to be in Ψ^*, thereby requiring preferences over continuation problems two periods ahead to be in Ψ^*, and so forth.[20] To falsify this requirement, one needs to find some finite history of states and sequence of menus after which continuation preferences are defined (each menu must be finite and have a unique optimal partition) and violate Axiom 0.

We conclude by sketching the proof of Theorem 3. Our starting point is Axiom 0, which provides us with the following representation:

$$V(x) = \max_{P\in\mathbb{Q}} \sum_{I\in P} \left[\max_{f\in x} \sum_{s\in I} \mathsf{E}^{f(s)} \left[u_s(c) + v_s(y,P) \right] \pi(s \mid I) \right] \pi(I)$$

We say that V is implemented by $((u_s), \mathbb{Q}, (v_s(\cdot, P)), \pi)$. This representation already has all the features we need to establish, except that it is static; it does not exploit the recursive structure of X.

Adapting the terminology of Abreu, Pearce, and Stacchetti (1990), we define the set Φ^* of *self-generating* value functions, where each $v \in \Phi^*$ is implemented by some tuple $((u'_s), \mathbb{Q}', (v'_s(\cdot, P)), \pi')$ in a way that each $v'_s(\cdot, P)$ is itself in Φ^* (see Appendix E.1). In Lemma E.5 in Appendix E, we show that for every P that is part of an undominated information plan, there is a menu x for which P is the unique optimal partition. We then rely on Axiom 1 and the recursive application of Axiom 0 to show that the representation V of $\succsim$ can be made self-generating.

(20) Due to L-Stationarity (Axiom 1(d)) preferences over continuation problems also satisfy Axiom 1; see Lemma E.8.

The remainder of our construction in Appendix E.3 has two main components. First, we construct an ICP $\mathscr{M} = (\Theta, \Gamma, \tau, \theta_0)$ from the self-generating representation. According to V, we let $\Gamma(\theta_0) = \mathbb{Q}$. Next, if $v_s(\cdot, P)$ in the representation is implemented by $((u'_s), \mathbb{Q}', (v'_s(\cdot, P)), \pi')$, we let $\Gamma(\tau(P, \theta_0, s)) = \mathbb{Q}'$, and so on.

Second, we apply Axioms 0 and 1 to establish a Recursive Anscombe-Aumann representation $V^*_L(\cdot, \pi_0)$ of $\succsim|_L$, where $V^*_L(\cdot, \pi_0) := \sum_s \pi_0(s) V^*_L(\cdot, s)$, and

$$V^*_L(\ell, s) := \sum_r \Pi(s, r)\, \mathsf{E}^{\ell(r)} \left[u_r(c) + \delta V^*_L(\ell', r)\right]$$

for some tuple $\big((u_s), \delta, \Pi\big)$.[21] We then note that the self-generating representation V above and the recursive representation V^*_L must agree on L. This lets us conclude that all the utilities in V can be taken to be (u_s) and all the subjective beliefs can be taken to be the Markovian beliefs Π. We can then pair the parameters $((u_s), \delta, \Pi)$ with $\mathscr{M}$ to find the ICP representation, $((u_s), \delta, \Pi, \mathscr{M})$, which is recursive on all of X. Intuitively, the lack of recursivity in the self-generating representation, which conditions only on the objective state s, is absorbed by the evolution of the subjective state θ in our representation, so that the representation becomes recursive when conditioning on both s and θ.

5.3. Special Cases

The special case of the ICP representation in which DM faces the same information constraint each period (independent of past information choices) is of interest due to its simplicity and its frequent use in dynamic models of rational inattention, where there is a periodic time invariant upper bound on information gain. Given an ICP representation of $\succsim$, one can verify that the behavioral content of this restriction is the following axiom:

Axiom 2. $x \succsim x'$ for all $x' \in X$ if, and only if, $x \succsim_{(y,s)} x'$ for all $x', y \in X$ and all $s \in S$.

Note that $x \succsim x'$ for all $x' \in X$ implies that x is one of the $\succsim$-maximal menus. In particular, x must be as good as the best consumption stream ℓ^*. Therefore, DM must be able to process enough information to choose from x in a way that also

(21) This is the Recursive Anscombe-Aumann representation in Krishna and Sadowski (2014). See Appendix D for a discussion of why the parameters (δ, Π) are unique, and the collection (u_s) is unique up to a common positive affine transformation.

guarantees c_s^+ in every period and any state. It is then clear from our identification strategy that two ICPs that give rise to the same collection of $\succsim$-maximal menus must be dynamically Blackwell equivalent. The axiom ensures stability of this collection, and hence of ICPs up to dynamic Blackwell equivalence, along three dimensions: stationarity, because between $\succsim$ and $\succsim_{(y,s)}$ a period has passed; menu independence, through the comparison of $\succsim_{(y,s)}$ and $\succsim_{(y',s)}$; and history independence, through the comparison of $\succsim_{(y,s)}$ and $\succsim_{(y,s')}$.

One can also confirm that, as in the static setting of Dillenberger et al. (2014), imposing full Independence on $\succsim$ corresponds to the special case where information is determined by a trivial choice process: it arrives exogenously over time.

6. Related Literature

We now comment on the most related menu choice literature.

Kreps (1979) studies choice between menus of prizes. He rationalizes monotonic preferences — those that exhibit preference for flexibility — via uncertain utility functions that are yet to be realized. Dekel, Lipman, and Rustichini (2001) show that by considering menus of lotteries over prizes, those utilities can be taken to be expected utility functions. Dillenberger et al. (2014) subsequently show that preference for flexibility over menus of acts corresponds to uncertainty about future beliefs about the objective state of the world. Ergin and Sarver (2010) and de Oliveira et al. (2017) replace Independence with Aversion to Randomization to model subjective uncertainty that is not fixed, but a choice variable. The former studies costly contemplation about future tastes, while the latter studies rational inattention to information about the state.

None of the models discussed so far are dynamic or let DM react to information arriving over multiple periods. Krishna and Sadowski (2014) provide a dynamic extension of Dekel, Lipman, and Rustichini (2001), where the flow of information is taken as given by DM. In particular, Krishna and Sadowski (2014) assume Independence, which means that there are no subjective or unobservable controls. Their subjective state space in each period is the space of vN-M utility functions. Their recursive domain consists of acts that yield a menu of lotteries over consumption and a new act for the next period. When all menus are degenerate, their domain reduces to the set of consumption streams L, as it does here. The key difference between the two domains lies in the timing of events: Instead of acts over menus of lotteries, we consider menus of acts over lotteries, which are appropriate for a dynamic extension of Dillenberger et al. (2014). Our model also extends de Oliveira et al. (2017), in the sense that the choice of information in a period now affects the feasible choices of information in the future.[22] Table 1 summarizes the position of our model with respect to the papers discussed thus far.

Information	***Uncertainty about vN-M taste***	***Uncertainty about state of world***
Static, fixed	Dekel, Lipman, and Rustichini (2001)	Dillenberger et al. (2014)
Static, choice	Ergin and Sarver (2010)	de Oliveira et al. (2017)
Dynamic, fixed	Krishna and Sadowski (2014)	*This paper*
Dynamic, choice process	★	*This paper*

Table 1: *Summary of Literature*

In our model, DM controls the flow of information over time. Thus, his preferences are interdependent across time and we can no longer appeal to the stationarity assumptions of Krishna and Sadowski (2014). To deal with this complication, we first assume that preferences over consumption streams, $\succsim|_L$, satisfy the standard axioms, including Stationarity, because future information plays no role when there is no consumption choice to be made in the future. We then use the ranking of consumption streams to 'calibrate' preferences over all dynamic choice problems.

Our recursive application of axioms can be compared to standard stationarity assumptions, as in Gul and Pesendorfer (2004) or in Krishna and Sadowski (2014). There, instantaneous and continuation preferences are required to be identical, rather than merely be of the same class. It is easy to see that if stationary preferences satisfy Axiom 0, then they also satisfy it recursively. A related axiomatization of dynamic preferences is in Maccheroni, Marinacci, and Rustichini (2006), who provide foundations for dynamic variational preferences. Their continuation preferences, like ours, satisfy the same type of axioms at any instant as the original preferences. The main difference is that their continuation preferences need only condition on the observable state. In contrast, in our model of endogenous information choice, past menus affect past information choices which, in turn, can affect the current information

(22) Both Dillenberger et al. (2014) and de Oliveira et al. (2017) permit more general information structures than partitions, and the latter also allows for explicit costs of acquiring information.

constraint. Also, their horizon is finite, so that they must (and more easily can) directly impose their axioms for every history of the observable state.

Additional works address unobservable information acquisition in different formal contexts. Piermont, Takeoka, and Teper (2016) study a decision maker who learns about his uncertain, but time invariant, consumption taste (only) through consumption, and so has some control over the flow of information. For static choice situations, the literature based on ex post choice partly parallels the menu-choice approach. Lu (2016) shows how information that, like in Dillenberger et al. (2014), is subjective yet not subject to choice, can be fully identified from random choice. Ellis (2018) identifies a partitional information constraint from ex post choice data. Caplin and Dean (2015) use random choice data to characterize a representation of costly information acquisition with more general information structures. They then proceed to consider stochastic choice data under the assumption that attention entails entropy costs, as do Matějka and McKay (2015). To our knowledge, there is no counterpart to our recursive analysis of information choice in the random-choice literature.

Appendices

A. Preliminaries

Appendix A.1 describes the relevant metric on the space of probability measures. Appendix A.2 describes our (recursive) domain of infinite horizon dynamic choice problems. Appendix A.3 describes canonical ICPs and shows that every ICP is isomorphic to a canonical ICP.

A.1. Metrics on Probability Measures

Let (Y, d_Y) be a metric space and let $\Delta(Y)$ denote the space of probability measures defined on the Borel σ-algebra of Y. For a function $\varphi \in \mathbb{R}^Y$, the Lipschitz seminorm is defined by $\|\varphi\|_{\mathrm{L}} := \sup_{y \neq y'} |\varphi(y) - \varphi(y')| / d_Y(y, y')$, and the supremum norm is $\|\varphi\|_\infty := \sup_y |\varphi(y)|$. This allows us to define the bounded Lipschitz norm $\|\varphi\|_{\mathrm{BL}} := \|\varphi\|_{\mathrm{L}} + \|\varphi\|_\infty$. Then, $\mathrm{BL}(Y) := \{\varphi \in \mathbb{R}^Y : \|\varphi\|_{\mathrm{BL}} < \infty\}$ is the space of real-valued, bounded, and Lipschitz functions on Y.

For $\alpha, \beta \in \Delta(Y)$, define $d_D(\alpha, \beta) := \frac{1}{2} \sup \left\{ \left| \int \varphi \, \mathrm{d}\alpha - \int \varphi \, \mathrm{d}\beta \right| : \|\varphi\|_{\mathrm{BL}} \leq 1 \right\}$, which is the Dudley metric on $\Delta(Y)$. Theorem 11.3.3 in Dudley (2002) says that for separable Y, d_D induces the topology of weak convergence on $\Delta(Y)$. The role of the factor $\frac{1}{2}$ is only to ensure that for all $\alpha, \beta \in \Delta(Y)$, $d_D(\alpha, \beta) \leq 1$.

A.2. Recursive Domain

Let $X_1 := \mathscr{K}\big(\mathscr{F}(\Delta(C))\big)$. For acts $f^1, g^1 \in \mathscr{F}(\Delta(C))$, define the metric $d^{(1)}$ on $\mathscr{F}(\Delta(C))$ by $d^{(1)}(f^1, g^1) := \max_s d_D(f^1(s), g^1(s)) \leq 1$. For any $f^1 \in \mathscr{F}(\Delta(C))$ and $x_1 \in X_1$, the distance of f^1 from x_1 is $d^{(1)}(f^1, x_1) := \min_{g^1 \in x_1} d^{(1)}(f^1, g^1)$ (where the minimum is achieved because x_1 is compact).

Define the Hausdorff metric $d_H^{(1)}$ on X_1 as

$$d_H^{(1)}(x_1, y_1) := \max\left\{\max_{f^1 \in x_1} d^{(1)}(f^1, y_1), \max_{g^1 \in y_1} d^{(1)}(g^1, x_1)\right\} \leq 1$$

Intuitively, X_1 consists of all one-period Anscombe-Aumann (AA) choice problems.

Now define recursively, for $n > 1$, $X_n := \mathscr{K}\big(\mathscr{F}(\Delta(C \times X_{n-1}))\big)$. The metric on $C \times X_{n-1}$ is the product metric, that is, $d_{C\times X_{n-1}}\big((c, x_{n-1}), (c', x'_{n-1})\big) = \max[d_C(c, c'), d^{(n-1)}(x_{n-1}, x'_{n-1})]$. This induces the Dudley metric on $\Delta(C \times X_{n-1})$.

We then define the distance between any two acts $f^n, g^n \in \mathscr{F}(\Delta(C \times X_{n-1}))$ as $d^{(n)}(f^n, g^n) := \max_s d_D(f^n(s), g^n(s))$, and the Hausdorff metric $d_H^{(n)}$ on X_n as

$$d_H^{(n)}(x_n, y_n) := \max\left\{\max_{f^n \in x_n} d^{(n)}(f^n, y_n), \max_{g^n \in y_n} d^{(n)}(g^n, x_n)\right\}$$

Here, X_n consists of all n-period AA choice problems. The agent faces a menu of acts which pay off in lotteries over consumption and $(n-1)$-period AA choice problems that begin the next period.

Finally, endow $\bigtimes_{n=1}^{\infty} X_n$ with the product topology. The Tychonoff metric induces this topology and is given as follows: for $x = (x_1, x_2, \dots), y = (y_1, y_2, \dots) \in \bigtimes_{n=1}^{\infty} X_n$,

$$d(x, y) := \sum_n \frac{d_H^{(n)}(x_n, y_n)}{2^n}$$

which is also bounded above by 1. Moreover, because it holds for $d_H^{(n)}$ for each n, $d\big(\frac{1}{2}x + \frac{1}{2}y, y\big) = \frac{1}{2}d(x, y)$.

The space of choice problems (menus) we consider X is all members of $\bigtimes_{n=1}^{\infty} X_n$ that are *consistent*. Intuitively, $x = (x_1, x_2, \dots)$ is consistent if deleting the last period in the n-period problem x_n results in the $(n-1)$-period problem x_{n-1}.[23] It follows

(23) See Gul and Pesendorfer (2004) for a more formal definition in a related setting.

from standard arguments that X is (linearly) homeomorphic to $\mathscr{K}\big(\mathscr{F}(\Delta(C\times X))\big)$, denoted $X \simeq \mathscr{K}\big(\mathscr{F}(\Delta(C\times X))\big)$. In what follows, we shall abuse notation and use d as a metric both on X and on $\mathscr{K}\big(\mathscr{F}(\Delta(C\times X))\big)$. (It will be clear from the context precisely which space we are interested in.)

There is a natural notion of inclusion in the space of menus: For $x, y \in X$, $y \subset x$ if $y_n \subset x_n$ for all $n \geq 1$.

A.3. Canonical Information Choice Processes

In terms of behavior, all that is relevant for the description of the ICP $\mathscr{M}$ is the set of partitions that are available at each moment in time. We therefore identify ICPs that permit the same choice of partition after every history as indistinguishable (see below), which allows us to metrize the set of all ICPs $\mathbf{M}$ (as in Section 3.2) and prove that finite horizon ICPs approximate arbitrary ICPs. Because all payoffs are bounded, the construction described above then establishes Theorem 1 via a simple continuity argument. We now describe the (pseudo-) metrization of $\mathbf{M}$.

Two ICPs $\mathscr{M}$ and $\mathscr{M}'$ are *indistinguishable* if they afford the same choices of partition in the first period and, for any choice in the first period, the same state-contingent choices in the second period, and so on. Intuitively, indistinguishable ICPs differ only up to a relabeling of the control states, and up to the addition of control states that can never be reached.

Let $\mathscr{M} = (\Theta, \Gamma, \tau, \theta_0)$ and $\mathscr{M}' = (\Theta', \Gamma', \tau', \theta_0')$ be two ICPs in $\mathbf{M}$. A choice of $P \in \Gamma(\theta_0)$ and a realization of state s results in a new ICP $(\Theta, \Gamma, \tau, \tau(\theta_0, P, s))$. To simplify notation, we denote this new ICP by $\mathscr{M}(\tau(\theta_0, P, s))$. Further abusing notation, $\mathscr{M}(\theta)$ denotes the ICP $(\Theta, \Gamma, \tau, \theta)$ with initial state θ. Define $\mathsf{D} : \mathbf{M} \times \mathbf{M} \to \mathbb{R}$ as:

$$\begin{aligned} &\mathsf{D}\left(\mathscr{M}(\theta_0), \mathscr{M}'(\theta_0')\right) := \\ &\qquad \max\left\{ d_H\left(\Gamma(\theta_0), \Gamma'(\theta_0')\right), \tfrac{1}{2} \max_{P \in \Gamma(\theta_0), s \in S} \mathsf{D}\left(\mathscr{M}(\tau(\theta_0, P, s)), \mathscr{M}'(\tau'(\theta_0', P, s))\right) \right\} \end{aligned} \tag{A.1}$$

We endow Θ with the discrete metric, which means that the Hausdorff distance $d_H(A, B) \leq 1$ for all $A, B \subset \Theta$. The function D captures the discrepancy between $\mathscr{M}$ and $\mathscr{M}'$. In what follows, let $\mathbf{B}(\mathbf{M} \times \mathbf{M})$ denote the space of real-valued bounded functions defined on $\mathbf{M} \times \mathbf{M}$ with the supremum norm.

Lemma A.1. There is a unique function $\mathsf{D} \in \mathbf{B}(\mathbf{M} \times \mathbf{M})$ that satisfies equation [A.1].

Proof. Consider the operator $T : \mathbf{B}(\mathbf{M} \times \mathbf{M}) \to \mathbf{B}(\mathbf{M} \times \mathbf{M})$ defined as

$$T\mathsf{D}'\left(\mathscr{M}(\theta_0), \mathscr{M}'(\theta_0')\right) :=$$
$$\max\left\{ d_H\left(\Gamma(\theta_0), \Gamma'(\theta_0')\right), \tfrac{1}{2} \max_{P\in\Gamma(\theta_0), s\in S} \mathsf{D}'\left(\mathscr{M}(\tau(\theta_0, P, s)), \mathscr{M}'(\tau'(\theta_0', P, s))\right) \right\}$$

for all $\mathsf{D}' \in \mathbf{B}(\mathbf{M}\times\mathbf{M})$. Observe that T is monotone in the sense that $\mathsf{D}_1 \le \mathsf{D}_2$ implies $T\mathsf{D}_1 \le T\mathsf{D}_2$. It also satisfies discounting, ie, $T(\mathsf{D} + a) \le T\mathsf{D} + \frac{1}{2}a$ for all $a \ge 0$. This implies that T has a unique fixed point in $\mathbf{B}(\mathbf{M} \times \mathbf{M})$, which satisfies [A.1]. □

We can now define an isomorphism between ICPs. In terms of the discrepancy function D, two ICPs $\mathscr{M}$ and $\mathscr{M}'$ are *indistinguishable* if $\mathsf{D}\left(\mathscr{M}(\theta_0), \mathscr{M}'(\theta_0')\right) = 0$. The definition of D immediately implies the following recursive characterization of indistinguishability whose proof is omitted.

Lemma A.2. Let $\mathscr{M}, \mathscr{M}' \in \mathbf{M}$. Then, $\mathscr{M}$ is indistinguishable from $\mathscr{M}'$ if, and only if, (i) $\Gamma(\theta_0) = \Gamma'(\theta_0')$, and (ii) for all $P \in \Gamma(\theta_0) \cap \Gamma'(\theta_0')$ and $s \in S$, the ICP $\mathscr{M}(\tau(\theta_0,), s))$ is indistinguishable from the ICP $\mathscr{M}'(\tau'(\theta_0', P, s))$.

We now construct a set of *canonical* ICPs.

Recall that $\mathscr{P}$ is the space of all partitions of S, where a typical partition is P. Then, $(\mathscr{P}, \mathsf{d})$ is a metric space, where d is the discrete metric.

For metric spaces X and Y, we denote by $\mathscr{K}_\flat(X \times Y)$ the space of all non-empty closed subsets of $X \times Y$ with the property that a subset contains distinct (x, y) and (x', y') only if $x \ne x'$.

Let $\Omega_1 := \mathscr{K}(\mathscr{P})$, and define recursively for $n > 1$, $\Omega_n := \mathscr{K}_\flat\left(\mathscr{P} \times \Omega_{n-1}^S\right)$. Set $\Omega' := \bigtimes_{n=1}^\infty \Omega_n$. A typical member of Ω_n is ω_n, while $\boldsymbol{\omega}_n = (\omega_{n,s})_{s\in S}$ denotes a typical member of Ω_n^S.

Let $\psi_1 : \mathscr{P} \times \Omega_1^S \to \mathscr{P}$ be given by $\psi_1(P, \boldsymbol{\omega}_1) = P$, and define $\Psi_1 : \Omega_2 \to \Omega_1$ as $\Psi_1(\omega_2) := \{\psi_1(P, \boldsymbol{\omega}_1) : (P, \boldsymbol{\omega}_1) \in \omega_2\}$. Now define recursively, for $n > 1$, $\psi_n : \mathscr{P} \times \Omega_n^S \to \mathscr{P} \times \Omega_{n-1}^S$ as $\psi_n(P, \boldsymbol{\omega}_n) := \left(P, (\Psi_{n-1}(\omega_{n,s}))_s\right)$, and the function (because Ω_n is a space of sets) $\Psi_n : \Omega_{n+1} \to \Omega_n$ by $\Psi_n(\omega_{n+1}) := \{\psi_n(P, \boldsymbol{\omega}_n) : (P, \boldsymbol{\omega}_n) \in \omega_{n+1}\}$.

An $\omega \in \Omega'$ is consistent if $\omega_{n-1} = \Psi_{n-1}(\omega_n)$ for all $n > 1$. The set of canonical ICPs Ω is the set of all consistent elements of Ω',

$$\Omega := \left\{ \omega \in \Omega' : \omega \text{ is consistent} \right\}$$

Notice that Ω_1 is a compact metric space when endowed with the Hausdorff metric. Then, inductively, $\mathscr{P} \times \Omega_{n-1}^S$ with the product metric is a compact metric space, so that endowing Ω_n with the Hausdorff metric in turn makes it a compact metric space. Thus, Ω endowed with the product metric is a compact metric space. (Moreover, Ω is isomorphic to the Cantor set, ie, it is separable and completely disconnected.)

It follows that for $\omega, \omega' \in \Omega$, where $\omega := (\omega_n)_{n=1}^{\infty}$ and $\omega' := (\omega'_n)_{n=1}^{\infty}$, $\omega \neq \omega'$ if, and only if, there is a smallest $N \geq 1$ such that for all $n < N$, $\omega_n = \omega'_n$ but $\omega_N \neq \omega'_N$.

Theorem 4. *The set* Ω *is homeomorphic to* $\mathscr{K}_\flat\big(\mathscr{P} \times \Omega^S\big)$.

We write $\Omega \simeq \mathscr{K}_\flat\big(\mathscr{P} \times \Omega^S\big)$. The theorem is not proved, but this can be done by adapting the arguments in Mariotti, Meier, and Piccione (2005).

The homeomorphism $\Omega \simeq \mathscr{K}_\flat\big(\mathscr{P} \times \Omega^S\big)$ suggests a recursive way to think of Ω: Each $\omega \in \Omega$ describes the set of feasible partitions available for choice in the first period, and how a choice of partition P and the realized state s determine a new $\omega'_s \in \Omega$ in the next period. That is, ω can be identified with a finite collection of pairs $(P, \boldsymbol{\omega}')$, where $\boldsymbol{\omega}' = (\omega'_s)_{s\in S}$. To see that every $\omega \in \Omega$ is indeed an ICP, set $\Gamma^*(\omega) = \{P : (P, \boldsymbol{\omega}') \in \omega\}$ and $\tau^*(\omega, P, s) = \omega'_s$ to obtain the ICP $\mathscr{M}_\omega = (\Omega, \Gamma^*, \tau^*, \omega)$ which is indistinguishable from ω. Also note that for $\omega, \omega' \in \Omega$, $\omega \neq \omega'$ implies $\mathsf{D}(\omega, \omega') > 0$.

Proposition A.3. The space $\mathbf{M}$ of ICPs is isomorphic to Ω in the following sense.

(a) Every $\mathscr{M} \in \mathbf{M}$ is indistinguishable from a unique $\omega_{\mathscr{M}} \in \Omega$.

(b) Every $\omega \in \Omega$ induces an $\mathscr{M}_\omega \in \mathbf{M}$ that is indistinguishable from ω.

Proof. To show (a), let $\mathscr{M} = (\Theta, \Gamma, \tau, \theta_0)$ be an ICP. Recall the definition of the space Ω_n and define the maps $\Phi_n : \Theta \to \Omega_n$ as follows. Let

- $\Phi_1(\theta) := \Gamma(\theta)$,
- $\Phi_2(\theta) := \big\{\big(P, (\Phi_1(\tau(P, \theta, s)))_{s\in S}\big) : P \in \Gamma(\theta)\big\}$,
- $\vdots$
- $\Phi_{n+1}(\theta) := \big\{\big(P, (\Phi_n(\tau(P, \theta, s)))_{s\in S}\big) : P \in \Gamma(\theta)\big\}$,

It is easy to see that for each $\theta \in \Theta$, $\Phi_n(\theta) \in \Omega_n$, ie, Φ_n is well defined.

Now, given θ_0, set $\Phi_n(\theta_0) =: \omega_n \in \Omega_n$. It can be verified that the sequence $(\omega_1, \omega_2, \ldots, \omega_n, \ldots) \in \bigtimes_{n\in\mathbb{N}} \Omega_n$ is consistent in the sense described above. Therefore, there exists $\omega \in \Omega$ such that $\omega = (\omega_1, \omega_2, \ldots, \omega_n, \ldots)$, ie, the ICP $\mathscr{M}$ corresponds to a canonical ICP ω. Observe that if $\omega' \in \Omega$ is indistinguishable form $\mathscr{M}$, then it must also be indistinguishable from $\omega_{\mathscr{M}}$ (because D is a metric), which proves that $\omega_{\mathscr{M}} = \omega$.

To show (b), let $\omega \in \Omega$. A partition P is supported by ω if there exists $\boldsymbol{\omega}' \in \Omega^S$ such that $(P, \boldsymbol{\omega}') \in \omega$. Now set $\Theta = \Omega$, $\theta_0 = \omega$, $\Gamma^*(\theta) = \{P : P \text{ is supported by } \theta\}$, and $\tau^*(P, \omega, s) = \omega'_s$ where $\boldsymbol{\omega}' \in \Omega^s$ is the unique collection of canonical ICPs such that $(P, \boldsymbol{\omega}') \in \omega$. This results in the ICP $\mathscr{M}_\omega = (\Theta, \Gamma^*, \tau^*, \theta_0 = \omega)$ that is uniquely determined by ω. $\square$

B. Value Function and Dynamic Blackwell Order: Proofs from Section 3

Appendix B.1 proves the existence of the value function V that satisfies [**Val**]. Appendix B.2 formally defines the dynamic Blackwell order and describes some of its properties.

B.1. Value Function

We now prove Proposition 3.2 for the case of canonical ICPs. The extension to the case of general ICPs is immediate. In what follows, let $\mathbf{C}(X \times \Omega \times (S \cup \{0\}))$ be the space of continuous functions over $X \times \Omega \times (S \cup \{0\})$ endowed with the supremum norm.

Proof of Proposition 3.2. Define the operator $T : \mathbf{C}(X \times \Omega \times (S \cup \{0\})) \to \mathbf{C}(X \times \Omega \times (S \cup \{0\}))$ as follows:

$$TW(x, \omega, s') = \max_{(P, \boldsymbol{\omega}') \in \omega} \sum_{I \in P} \left[\max_{f \in x} \sum_{s \in S} \mathsf{E}^{f(s)} \Big[u_s(c) + \delta W(y, \omega'_s, s) \Big] \pi_{s'}(s \mid I) \right] \pi_{s'}(I)$$

Because ω is a finite union of $(P, \boldsymbol{\omega}')$, the outer maximum of TW is achieved. The compactness of x implies that a maximizer $f \in x$ exists. It follows from the Theorem of the Maximum (using standard arguments) that T is well defined, ie, $TW \in \mathbf{C}(X \times \Omega \times (S \cup \{0\}))$. It is also easy to see that T is monotone (ie, $W \leq W'$ implies $TW \leq TW'$) and satisfies discounting (ie, $T(W + a) \leq TW + \delta a$), so T is a contraction mapping with modulus $\delta \in (0, 1)$. Using Proposition A.3, it follows that for each tuple $\big((u_s)_{s \in S}, \Pi, \delta, \omega\big)$, there exists a unique $V \in \mathbf{C}(X \times \Omega \times (S \cup \{0\}))$ that satisfies the functional equation [**Val**].

The compactness of C and the discounting of payoffs ensure that the plan that achieves the maxima in TW is actually optimal, that is, the optimal plan is Markovian in (x, ω, s') — see Orkin (1974) or Proposition A6.8 of Kreps (2012). $\square$

B.2. Dynamic Blackwell Order

In this section, we construct the dynamic Blackwell order for canonical ICPs. Appendix A.3 exhibits an isomorphism between canonical ICPs and ICPs. The isomorphism now induces the dynamic Blackwell order on ICPs.

Let $\hat{\omega} \in \Omega$ denote the canonical ICP that delivers the coarsest partition in each period in every state. Define $\hat{\Omega}_0 := \mathscr{K}_\flat\big(\mathscr{P} \times \{\hat{\omega}\}\big)$, and inductively define $\hat{\Omega}_{n+1} := \mathscr{K}_\flat\big(\mathscr{P} \times \hat{\Omega}_n\big)$ for all $n \geq 0$. Notice that for all $n \geq 0$, $\hat{\Omega}_n \subset \hat{\Omega}_{n+1}$. We now define an order $\gtrsim_0$ on $\hat{\Omega}_0$ as follows: $\omega_0 \gtrsim_0 \omega_0'$ if for all $(P', \hat{\boldsymbol{\omega}}) \in \omega_0'$, there exists $(P, \hat{\boldsymbol{\omega}}) \in \omega_0$ such that P is finer than P'. This allows us to define inductively, for all $n \geq 1$, $\gtrsim_n$ on $\hat{\Omega}_n$: For all $\omega_n, \omega_n' \in \hat{\Omega}_n$, $\omega_n \gtrsim_n \omega_n'$ if for all $(P', \boldsymbol{\omega}_{n-1}') \in \omega_n'$, there exists $(P, \boldsymbol{\omega}_{n-1}) \in \omega_n$ such that (i) P is finer than P', and (ii) $\omega_{n-1,s} \gtrsim_{n-1} \omega_{n-1,s}$ for all $s \in S$.

It is easy to see that $\gtrsim_n$ is reflexive and transitive for all n. There is a natural sense in which $\gtrsim_{n+1}$ extends $\gtrsim_n$, as we show next.

Lemma B.1. For all $n \geq 0$, $\gtrsim_{n+1}$ extends $\gtrsim_n$, ie, $\gtrsim_{n+1} |_{\hat{\Omega}_n} = \gtrsim_n$.

Proof. As observed above, $\hat{\Omega}_n \subset \hat{\Omega}_{n+1}$ for all n. First consider the case of $n = 0$ and recall that by construction $\hat{\omega} \in \hat{\Omega}_0$. Let $\omega_0 \gtrsim_0 \omega_0'$. Then, for $(P', \hat{\boldsymbol{\omega}}) \in \omega_0'$, there exists $(P, \hat{\boldsymbol{\omega}}) \in \omega_0$ such that P is finer than P'. Moreover, because $\gtrsim_0$ is reflexive, $\hat{\boldsymbol{\omega}} \gtrsim_0 \hat{\boldsymbol{\omega}}$. But this implies $\omega_0 \gtrsim_1 \omega_0'$. Conversely, let $\omega_0 \gtrsim_1 \omega_0'$. Then, for all $(P', \hat{\boldsymbol{\omega}}) \in \omega_0'$, there exists $(P, \hat{\boldsymbol{\omega}}) \in \omega_0$ such that (i) P is finer than P', and (ii) $\hat{\omega} \gtrsim_0 \hat{\omega}$ for all $s \in S$. But this implies $\omega_0 \gtrsim_0 \omega_0'$, which proves that $\gtrsim_{n+1} |_{\hat{\Omega}_n} = \gtrsim_n$ when $n = 0$.

As our inductive hypothesis, we suppose that $\gtrsim_n |_{\hat{\Omega}_{n-1}} = \gtrsim_{n-1}$. Let $\omega_n \gtrsim_n \omega_n'$. Then, for all $(P', \tilde{\boldsymbol{\omega}}_{n-1}') \in \omega_n'$, there exists $(P, \tilde{\boldsymbol{\omega}}_{n-1}) \in \omega_n$ such that (i) P is finer than P', and (ii) $\tilde{\omega}_{n-1,s} \gtrsim_{n-1} \tilde{\omega}_{n-1,s}'$ for all $s \in S$. But by the induction hypothesis, this is equivalent to $\tilde{\omega}_{n-1,s} \gtrsim_n \tilde{\omega}_{n-1,s}'$ for all $s \in S$, which implies that $\omega_n \gtrsim_{n+1} \omega_n'$.

Conversely, let $\omega_n \gtrsim_{n+1} \omega_n'$. Then, for all $(P', \tilde{\boldsymbol{\omega}}_{n-1}') \in \omega_n'$, there exists $(P, \tilde{\boldsymbol{\omega}}_{n-1}) \in \omega_n$ such that (i) P is finer than P', and (ii) $\tilde{\omega}_{n-1,s} \gtrsim_n \tilde{\omega}_{n-1,s}'$ for all $s \in S$. However, the induction hypothesis implies $\tilde{\omega}_{n-1,s} \gtrsim_{n-1} \tilde{\omega}_{n-1,s}'$ for all $s \in S$, proving that $\omega_n \gtrsim_n \omega_n'$ and therefore $\gtrsim_{n+1} |_{\hat{\Omega}_n} = \gtrsim_n$. $\square$

Let $\hat{\Omega} := \bigcup_{n \geq 0} \hat{\Omega}_n$. Let $\gtrsim$ be a partial order defined on $\hat{\Omega}$ as follows: $\omega \gtrsim \omega'$ if there is $n \geq 1$ such that $\omega, \omega' \in \hat{\Omega}_n$ and $\omega \gtrsim_n \omega'$.

By definition of $\hat{\Omega}$ there is some n such that $\omega, \omega' \in \hat{\Omega}_n$, and by Lemma B.1 the precise choice of this n is irrelevant. Hence $\gtrsim$ is well defined. We now show that $\gtrsim$ also has a recursive definition.

Proposition B.2. For any $\omega, \omega' \in \hat{\Omega}$, the following are equivalent.

(a) $\omega \gtrsim \omega'$.

(b) for all $(P', \tilde{\boldsymbol{\omega}}') \in \omega'$, there exists $(P, \tilde{\boldsymbol{\omega}}) \in \omega$ such that (i) P is finer than P', and (ii) $\tilde{\omega}_s \gtrsim \tilde{\omega}_s'$ for all $s \in S$.

Therefore, $\gtrsim$ is the unique partial order on $\hat{\Omega}$ defined as $\omega \gtrsim \omega'$ if (b) holds.

Proof. (a) implies (b). Suppose $\omega \gtrsim \omega'$. Then, by definition, there exists n such that $\omega, \omega' \in \hat{\Omega}_n$ and $\omega \gtrsim_n \omega'$. This implies that for all $(P', \tilde{\boldsymbol{\omega}}'_{n-1}) \in \omega'_n$, there exists $(P, \tilde{\boldsymbol{\omega}}_{n-1}) \in \omega_n$ such that (i) P is finer than P', and (ii) $\tilde{\omega}_{n-1,s} \gtrsim_{n-1} \tilde{\omega}'_{n-1,s}$ for all $s \in S$. But the latter property implies $\tilde{\omega}_s \gtrsim \tilde{\omega}'_s$ for all $s \in S$, which establishes (b). The proof that (b) implies (a) is similar and is therefore omitted.

The uniqueness of $\gtrsim$ on $\hat{\Omega}$ follows from the uniqueness of $\gtrsim_n$ for all $n \geq 0$. $\square$

We can now prove the existence of a recursive order on Ω. In particular, for all $\omega, \omega' \in \Omega$, we say that ω *dynamically Blackwell dominates* ω' if for all $(P', \tilde{\boldsymbol{\omega}}') \in \omega'$, there exists $(P, \tilde{\boldsymbol{\omega}}) \in \omega$ such that (i) P is finer than P', and (ii) $\tilde{\omega}_s$ dynamically Blackwell dominates $\tilde{\omega}'_s$ for all $s \in S$. The following proposition characterizes the dynamic Blackwell order.

Proposition B.3. The order $\gtrsim$ on $\hat{\Omega}$ has a unique continuous extension to Ω. Moreover, on Ω, $\gtrsim$ is the unique non-trivial and continuous dynamic Blackwell order.

Proof. Notice that $\Omega = \mathrm{cl}(\hat{\Omega})$, where $\mathrm{cl}(\hat{\Omega})$ is the closure of $\hat{\Omega}$, and hence we simply extend $\gtrsim$ to Ω by taking its closure, namely $\mathrm{cl}(\gtrsim)$. Abusing notation, this extension is also denoted by $\gtrsim$. It is easy to see that $\gtrsim$ so defined is continuous and non-trivial. That $\gtrsim$ is a unique dynamic Blackwell order follows from the facts that $\hat{\Omega}$ is dense in Ω, the continuity of $\gtrsim$, and Proposition B.2. $\square$

Let $\mathrm{proj}_n : \Omega \to \hat{\Omega}_n$ be the map associating with each ω the 'truncated and concatenated' version ω_n, which offers the same choices of partition as ω for n stages, but then offers $\hat{\omega}$ (the coarsest partition) forever. Given $\omega \in \Omega$, the sequence (ω_n) is Cauchy, and converges to ω. The next corollary gives us an easy way to establish dominance.

Corollary B.4. For $\omega, \omega' \in \Omega$, $\omega \gtrsim \omega'$ if, and only if, for all $n \in \mathbb{N}$, $\omega_n \gtrsim \omega'_n$.

Proof. The 'only if' part is obvious. The 'if' part follows from the continuity of $\gtrsim$. $\square$

Notice that if $m \geq n$, then $\omega_n = \mathrm{proj}_n\, \omega = \mathrm{proj}_n\, \omega_m$. This observation implies the following corollary.

Corollary B.5. For all $\omega, \omega' \in \Omega$ and $m \geq 1$, $\omega_m \gtrsim \omega'_m$ implies $\omega_n \gtrsim \omega'_n$ for all $1 \leq n \leq m$.

Proof. Notice that $\omega_m, \omega'_m \in \Omega$. Therefore, by Corollary B.4, for all $n \geq 1$, $\text{proj}_n \omega_m \gtrsim \text{proj}_n \omega'_m$. For $n \geq m$, $\text{proj}_n \omega_m = \omega_m$, but for $n \leq m$, $\text{proj}_n \omega_m = \omega_n$, which implies that for all $n \leq m$, $\omega_n \gtrsim \omega'_n$. $\square$

Corollary B.6. Let $\omega^1, \omega^2 \in \Omega$ be such that $\text{proj}_n(\omega^1) \not\gtrsim \text{proj}_n(\omega^2)$ for some $n \geq 1$, but for all $m < n$, $\text{proj}_m(\omega^1) \gtrsim \text{proj}_m(\omega^2)$. Then, there exist finite sequences $(P_k)_1^{n-1}$ and $(s_k)_1^{n-1}$ which induce canonical ICPs $\omega^i_{(n-k)} := \tau^*(\omega^i_{(n-k+1)}, P_k, s_k) \in \hat{\Omega}_{n-k}$ where $P_k \in \Gamma^*(\omega^i_{n-k+1})$, such that $\Gamma^*(\omega^1_1)$ does not setwise Blackwell dominate $\Gamma^*(\omega^2_1)$.[24]

Proof. If not, we would have $\omega^1_n \gtrsim \omega^2_n$, a contradiction. $\square$

Let $\trianglerighteq$ be a dynamic order on $\mathbf{M}$ that satisfies the following recursive criterion: For ICPs $\mathcal{M} = (\Theta, \Gamma, \tau, \theta_0)$ and $\mathcal{M}' = (\Theta', \Gamma', \tau', \theta'_0)$,

[✠] $\mathcal{M} \trianglerighteq \mathcal{M}'$ if for every $P' \in \Gamma'(\theta'_0)$, there exists $P \in \Gamma(\theta_0)$ such that (i) P is finer than P', and (ii) $(\Theta, \Gamma, \tau, \tau(\theta_0, P, s)) \trianglerighteq (\Theta', \Gamma', \tau', \tau'(\theta'_0, P', s))$ for all $s \in S$.

To see that such a dynamic order exists, for any ICP $\mathcal{M}$, denote by $\omega_{\mathcal{M}}$ the canonical ICP that is indistinguishable from it (by part (a) of Proposition A.3, there is a unique such $\omega_{\mathcal{M}}$.) Then the order $\trianglerighteq^*$ on $\mathbf{M}$ defined by $\mathcal{M} \trianglerighteq^* \mathcal{M}'$ if, and only if, $\omega_{\mathcal{M}} \gtrsim \omega_{\mathcal{M}'}$ satisfies condition [✠].

We are now ready to prove Proposition 3.3 that demonstrates that there is a largest order that satisfies [✠]. The proof will establish that this order is $\trianglerighteq^*$. We will then term $\trianglerighteq^*$ the dynamic Blackwell order on $\mathbf{M}$, which is precisely the one we refer to in the main text.

Proof of Proposition 3.3. We will show that for any $\trianglerighteq$ that satisfies condition [✠], if $\mathcal{M} \not\trianglerighteq^* \mathcal{M}'$ then $\mathcal{M} \not\trianglerighteq \mathcal{M}'$. Suppose that $\omega_{\mathcal{M}} \not\gtrsim \omega_{\mathcal{M}'}$. Corollaries B.4 and B.5 imply that there is a smallest n such that $\omega_{\mathcal{M},n} \not\gtrsim \omega_{\mathcal{M}',n}$ but that for all $m < n$, $\omega_{\mathcal{M},m} \gtrsim \omega_{\mathcal{M}',m}$ (where $\omega_{\mathcal{M},n} = \text{proj}_n \omega_{\mathcal{M}}$ as defined in Appendix B.2). By Corollary B.6, there exists a finite sequence of partitions (P_k) and states (s_k) such that $\Gamma^*\big(\tau^{*(n)}(\theta_0, (P_k), (s_k))\big)$ does not setwise Blackwell dominate $\Gamma^*\big(\tau^{*(n)}(\theta'_0, (P_k), (s_k))\big)$, where $\tau^{*(n)}(\theta_0, (P_k), (s_k))$ represents the n-stage transition following the sequence of choices (P_k) and states (s_k). Recall that $\mathcal{M}$ is indistinguishable from $\omega_{\mathcal{M}}$, and so is $\mathcal{M}'$ from $\omega_{\mathcal{M}'}$. Therefore, $\Gamma\big(\tau^{(n)}(\theta_0, (P_k), (s_k))\big)$ does not setwise Blackwell dominate $\Gamma'\big({\tau'}^{(n)}(\theta'_0, (P_k), (s_k))\big)$, which implies that $\mathcal{M} \not\trianglerighteq \mathcal{M}'$. $\square$

(24) A set of partitions $\mathscr{P}$ setwise Blackwell dominates another set of partitions $\mathbb{Q}$ if for any $Q \in \mathbb{Q}$ there is a $P \in \mathscr{P}$ which is finer than Q.

C. Identification and Behavioral Comparison: Proofs from Section 4

Based on the previous results and notation, we now establish Theorems 1 and 2.

In accordance with the discussion in Section 4.1, x is strongly aligned with ω if (i) $V(x,\omega,0) \geq V(x,\omega',0)$ for all $\omega' \in \Omega$, and (ii) ω' does not dynamically Blackwell dominate ω implies $V(x,\omega,0) > V(x,\omega',0)$. We say that P is supported by ω if there exists $\boldsymbol{\omega}' \in \Omega^S$ such that $(P,\boldsymbol{\omega}') \in \omega$.

Recall that ℓ^* and ℓ_* are the best and worst consumption streams, respectively, while c_s^+ and c_s^- denote the best and worst instantaneous consumption in state s. It follows immediately from the representation that ℓ^* (respectively, ℓ_*) consists of the prize (ie, the degenerate lottery) c_s^+ (respectively, c_s^-) in state s in any period.

Lemma C.1. Let $(P,\boldsymbol{\omega}') \in \omega$. There exists a menu $x(P,\boldsymbol{\omega}')$ recursively defined as $x(P,\boldsymbol{\omega}') = \{f_J : J \in P\}$ with

[★]
$$f_J(s) := \begin{cases} \left(c_s^+, \text{Unif}\left(\{x(Q,\tilde{\boldsymbol{\omega}}) : (Q,\tilde{\boldsymbol{\omega}}) \in \omega_s'\}\right)\right) & \text{if } s \in J \\ \ell_*(s) & \text{if } s \notin J \end{cases}$$

where $\text{Unif}(\cdot)$ is the uniform lottery over a finite set.

Proof. For a partition P with generic cell J, define the act

$$f_{1,J}(s) := \begin{cases} \ell^*(s) & \text{if } s \in J \\ \ell_*(s) & \text{if } s \notin J \end{cases}$$

and for each P that is supported by ω, define $x_1(P) := \{f_{1,J} : J \in P\}$.

Now, proceed inductively, and for $n \geq 2$, suppose we have the menu $x_{n-1}(P,\boldsymbol{\omega}')$ for each $(P,\boldsymbol{\omega}') \in \omega$, and define, for each cell $J \in P$, the act

$$f_{n,J}(s) := \begin{cases} (c_s^+, \text{Unif}_{n-1}\left(\{x_{n-1}(Q,\tilde{\boldsymbol{\omega}}) : (Q,\tilde{\boldsymbol{\omega}}) \in \omega_s'\}\right)) & \text{if } s \in J \\ \ell_*(s) & \text{if } s \notin J \end{cases}$$

Then, given $(P,\boldsymbol{\omega}') \in \omega$, we have the menu $x_n(P,\boldsymbol{\omega}') := \{f_{n,J} : J \in P\}$.

It is easy to see that for a fixed $(P,\boldsymbol{\omega}') \in \omega$, the sequence of menus $\big(x_n(P,\boldsymbol{\omega}')\big)$ is a Cauchy sequence. Because X is complete, this sequence must converge to some $x(P,\boldsymbol{\omega}') \in X$. Moreover, this means that the sequence of sets $(\{x_n(Q,\tilde{\boldsymbol{\omega}}) : (Q,\tilde{\boldsymbol{\omega}}) \in \omega_s'\})$ also converges to $(\{x(Q,\tilde{\boldsymbol{\omega}}) : (Q,\tilde{\boldsymbol{\omega}}) \in \omega_s'\})$. This allows us to denote the uniform lottery over this finite set of points in X by $\text{Unif}(\omega_s')$.

Thus, $x(P, \boldsymbol{\omega}')$ consists of the acts $\{f_J : J \in P\}$ where for each $J \in P$

$$f_J(s) := \begin{cases} \left(c_s^+, \text{Unif}\left(\{x(Q, \tilde{\boldsymbol{\omega}}) : (Q, \tilde{\boldsymbol{\omega}}) \in \omega_s'\}\right)\right) & \text{if } s \in J \\ \ell_*(s) & \text{if } s \notin J \end{cases}$$

as claimed. □

It is straightforward to verify that

$$V(\ell^*, \omega, 0) = V\big(x(P, \boldsymbol{\omega}'), \omega, 0\big) \geq V\big(x(P, \boldsymbol{\omega}'), \tilde{\omega}, 0\big)$$

for all $\tilde{\omega} \in \Omega$. Indeed, $V\big(x(P, \boldsymbol{\omega}'), \omega, 0\big) = V\big(x(P, \boldsymbol{\omega}'), (P, \boldsymbol{\omega}'), 0\big)$.

Lemma C.2. Let $P, Q \in \mathscr{P}$ and suppose Q is not finer than P. Then, for any $\boldsymbol{\omega}, \boldsymbol{\omega}' \in \Omega^S$, $V\big(x(P, \boldsymbol{\omega}), (P, \boldsymbol{\omega}), 0\big) > V\big(x(P, \boldsymbol{\omega}), (Q, \boldsymbol{\omega}'), 0\big)$.

Proof. Fix $(P, \boldsymbol{\omega}) \in \Omega$ and consider the menu $x(P, \boldsymbol{\omega})$ defined in [★]. As noted above, for all $\boldsymbol{\omega}'$, we have $V\big(x(P, \boldsymbol{\omega}), (P, \boldsymbol{\omega}), 0\big) = V\big(x(P, \boldsymbol{\omega}'), (P, \boldsymbol{\omega}'), 0\big)$. Moreover, it must be that for all $(Q, \boldsymbol{\omega}')$ (even for $Q = P$), we have $V\big(x(P, \boldsymbol{\omega}), (P, \boldsymbol{\omega}), 0\big) \geq V\big(x(P, \boldsymbol{\omega}'), (Q, \boldsymbol{\omega}'), 0\big)$ and in the case where Q is not finer than P and $Q \neq P$, $V\big(x(P, \boldsymbol{\omega}'), (P, \boldsymbol{\omega}'), 0\big) > V\big(x(P, \boldsymbol{\omega}'), (Q, \boldsymbol{\omega}'), 0\big)$ by construction of the menu $x(P, \boldsymbol{\omega}')$. (This is a version of Blackwell's theorem on comparison of experiments; see Theorem 1 on p59 of Laffont (1989).) □

Lemma C.3. Suppose ω^2 does not dynamically Blackwell dominate ω^1. Then, for some $(P, \tilde{\boldsymbol{\omega}}) \in \omega^1$, the menu $x(P, \tilde{\boldsymbol{\omega}})$ defined in [★] is such that $V\big(x(P, \tilde{\boldsymbol{\omega}}), \omega^1, 0\big) = V\big(x(P, \tilde{\boldsymbol{\omega}}), (P, \tilde{\boldsymbol{\omega}}), 0\big) > V\big(x(P, \tilde{\boldsymbol{\omega}}), \omega^2, 0\big)$.

Proof. Suppose ω^2 does not dynamically Blackwell dominate ω^1. Then, there exists a smallest $n \geq 1$ such that for all $m < n$, $\text{proj}_m(\omega^2)$ dynamically Blackwell dominates $\text{proj}_m(\omega^1)$, while $\text{proj}_n(\omega^2)$ does not dynamically Blackwell dominate $\text{proj}_n(\omega^1)$.

From Corollary B.6 it follows that there exist finite sequences (P_k) and (P_k') of partitions, and (s_k) of states, such that $\Gamma^*\big(\tau^{*(n)}(\omega^2, (P_k'), (s_k))\big)$ does not setwise Blackwell dominate the set $\Gamma^*\big(\tau^{*(n)}(\omega^1, (P_k), (s_k))\big)$, where (i) $\tau^{*(n)}(\theta_0, (P_k), (s_k))$ represents the n-stage transition following the sequence of choices (P_k) and states (s_k), (ii) $\omega_{n-k}^i = \tau^*(\omega_{n-k+1}^i, P_k, s_k)$ where $P_k \in \Gamma^*(\omega_{n-k+1}^i)$, and (iii) $\Gamma^*(\omega_1^1)$ does not setwise Blackwell dominate $\Gamma^*(\omega_1^2)$.

Let $(P_1, \tilde{\boldsymbol{\omega}}) \in \omega^1$ be the unique first period choice under ω^1 that makes the sequence (P_k) feasible. Then $x(P_1, \tilde{\boldsymbol{\omega}})$ defined in [★] is aligned with $(P_1, \tilde{\boldsymbol{\omega}})$. That is, after n stages of choice and a certain path of states we can appeal to Lemma C.2, which completes the proof. □

Proof of Theorem 1. For the case of finitely many prizes (ie, when C is a finite set), Corollary 5 of Krishna and Sadowski (2014) establishes that the collection $\big((u_s), \Pi, \delta\big)$ is unique in the sense of the Theorem. While we cannot directly appeal to their result, judicious and repeated applications of their corollary allows us to reach the same conclusion for a compact set of prizes (see the proof of Proposition D.1 below). Now, define $F_\omega := \{x(P, \tilde{\boldsymbol{\omega}}) : (P, \tilde{\boldsymbol{\omega}}) \in \omega\}$. It follows immediately from Lemma C.3 that F_ω is uniformly strongly aligned with ω. $\square$

This allows us to characterize the dynamic Blackwell order in terms of the instrumental value of information.

Corollary C.4. Let $\omega, \omega' \in \Omega$. Then, the following are equivalent.

(a) ω dynamically Blackwell dominates ω'.

(b) For any $\big((u_s), \Pi, \delta\big)$ that induces $\omega \mapsto V(\cdot, \omega, \cdot)$, $V(x, \omega, \cdot) \geq V(x, \omega', \cdot)$ for all $x \in X$.

Proof. That (a) implies (b) is easy to see. That (b) implies (a) is merely the contrapositive to Lemma C.3. $\square$

We are now in a position to prove Theorem 2.

Proof of Theorem 2. We first show the 'only if' part. On L, we have $\ell \succsim^\dagger \ell'$ implies $\ell \succsim \ell'$. This implies, by Lemma 34 of Krishna and Sadowski (2014), that $\succsim^\dagger|_L = \succsim|_L$. Together with the uniqueness of the Recursive Anscombe-Aumann (RAA) representation (see Appendix D below), this implies that $\big((u_s), \delta, \Pi\big) = \big((u_s^\dagger), \delta^\dagger, \Pi^\dagger\big)$ after a suitable (and behaviorally irrelevant) normalization of the state-dependent utilities. Thus, part (b) of Corollary C.4 holds, which establishes the claim.

The 'if' part follows immediately from Corollary C.4. $\square$

D. Representation on Consumption Streams

Before we proceed, it is useful to record a consequence of assuming Axioms 0 and 1.

Let $u_s \in \mathbf{C}(C)$ for all $s \in S$, $\delta \in (0, 1)$, Π represent the transition operator for a fully connected Markov process on S, and π_0 be the unique invariant distribution of Π. A preference on L has a *Recursive Anscombe-Aumann* (RAA) representation $((u_s)_{s \in S}, \Pi, \delta)$ if $W_0(\cdot) := \sum_s W(\cdot, s)\pi_0(s)$ represents it, where $W(\cdot, s)$ is defined recursively as

$$W(\ell; s) = \sum_{s' \in S} \Pi(s, s')\big[u_{s'}(\ell_1(s')) + \delta W(\ell_2(s'); s')\big]$$

and where u_s is non-trivial for some $s \in S$. Then, W_0 can also be written as

$$W_0(\ell) = \sum_{s \in S} \pi_0(s) \big[u_s(\ell_1(s)) + \delta W(\ell_2(s); s)\big]$$

because π_0 is the unique invariant distribution of Π and satisfies $\pi_0(s) = \sum_s \pi_0(s')\Pi(s', s)$. For some fixed $c_s^\dagger \in C$, we say that the preference on L has a *standard* RAA representation $((u_s)_{s \in S}, \Pi, \delta)$, if $u_s(c_s^\dagger) = 0$ for all $s \in S$.

Proposition D.1. The preference $\succsim|_L$ on L has a standard RAA representation if, and only if, it satisfies Axioms 0 and 1. Moreover, Π and δ are unique and the collection $(u_s)_{s \in S}$ is unique up to a common positive scaling.

Proposition D.1 is formally established in the end of this section.

To see that $L \simeq \mathscr{F}(\Delta(C \times L))$, note that we can define $L^{(1)} := \mathscr{F}(\Delta(C))$ and then recursively define $L^{(n)} := \mathscr{F}(\Delta(C \times L^{(n-1)}))$ as the space of consumption streams of length n. Just as with the definition of the space of choice problems X in Appendix A.2, we say that L is the space of all consistent sequences in $\bigtimes_{n=1}^{\infty} L^{(n)}$.

The support of a consumption stream $\ell \in L$ is a set $B \subset C$ such that at any period and in any state, the realized consumption is in B. A consumption stream has finite support if its support in C is finite. For any finite set $B \subset C$, we can define L_B as the space of all consumption streams with prizes in B. Formally, $L_B \simeq \mathscr{F}(\Delta(C \times L_B))$. Let L_0 be the space of all consumption streams with finite support. That is, $L_0 := \bigcup \{L_B : B \subset C,\ B \text{ finite}\}$.

Analogously to L_0, we can define $L_0^{(n)}$ as the space of consumption streams of length n with finite support. For any $\ell^{(n)} \in L_0^{(n)}$, $\ell^{(n)} \diamond \ell^\dagger \in L_0$, where $\ell^{(n)} \diamond \ell^\dagger$ is the concatenation of $\ell^\dagger$ to $\ell^{(n)}$. In other words, each $L^{(n)}$ is naturally embedded in L_0.

Proposition D.2. The space L_0 is dense in L.

Proof. Because probability measures on C with finite support are dense in $\Delta(C)$, it follows that for all $n \geq 1$, $L_0^{(n)}$ is dense in $L^{(n)}$. (The metrics defined on $L^{(n)}$ make this clear — see Appendix A.2 for a formal definition.) By the definition of the product metric, for any $\ell \in L$ and $\varepsilon > 0$, there exist an n and an $\ell^{(n)} \in L^{(n)}$ such that $d(\ell, \ell^{(n)} \diamond \ell^\dagger) < \varepsilon$, which completes the proof. $\square$

Axiom 1 stipulates that $\succsim|_L$ is non-trivial. We now show that $\succsim_s$ (defined in Section 5.1) is also non-trivial for each $s \in S$. First some notation. For each $I \subset S$, $f \in \mathscr{F}(\Delta(C \times X))$, and $\varepsilon \in [0, 1]$, define

$$(f \oplus_{\varepsilon, I} (c, y))(s) := f(s) + \varepsilon 1_{s \in I}\big((c, y) - f(s)\big)$$

That is, for any state $s \in I$, the act $f \oplus_{\varepsilon,I} (c, y)$ perturbs the continuation lottery with y. Recall that $L \subset \mathscr{F}(\Delta(C \times X))$.

Lemma D.3. Let $\ell^0, \ell^1 \in L$. Then, $\ell^0(s) \sim_s \ell^1(s)$ for all $s \in S$ implies $\ell^0 \sim |_L \ell^1$.

Proof. By definition of $\succsim_s$, $\ell^0(s) \sim_s \ell^1(s)$ if, and only if, $\ell_* \oplus_{(1,s)} \ell^0 \sim |_L \ell_* \oplus_{(1,s)} \ell^1$. Repeatedly applying L-Independence, we find

$$\frac{1}{n}\ell^0 + \frac{n-1}{n}\ell_* = \frac{1}{n}\sum_{s \in S} \ell_* \oplus_{(1,s)} \ell^0 \sim |_L \frac{1}{n}\sum_{s \in S} \ell_* \oplus_{(1,s)} \ell^1 = \frac{1}{n}\ell^1 + \frac{n-1}{n}\ell_*$$

By L-Independence, $\ell^0 \sim |_L \ell^1$. (This follows immediately once we note that, by the Mixture Space Theorem, $\succsim |_L$ has an affine representation.) $\square$

Lemma D.4. There exists $s \in S$ such that $\ell^*(s) \nsim_s \ell_*(s)$. For all $s \in S$, there exists $s' \in S$ such that $(c, \ell_* \oplus_{(1,s)} \ell^*) \nsim_{s'} (c, \ell_*)$.

Proof. Axiom a says that $\ell^* \succ |_L \ell_*$. Therefore, by (the contrapositive to) Lemma D.3, there must exists an s such that $\ell^*(s) \nsim_s \ell_*(s)$. In particular, $\ell_* \oplus_{(1,s)} \ell^* \nsim |_L \ell_*$.

To see the second part, suppose by way of contradiction that for all $s' \in S$, $(c, \ell_* \oplus_{(1,s)} \ell^*) \sim_{s'} (c, \ell_*)$. Set ℓ^0 and ℓ^1 such that $\ell^0(s') = (c, \ell_* \oplus_{(1,s)} \ell^*)$, while $\ell^1(s') = (c, \ell_*)$. It follows from Lemma D.3 that $(c, \ell_* \oplus_{(1,s)} \ell^*) \sim |_L (c, \ell_*)$.

Now, L-Stationarity (Axiom 1(d)) and the fact that $\ell_* \oplus_{(1,s)} \ell^* \nsim |_L \ell_*$ imply that we have $(c, \ell_* \oplus_{(1,s)} \ell^*) \nsim |_L (c, \ell_*)$, which yields the desired contradiction. $\square$

Proposition D.5. For all $s \in S$, $\succsim_s$ is non-trivial.

Proof. Lemma D.4 and (the contrapositive to) L-History Independence (Axiom 1(c)) imply $(c, \ell_* \oplus_{(1,s)} \ell^*) \nsim_{s''} (c, \ell_*)$ for all $s'' \in S$, as claimed. $\square$

Lemma D.6. For any finite set $B \subset C$, the induced preference $\succsim |_{L_B}$ satisfies the Axioms stated in Corollary 5 of Krishna and Sadowski (2014, henceforth KS).

Proof. It follows from Proposition D.5 that each $\succsim_s$ is non-trivial. That is, $\succsim |_L$ is state-wise nontrivial. In addition, $\succsim |_L$ is continuous, satisfies Independence, and is separable in ℓ_1 and ℓ_2, thereby satisfying Axioms 2, 3, and 5 in KS. Axioms 6, 7, and 9 in KS correspond to properties (c), (d), and (b) of our Axiom 1. $\square$

We now proceed to the proof of Proposition D.1.

Proof of Proposition D.1. Let $\succsim_s^C$ denote the induced preference over $\Delta(C)$ in state s. It is clear that $\succsim_s^C$ is well defined, continuous on $\Delta(C)$, satisfies Independence, and that $c^+(s) \succsim_s^C \alpha \succsim_s^C c^-(s)$ for all $\alpha \in \Delta(C)$. Let F_0 be the finite set of consumption defined as

$$F_0 := \{c^-(s), c^\dagger(s), c^+(s) : s \in S\}$$

Let $B \subset C$ be finite. By Lemma D.6, $\succsim |_{L_B}$ satisfies the Axioms in Corollary 5 of KS. This implies there exists a tuple $\big((u_s^B)_{s\in S}, \delta^B, \Pi^B\big)$ that is an RAA representation of $\succsim |_{L_B}$. If $F_0 \subset B$, then we may assume, without loss of generality, that $u_s^B(c^\dagger(s)) = 0$ for all $s \in S$. Then, Corollary 5 in KS says that the collection of utilities (u_s^B) is uniquely identified up to a joint scaling, and that Π^B and δ^B are also uniquely determined.

Now, consider any other finite set D such that $F_0 \subset B \subset D$. By Lemma D.6, $\succsim |_{L_D}$ also has an RAA representation $\big((u_s^D)_{s\in S}, \delta^D, \Pi^D\big)$. As before, if we set $u_s^D(c^\dagger(s)) = 0$ for all $s \in S$, then the collection of utilities (u_s^D) is identified up to a common scaling. Now, because $B \subset D$, we have $L_B \subset L_D$. Therefore, the RAA representation $\big((u_s^D)_{s\in S}, \delta^D, \Pi^D\big)$ of $\succsim |_{L_D}$ when restricted to L_B is also a representation of $\succsim |_{L_B}$, with the feature that $u_s^D(c^\dagger(s)) = 0$ for all $s \in S$. Once again, the uniqueness of the RAA representation implies that a single joint scaling of the collection (u_s^D) results in $u_s^D|_B = u_s^B$ for all $s \in S$, $\Pi^B = \Pi^D$, and $\delta^B = \delta^D$.

Recall that $c^+(s) \succsim_s^C \alpha \succsim_s^C c^-(s)$ for all $\alpha \in \Delta(C)$. Because u_s^B and u_s^D represent, respectively, $\succsim_s^C |_{\Delta(B)}$ and $\succsim_s^C |_{\Delta(D)}$, it must be that $\lambda^*(s) := u_s^B(c^+(s)) = u_s^D(c^+(s))$ and $\lambda_*(s) := u_s^B(c^-(s)) = u_s^D(c^+(s))$. Since B and D are arbitrary, the same must be true for all finite B that contains F_0.

It is easy to extend this argument to any two finite sets that contain F_0, and thus to all finite sets that contain F_0. In particular, for RAA representations on all such finite sets, the Markov transition operator Π has been identified uniquely, as has the discount factor $\delta \in (0, 1)$.

Let $u_s \in \mathbf{C}(C)$ be a vN-M utility representation of $\succsim_s^C$ such that $u_s(c^\dagger(s)) = 0$. Both $u_s|_{\Delta(B)}$ and u_s^B are vN-M representations of $\succsim^C |_{\Delta(B)}$ and by the Mixture Space Theorem, differ at most by a positive affine transformation. Because they agree on $c^\dagger(s)$, they differ at most by a positive scaling. Therefore, if we scale u_s so that $u_s(c^+(s)) = \lambda^*(s)$, we have $u_s(c^-(s)) = \lambda_*(s)$ for all $s \in S$.

Consider the tuple $((u_s), \Pi, \delta)$, and the functional $W_0 : L \to \mathbb{R}$ defined as $W_0(\ell) := \sum_s \pi_0(s) W(\ell, s)$, where

$$W(\ell, s) := \sum_{s'} \Pi(s, s') \big[u_{s'}(\ell_1(s')) + \delta W(\ell_2(s'), s')\big]$$

Note that the function W_0 is uniquely determined by the tuple $\big((u_s)_{s\in S}, \delta, \Pi\big)$. As established above, W_0 represents $\succsim|_{L_B}$ for every finite B, that is, W_0 represents $\succsim|_{L_0}$. Proposition D.2 says that L_0 is dense in L, and because W_0 is (uniformly) continuous, it also represents $\succsim$ on L. The uniqueness of the RAA representation of $\succsim|_L$ (given our normalizations) follows immediately, which concludes the proof. $\square$

E. Existence

By Axiom 0 we have the representation in display [5.1]. We also take as given the existence of the Recursive Anscombe-Aumann Representation for $\succsim|_L$, described in Appendix D.

In the rest of this Appendix, we exhibit the existence of an ICP representation. We first define self-generating representations and the concomitant dynamic plans in Appendix E.1. In Appendix E.1.1 we show that continuation preferences $\succsim_{(x,s)}$ are well defined when x is finite and has a unique optimal partition. In Appendix E.1.2, we show how menus are valued relative to consumption streams based on the recursive application of our axioms. Finally, we establish the existence of self-generating representations (by proving Proposition E.3) in Appendix E.1.3.

Given a self-generating representation, Appendix E.2 shows that it is possible to extract the underlying canonical ICP. In Appendix E.3, we note that $\succsim$ has a unique representation on L. We use this fact, and the canonical ICP extracted from the self-generating representation, to prove Theorem 3.

E.1. Self-Generating Representations and Dynamic Plans

Recall that $\mathbf{C}(X)$ is the space of all real-valued continuous functions on X. Let $\ell^\dagger \in L$ be the consumption stream that delivers $c_s^\dagger$ in state s at every date.

Suppose $\big((u_s), \mathbb{Q}, (v_s(\cdot, P)), \pi\big)$ is a tuple where

- $u_s \in \mathbf{C}(C)$ for all $s \in S$,
- $\mathbb{Q} \subset \mathscr{P}$,
- $v_s(\cdot, P) \in \mathbf{C}(X)$ for all $s \in S$ and $P \in \mathbb{Q}$,
- $\pi \in \Delta(S)$,
- $u_s(c_s^\dagger) = v_s(\ell^\dagger, P) = 0$ for all $s \in S$ and $P \in \mathbb{Q}$,
- $v_s(\cdot, P)$ is independent of P on L, and
- $v_s(\cdot, P)$ is non-trivial on L, and hence on X, for all $s \in S$ and $P \in \mathbb{Q}$,

and $v \in \mathbb{R}^X$ is such that

$$v(x) = \max_{P \in \mathbb{Q}} \sum_{I \in P} \pi(I) \max_{f \in x} \sum_{s \in S} \pi(s \mid I) \left[u_s(f_1(s)) + v_s(f_2(s), P)\right]$$

In that case, we say that the tuple $\big((u_s), \mathbb{Q}, (v_s(\cdot, P)), \pi\big)$ is a *separable and partitional implementation* of v, or in short, an *implementation* of v. (By definition, the implementation takes value 0 on $\ell^\dagger(s)$ for all $s \in S$ and is linear on L.)

More generally, for any subset $\Phi \subset \mathbf{C}(X)$, define the operator $\mathbf{A} : 2^{\mathbf{C}(X)} \to 2^{\mathbf{C}(X)}$ as follows:

$$\mathbf{A}\Phi := \Big\{ v \in \mathbf{C}(X) : \exists \big((u_s), \mathbb{Q}, (v_s(\cdot, P)), \pi\big) \text{ that implements } v \\ \text{and } v_s(\cdot, P) \in \Phi \text{ for all } s \in S \text{ and } P \in \mathbb{Q} \Big\}$$

Proposition E.1. The operator $\mathbf{A}$ is well defined and has a largest fixed point $\Phi^* \neq \{\mathbf{0}\}$. Moreover, Φ^* is a cone.

Proof. For all nonempty $\Phi \subset \mathbf{C}(X)$, $\mathbf{A}\Phi$ is nonempty. (Simply take any $\mathbb{Q}$, any $\mathbf{0} \neq v_s(\cdot, P) \in \Phi$ for each $P \in \mathbb{Q}$, and any u_s for each $s \in S$, so that $\mathbf{A}\Phi \neq \varnothing$.) The operator $\mathbf{A}$ is monotone in the sense that $\Phi \subset \Phi'$ implies $\mathbf{A}\Phi \subset \mathbf{A}\Phi'$. Thus, it is a monotone mapping from the lattice $2^{\mathbf{C}(X)}$ to itself, where $2^{\mathbf{C}(X)}$ is partially ordered by inclusion. The lattice $2^{\mathbf{C}(X)}$ is complete because any collection of subsets of $2^{\mathbf{C}(X)}$ has an obvious least upper bound: the union of this collection of subsets. Similarly, a greatest lower bound is the intersection of this collection of subsets (which may be empty). Therefore, by Tarski's fixed point theorem, $\mathbf{A}$ has a largest fixed point $\Phi^* \in 2^{\mathbf{C}(X)}$.

To see that $\Phi^* \neq \{\mathbf{0}\}$, ie, Φ^* does not contain only the trivial function $\mathbf{0}$, fix $\mathbb{Q} = \big\{ \{\{s\} : s \in S\} \big\}$ so that it contains only the finest partition of S. For the value function V in [Val], take any $u_s \in \mathbf{C}(C) \setminus \{\mathbf{0}\}$ with $u_s(c_s^\dagger) = 0$ for all $s \in S$, a discount factor $\delta \in (0, 1)$, and π as the uniform distribution over S. Then V is implemented by $\big((u_s), \mathbb{Q}, \delta V), \pi\big)$, while δV is implemented by $\big((\delta u_s), \mathbb{Q}, \delta^2 V), \pi\big)$, and so on. Therefore, the set $\Phi_V := \{\delta^n V : n \geq 0\}$ is a fixed point of $\mathbf{A}$. Because $\Phi_V \subset \Phi^*$, it must be that Φ^* is nonempty.

Finally, to see that Φ^* is a cone, let $v \in \Phi^*$ and suppose $\big((u_s), \mathbb{Q}, (v_s(\cdot, P)), \pi\big)$ implements v. Then, for all $\lambda \geq 0$, $\big((\lambda u_s), \mathbb{Q}, (\lambda v_s(\cdot, P)), \pi\big)$ implements λv, ie, $\lambda \Phi^*$ is also a fixed point of $\mathbf{A}$. Because Φ^* is the largest fixed point, it must be a cone. $\square$

Notice that each $v \in \Phi^*$ is implemented by a tuple $\big((u_s), \mathbb{Q}, (v_s(\cdot, P)), \pi\big)$ with the property that each $v_s(\cdot, P) \in \Phi^*$. Adapting the terminology of Abreu, Pearce, and

Stacchetti (1990), the set Φ^* consists of *self-generating* preference functionals that have a separable and partitional implementation. (Notice that unlike Abreu, Pearce, and Stacchetti (1990), our self-generating set lives in an infinite dimensional space.) In what follows, if $\succsim$ is represented by $V \in \Phi^*$, we shall say that V is a *self-generating representation* of $\succsim$.

Lemma E.2. Let $x \in X'_P$. Then, there exists $\varepsilon > 0$ such that $B(x;\varepsilon) \subset X'_P$. In particular, for any $y \in X$ and $s \in S$, there exists $\varepsilon' \in (0,1]$ such that for all $\varepsilon \in (0,\varepsilon')$, $x \oplus_{\varepsilon,s} y \in X'_P$.

Proof. The function V_0 in [**5.1**] is continuous, which in turn implies that each X'_P is open. The result now follows immediately. □

In Appendix E.1.1, we show that $\succsim_{(x,s)}$ is a well-defined binary relation on X, and that $\succsim_{(x,s)}=\succsim_{(P,s)}$, where P is an optimal information choice given x. Conversely, for every P and s there is a finite $x \in X$, such that $\succsim_{(x,s)}=\succsim_{(P,s)}$ on X. With $\succsim_{(x,s)}$ now defined, we can recursively apply Axiom 0 as described in Section 5.2.

Proposition E.3. Let $\succsim$ be a binary relation on X. The following are equivalent.

(a) $\succsim$ recursively satisfies Axioms 0 and satisfies Axiom 1.

(b) $\succsim$ has a self-generating representation, ie, there exists a function $V \in \Phi^*$ that represents $\succsim$ and has a standard RAA representation on L.

As steps to prove this result, in Appendix E.1.1 we show that $\succsim_{(P,s)}$ recursively satisfies Axioms 0 and satisfies Axiom 1. In Appendix E.1.2, we derive some useful properties of consumption streams based on the recursive application of Axiom 0. Finally, we prove Proposition E.3 in Appendix E.1.3.

E.1.1. *Continuation Preferences*

Lemma E.4. Let $x = \{f_1, \dots, f_m\}$, and $x' = \{f'_1, f'_2, \dots, f'_m\}$. Suppose $d(f_i, f'_i) < \varepsilon$. Then, $d(x, x') < \varepsilon$.

Proof. Recall that $d(f_i, x') := \min_{f'_j \in x'} d(f_i, f'_j) < \varepsilon$. Therefore, $\max_{f_i \in x} d(f_i, x') < \varepsilon$. A similar calculation yields $\max_{f'_j \in x'} d(f'_j, x') < \varepsilon$, which implies that $d(x, x') < \varepsilon$ from the definition of the Hausdorff metric. □

Notice that $\mathbb{Q}$ in [**5.1**] is finite and can be taken to be minimal (in the sense that if $\mathbb{Q}'$ is another set that represents V_0 as in [**5.1**], then $\mathbb{Q} \subset \mathbb{Q}'$) without affecting the representation. Recall that $X^* := \{(1-t)x + t\ell^* : x \in X \text{ is finite}, t \in (0,1)\}$.

Lemma E.5. Let $\succsim$ have a representation as in [5.1]. For all $P \in \mathbb{Q}$, there exists a finite $x \in X'_P \cap X^*$. In particular, $x = \frac{1}{2}x' + \frac{1}{2}x_1(P)$ for some $x' \in X$.

Proof. The finiteness and minimality of $\mathbb{Q}$ in [5.1] implies that for any $P \in \mathbb{Q}$, there exists an open set $O \subset X'_P$. Because the space X^* is dense in X, there exists $x' \in O \cap X^*$. It follows from [5.1] that $x := \frac{1}{2}x' + \frac{1}{2}x_1(P) \in X'_P \cap X^*$, as claimed. □

Lemma E.6. Let $\succsim$ have a representation as in [5.1]. Fix $P \in \mathbb{Q}$. For any finite $x \in X'_P$ and $s \in S$, $\succsim_{(x,s)}$ is independent of the choice of $\varepsilon \in (0,1)$ for which Definition 5.3 applies and is independent of the choice of c in the definition of $x \oplus_{\varepsilon,s} y$. Further, $\succsim_{(x,s)}$ is represented by $v_s(\cdot, P)$. Finally, if $x' \in X_P$ is finite, then $\succsim_{(x,s)} = \succsim_{(x',s)}$.

Proof. Let $x \in X'_P$ be finite, so that $V_0(x) = V_0(x, P)$. Fix $s \in S$. By Lemma E.2, there exists $\varepsilon > 0$ such that $B(x;\varepsilon) \in X'_P$. By Lemmas E.2 and E.4, $[x \oplus_{\varepsilon,s} y]$, $[x \oplus_{\varepsilon,s} y'] \in B(x;\varepsilon)$ for all $y, y' \in X$, establishing the completeness of $\succsim_{(x,s)}$. Then, $[x \oplus_{\varepsilon,s} y] \succsim [x \oplus_{\varepsilon,s} y']$ if, and only if, $V_0\left(x \oplus_{(\varepsilon,s)} y\right) \geq V_0\left(x \oplus_{\varepsilon,s} y'\right)$.

Suppose $f \oplus_{\varepsilon,s} y$ is optimally chosen from $x \oplus_{\varepsilon,s} y$ in state s. Then,

$$\begin{aligned}(1-\varepsilon)[u_s(f_1(s)) + v_s(f_2(s), P)] + \varepsilon\,[u_s(c) + v_s(y, P)] \\ \geq (1-\varepsilon)\,[u_s(f_1(s)) + v_s(f_2(s), P)] + \varepsilon\,[u_s(c) + v_s(y', P)]\end{aligned}$$

which implies $v_s(y, P) \geq v_s(y', P)$. Conversely, $v_s(y, P) \geq v_s(y', P)$ implies that if $f \oplus_{\varepsilon,s} (c, y')$ is optimally chosen from $x \oplus_{\varepsilon,s} y'$ in state s, then the inequality displayed above holds, which implies $\left[x \oplus_{\varepsilon,s} y\right] \succsim \left[x \oplus_{\varepsilon,s} y'\right]$. But this is independent of our choice of $\varepsilon > 0$ as long as the menus remain in X'_P. Notice also that the argument above is independent of the choice of $c \in C$.

Finally, observe that the choice of $x \in X'_P$ in the argument above is irrelevant, in that $v_s(\cdot, P)$ represents *all* $\succsim_{(x,s)}$ with $x \in X'_P$. This completes the proof. □

Axiom 0 implies that we can define the function $v_s \in \mathbf{C}(L)$ such that $v_s(\cdot) = v_s(\cdot, P)|_L$ for all $P \in \mathbb{Q}$ in [5.1].

Lemma E.7. For all $x \in X$, $\ell \succsim_{(x,s)} \ell'$ if, and only if, $v_s(\ell) > v_s(\ell')$.

Proof. Evaluate $x \oplus_{\varepsilon,s} \ell$ according to representation [5.1] in Axiom 0. By construction, for $\varepsilon > 0$ small enough, neither the optimal partition P nor the subsequent optimal choice of act f depend on ℓ. Therefore, $\ell \succsim_{(x,s)} \ell'$ if, and only if, $V_0(x \oplus_{\varepsilon,s} \ell) - V_0(x \oplus_{\varepsilon,s} \ell') = \varepsilon\pi(s)(v_s(\ell) - v_s(\ell')) \geq 0$. □

Lemma E.8. The binary relation $\succsim_{(P,s)}$ on X which is represented by $v_s(\cdot, P)$ recursively satisfies Axiom 0 and satisfies Axiom 1.

Proof. By Lemma E.6, $\succsim_{(P,s)}=\succsim_{(x,s)}$ for some $x \in X'_P$. Since $\succsim_{(x,s)}\in \Psi^*$, it recursively satisfies Axiom 0. Further, by Lemma E.7, $\ell \succsim_{(x,s)} \ell'$ if, and only if $\ell \succsim_{(\ell^*,s)} \ell'$ for some $\ell^* \in L$. By L-stationarity (Axiom 1d), $\ell \succsim_{(\ell^*,s)} \ell'$ if, and only if $\ell \succsim \ell'$. Since $\succsim$ satisfies Axiom 1, $\succsim_{(x,s)}$ must then do so as well. $\square$

E.1.2. *Comparing Menus and Consumption Streams*

We now relate preferences on L to those on X.

Let $\tilde{X}_1 := \mathscr{K}\big(\mathscr{F}(\Delta(C \times \{\ell_*\}))\big)$ be the space of one-period problems that always give ℓ_* at the beginning the second period. Inductively define $\tilde{X}_{n+1} := \mathscr{K}\big(\mathscr{F}(\Delta(C \times \tilde{X}_n))\big)$ for all $n \geq 1$, and note that for all such n, $\tilde{X}_n \subset X$. Finally, let $\tilde{X} := \bigcup_n \tilde{X}_n$.

Lemma E.9. The set $\tilde{X} \subset X$ is dense in X.

Proof. Recall that X is the space of all consistent sequences in $\bigtimes_{n=1}^{\infty} X_n$, where $X_1 := \mathscr{K}\big(\mathscr{F}(\Delta(C))\big)$ and $X_{n+1} := \mathscr{K}\big(\mathscr{F}(\Delta(C \times X_n))\big)$. As noted in the construction in Appendix A.2, every $x \in X$ is a sequence of the form $x = (x_1, x_2, \ldots, x_n, \ldots)$ where $x_n \in X_n$, and the metric on X is the product metric.

For any $x = (x_1, x_2, \ldots) \in X$ and $n \geq 1$ set $\tilde{x}_n \in \tilde{X}_n$ to be x_n concatenated with ℓ_*. It follows from the product metric on X — see Appendix A.2 — that for any $\varepsilon > 0$, there exists $n \geq 1$ such that $d(x, \tilde{x}_n) < \varepsilon$, as claimed. $\square$

Lemma E.10. Let $\succsim$ satisfy Axioms 0 and 1. Then, for any $s \in S$ and $P \in \mathbb{Q}$, $\ell^* \succsim_{(P,s)} \ell \succsim_{(P,s)} \ell_*$ for all $\ell \in L$.

Proof. The preference $\succsim$ has a separable and partitional representation as in [**5.1**]. Therefore, $\succsim_s$ on L is represented by $u_s(\cdot) + v_s(\cdot, Q)$ for all Q. Moreover, $\succsim|_L$ has an RAA representation. As observed in Section D, $\succsim_s$ on L is separable and has the property that for all $c \in C$, $\ell \in L$ and $s \in S$, $(c, \ell^*) \succsim_s (c, \ell) \succsim_s (c, \ell_*)$. This implies that for all $\ell \in L$, $v_s(\ell^*, Q) \geq v_s(\ell, Q) \geq v_s(\ell_*, Q)$ for all partitions $Q \in \mathbb{Q}$ in [**5.1**]. But since $v_s(\cdot, P)$ represents $\succsim_{(P,s)}$, $\ell^* \succsim_{(P,s)} \ell \succsim_{(P,s)} \ell_*$ for all $\ell \in L$, $s \in S$. $\square$

Proposition E.11. Let $\succsim$ recursively satisfy Axiom 0 and satisfy Axiom 1. Then, for all $x \in X$, $\ell^* \succsim x \succsim \ell_*$. In particular, for any $P \in \mathbb{Q}$ and $s \in S$, we have $\ell^* \succsim_{(P,s)} x \succsim_{(P,s)} \ell_*$.

Proof. We shall only prove that $\ell^* \succsim x$ and that $\ell^* \succsim_{(P,s)} x$ for all $x \in X$, $P \in \mathbb{Q}$, and $s \in S$. The other bounds are proved analogously.

By the continuity of $\succsim$ and by Lemma E.9, it suffices to show that for all $\tilde{x} \in \tilde{X}$, $\ell^* \succsim \tilde{x}$.

Suppose $\tilde{x} \in \tilde{X}_n$. We first consider the case $n = 1$. It follows immediately from the representation [**5.1**] in Axiom 0 that $V(\tilde{x}_1) \leq V(\ell^*)$ for all $\tilde{x}_1 \in \tilde{X}_1$. Since $\succsim \in \Psi^*$, by Lemma E.8 $\succsim_{(P,s)}$ also satisfies Axioms 0 and 1 for any $P \in \mathbb{Q}$, which implies that there exists $\ell^*_{(P,s)}$ such that $v_s(\ell^*_{(P,s)}) \geq v_s(\tilde{x}_1, P)$ for all $\tilde{x}_1 \in \tilde{X}_1$. By Lemma E.10, we may take $\ell^*_{(P,s)} = \ell^*$, so that $v_s(\ell^*, P) \geq v_s(\tilde{x}_1, P)$ for all $\tilde{x}_1 \in \tilde{X}_1$.

Now consider the induction hypothesis: If $\succsim$ recursively satisfies Axiom 0 and satisfies Axiom 1, then for all $\tilde{x}_n \in \tilde{X}_n$, $\ell^* \succsim \tilde{x}_n$. Suppose the induction hypothesis is true for some $n \geq 1$. We shall now show that it is also true for $n + 1$.

By Lemma E.8, $\succsim_{(P,s)}$ recursively satisfies Axiom 0 and satisfies Axiom 1 on X, so that $v_s(\ell^*, P) \geq v_s(\tilde{x}_n, P)$ for all $\tilde{x}_n \in \tilde{X}_n$ (where we have appealed to Lemma E.10 to establish that ℓ^* is the $v_s(\cdot, P)$-best consumption stream). In particular, this implies that for any lottery $\alpha_2 \in \Delta(\tilde{X}_n)$, $v_s(\ell^*, P) \geq v_s(\alpha_2, P)$, which along with Axiom 0 proves the second part of the proposition because n is arbitrary.

Now consider any $\tilde{x}_{n+1} \in \tilde{X}_{n+1}$. We have, for any choice of P,

$$\begin{aligned}
V(\tilde{x}_{n+1}, P) &= \max_{f \in \tilde{x}_{n+1}} \sum_{J \in P} \pi_0(s \mid J)\big[u_s(f_1(s)) + v_s(f_2, P)\big] \\
&\leq \sum_{J \in P} \pi_0(s \mid J)\big[u_s(c_s^+) + v_s(\ell^*, P)\big] \\
&= V(\ell^*, P) = V(\ell^*)
\end{aligned}$$

where we have used the facts that $f_1(s) \in \Delta(C)$ and $f_2(s) \in \Delta(\tilde{X}_n)$, and that $u_s(c_s^+)$ and $v_s(\ell^*; P)$ respectively dominate all such lotteries, as established above. Thus, for all $\tilde{x}_{n+1} \in \tilde{X}_{n+1}$, $\ell^* \succsim \tilde{x}_{n+1}$, which completes the proof. □

E.1.3. *Proof of Proposition E.3*

Proof. To see that (b) implies (a), suppose $\succsim$ has the representation [**5.1**], thereby meeting Axiom 0. We now establish is that $\succsim$ recursively satisfies the axiom.

Given a representation as in [**5.1**] that is also self-generating, let $x \in X$ be finite and $P \in \mathbb{Q}$ be the unique optimal partition for x. By Lemma E.6, $\succsim_{(x,s)}$ is represented by $v_s(\cdot, P)$ on X. Because the representation is self-generating, $\succsim_{(x,s)}$ must satisfy Axiom 0 on X. Because $V \in \Phi^*$, the same argument applies to preferences induced by

$\succsim_{(x,s)}$, and so on, ad infinitum, which establishes that $\succsim$ recursively satisfies Axiom 0 . Moreover, $\succsim$ has an RAA representation on L, which means it satisfies Axiom 1 on L.

To see that (a) implies (b), note that Lemma E.8 has two implications. First, $\succsim_{(P,s)}$ has a separable and partitional representation $v'_s(\cdot, P)$ as in [5.1]. Because $v_s(\cdot, P)$ also represents $\succsim_{(P,s)}$, $v_s(\cdot, P)$ and $v'_s(\cdot, P)$ are identical up to a monotone transformation. But, by L-Indifference to Timing (Axiom 1(e)), it must be that $v_s(\cdot, P)$ and $v'_s(\cdot, P)$ are unique up to a positive affine transformation on L. Let us re-normalize $v'_s(\cdot, P)$ so that $v_s(\cdot, P) = v'_s(\cdot, P)$ on L.

Second, because $\succsim_{(P,s)}$ recursively satisfies Axiom 0 and satisfies Axiom 1, it satisfies the hypotheses of Proposition E.11. By Lemma E.11, we have $\ell^* \succsim_{(P,s)} y \succsim_{(P,s)} \ell_*$ for all $y \in X$. Because $v_s(\cdot, P)$ and $v'_s(\cdot, P)$ both represent $\succsim_{(P,s)}$, they agree on L and, thus, as any element in X has an indifferent alternative in L, must agree on X. It follows that $v_s(\cdot, P)$ also has a representation as in [5.1], that is, it can be written as

$$v_s(x, P) = \max_{P' \in \mathbb{Q}'(P)} \sum_{J \in P'} \pi_0(J) \max_{f \in x} \sum_{s} \pi_0(s \mid J) \left[u'_s(f_1(s)) + v'_s(f_2(s); P')\right]$$

Because $\succsim_{(x,s)}$ recursively satisfies Axiom 0 and satisfies Axiom 1, it follows from the reasoning above that each $v'_s(\cdot, P')$ in the above representation of $v_s(\cdot, P)$ also has a representation as in [5.1], and so on, ad infinitum, which demonstrates that $V \in \Phi^*$. $\square$

E.2. Extracting the Canonical ICP

Given a $V \in \Phi^*$ that is a self-generating representation of $\succsim$, we would like to extract the underlying (subjective) informational constraints. We show next that this is possible.

Proposition E.12. There is a unique map $\varphi^* : \Phi^* \to \Omega$ that for some implementation $\big((u_s), \mathbb{Q}, (v_s(\cdot, P)), \pi\big)$ of v satisfies

$$\varphi^*(v) := \Big\{ \big(P, \varphi^*(v_s(\cdot, P))\big) : P \in \mathbb{Q} \Big\}$$

and is independent of the implementation chosen.

Proof. Let $v^{(1)} \in \Phi_1$, and suppose $\big((u_s), \mathbb{Q}, (v_s(\cdot, P)), \pi\big)$ implements $v^{(1)}$. In this implementation, $\mathbb{Q}$ is unique. (The argument follows from our identification argument in Appendix C. On the other hand, (u_s), $(v_s(\cdot, P))$, and π will typically not be unique.) Then, define $\varphi_1 : \Phi_1 \to \Omega_1$ as

$$\varphi_1(v^{(1)}) := \mathbb{Q}, \qquad \text{where } \big((u_s), \mathbb{Q}, (v_s(\cdot, P)), \pi\big) \text{ implements } v^{(1)}$$

Proceeding iteratively, we define $\varphi_n : \Phi_n \to \Omega_n$ as

$$\varphi_n(v^{(n)}) := \Big\{ \big(P, \varphi_{n-1}(v_s^{(n-1)}(\cdot, P)\big) : \; \exists \, \big((u_s), \mathbb{Q}, (v_s^{(n-1)}(\cdot, P)), \pi\big) \text{ that} \\ \text{implements } v^{(n)} \text{ and } P \in \mathbb{Q} \Big\}$$

Notice that the same argument that established the uniqueness of φ_1 also applies here to provide the uniqueness of φ_n.

Suppose $v \in \Phi^*$. This implies that v has a partitional and separable implementation $\big((u_s), \mathbb{Q}, (v_s(\cdot, P)), \pi\big)$, where each $v_s(\cdot, P)$ also has a partitional and separable implementation, and so on, ad infinitum. Then, we may define, for all $n \geq 1$, $\omega^{(n)} := \varphi_n(v)$. Consider the infinite sequence

$$\omega_0 := \big(\omega^{(1)}, \omega^{(2)}, \dots, \omega^{(n)}, \dots\big) \in \Omega$$

and define the map $\varphi^* : \Phi^* \to \Omega$ as $\varphi^*(v) = \big(\varphi_1(v), \varphi_2(v), \dots\big)$, which extracts the underlying canonical ICP from any function $v \in \Phi^*$, independently of the other components of the implementation, as claimed. $\square$

To recapitulate, we can now extract a canonical ICP from a self-generating representation. In other words, the identification of the canonical ICP ω_0 doesn't depend on the recursivity of the value function. This stands in contrast to the identification of the other preference parameters, which relies on recursivity.

A dynamic plan consists of two parts: the first entails picking a partition for the present period (and the corresponding continuation constraint), and the second entails picking an act from x, whilst requiring that the choice of act, as a function of the state, be measurable with respect to the chosen partition. The first part is a dynamic information plan while the second is a dynamic consumption plan.

An n-period history is an (ordered) tuple

$$\mathfrak{h}_n = \big((x^{(0)}, \omega^{(0)}), \dots, (P^{(n-1)}, f^{(n-1)}, s^{(n-1)}, x^{(n-1)}, \omega^{(n-1)})\big)$$

Let $\mathfrak{H}_n$ denote the collection of all n-period histories.

Formally, a *dynamic information plan* is a sequence $\sigma_i = (\sigma_i^{(1)}, \sigma_i^{(2)}, \dots)$ of mappings where $\sigma_i^{(n)} : \mathfrak{H}_n \to \mathscr{P} \times \Omega^S$. Similarly, a *dynamic consumption plan* is a sequence $\sigma_c = (\sigma_c^{(1)}, \sigma_c^{(2)}, \dots)$ of mappings where $\sigma_c^{(n)} : \mathfrak{H}_n \to \mathscr{F}(\Delta(C \times X))$. A *dynamic plan* σ is a pair $\sigma = (\sigma_i, \sigma_c)$.

A dynamic plan $\sigma = (\sigma_i, \sigma_c)$ with initial states $x^{(0)} := x$ and $\omega^{(0)} := \omega_0$ is feasible if (i) $\sigma_i^{(n)}(\mathfrak{h}_n) \in \omega^{(n-1)}$, (ii) $\sigma_c^{(n)}(\mathfrak{h}_n) \in x^{(n-1)}$, and (iii) given the information plan $\sigma_i^{(n)}(\mathfrak{h}_n) = (P, \boldsymbol{\omega}') \in \omega^{(n-1)}$, $\sigma_c^{(n-1)}(\mathfrak{h}_n)$ is P-measurable, ie, for all $I \in P$ and for all $s, s' \in I$, $\sigma_c^{(n)}(\mathfrak{h}_n)(s) = \sigma_c^{(n)}(\mathfrak{h})(s')$.

Each dynamic plan along with initial states (x, ω_0, π_0) induces a probability measure over $\big(X \times \Omega \times S\big)^\infty$ or, put differently, an $X \times \Omega \times S$ valued process. Let $(x^{(n)}, \omega^{(n)}, s^{(n)})$ be the $X \times \Omega \times S$-valued stochastic process of menus, canonical ICPs, and objective states induced by a dynamic plan, where $x^{(n)} \in X$ is the menu beginning at period $n+1$, $\omega^{(n)} \in \Omega$ is the canonical ICP beginning at period $n+1$, and $s^{(n)} \in S$ is the state in period n. A dynamic plan is stationary if $\sigma^{(n)}(\mathfrak{h}_n)$ only depends on $(x^{(n-1)}, \omega^{(n-1)}, s^{(n-1)})$.[25]

For a fixed $V \in \Phi^*$, let $v^{(n)}(\cdot, \omega^{(n)}, s^{(n)}, \sigma)$ denote the value function that corresponds to the n-th period implementation of V when following the dynamic information plan σ, where $\omega^{(n)} = \varphi_n(V)$ as in Proposition E.12 and $s^{(n)}$ is the state in period n.

While we have shown that each $v \in \Phi^*$ can be written as the sum of some instantaneous utility and some continuation utility function that also lies in Φ^*, we still need to verify that the value that V obtains for any menu is indeed the infinite sum of consumption utilities. We verify this next.

Proposition E.13. Let $V \in \Phi^*$, and suppose $v^{(n)}(\cdot, \omega^{(n)}, s^{(n)}, \sigma)$ is defined as above. Then, for any feasible dynamic plan $\sigma = (\sigma_c, \sigma_i)$, we have

$$\lim_{n\to\infty} \left\| \mathsf{E}^{\sigma,\boldsymbol{\pi}}\, v^{(n)}(\cdot, \omega^{(n)}, s^{(n)}, \sigma) \right\|_\infty = 0 \tag{E.1}$$

Proof. Consider $V \in \Phi^*$ with Lipschitz rank λ. Recall that for any $x \in X$, $\ell^\dagger \diamond_n x \in X$ denotes the menu that delivers $\ell^\dagger$ in every period until period $n-1$ and then, in period n, in every state, delivers x. Recall further that X is an infinite product space, and by the definition of the product metric (see Appendix A.2), it follows that for any $\varepsilon > 0$, there exists an $N > 0$ such that for all $x, y \in X$ and $n \geq N$, $d(\ell^\dagger \diamond_n x, \ell^\dagger \diamond_n y) < \varepsilon/\lambda$. Lipschitz continuity of V then implies $\big|V(\ell^\dagger \diamond_n x) - V(\ell^\dagger \diamond_n y)\big| < \varepsilon$.

For a given n, $V(\ell^\dagger \diamond_n x) = 0 + \mathsf{E}^{\sigma,\boldsymbol{\pi}}[v^{(n)}(x, \omega^{(n)}, s^{(n)}, \sigma)]$, which implies

$$\big|\mathsf{E}^{\sigma,\boldsymbol{\pi}}\, v^{(n)}(x, \omega^{(n)}, s^{(n)}) - \mathsf{E}^{\sigma,\boldsymbol{\pi}}\, v^{(n)}(y, \omega^{(n)}, s^{(n)}, \sigma)\big| < \varepsilon$$

(25) The choice of plan clearly does not affect the evolution of the objective states $(s^{(n)})$.

for all $n \geq N$. Recall that

$$\left\| \mathsf{E}^{\sigma,\boldsymbol{\pi}}\, v^{(n)}(\cdot, \omega^{(n)}, s^{(n)}, \sigma) \right\|_\infty = \sup_x \left| \mathsf{E}^{\sigma,\boldsymbol{\pi}}\, v^{(n)}(x, \omega^{(n)}, s^{(n)}, \sigma) \right|$$

Moreover, we have

$$\begin{aligned} &\sup_x \left| \mathsf{E}^{\sigma,\boldsymbol{\pi}}\, v^{(n)}(x, \omega^{(n)}, s^{(n)}, \sigma) \right| \\ &\qquad = \sup_x \left| \mathsf{E}^{\sigma,\boldsymbol{\pi}} \left[v^{(n)}(x, \omega^{(n)}, s^{(n)}, \sigma) - v^{(n)}(\ell^\dagger, \omega^{(n)}, s^{(n)}, \sigma) \right] \right| < \varepsilon \end{aligned}$$

which completes the proof. □

Adapting the terminology of Dubins and Savage (1976), we say that a function $V \in \Phi^*$ is equalizing if [**E.1**] holds. (To be precise, if [**E.1**] holds, then every dynamic plan is equalizing in the sense of Dubins and Savage (1976).)

Given an initial $(x, \omega) \in X \times \Omega$, each σ induces a probability measure over $\bigtimes_n \mathfrak{H}_n$, the space of all histories. It also induces a unique consumption stream $\ell_{\sigma(x,\omega)}$ that delivers consumption $\sigma_c(\mathfrak{h}_n)(s')$ after history $\mathfrak{h}_n$ in state s' in period n. We show next that for any self-generating preference functional $V \in \Phi^*$, the utility from following the plan σ given the menu x is the same as the utility from the consumption stream $\ell_{\sigma(x,\omega)}$. (Recall that there are no consumption choices to be made for the consumption stream $\ell_{\sigma(x,\omega,s)}$.) Moreover, there is an optimal plan such that following this plan induces a consumption stream that produces the same utility as the menu x.

Let Σ denote the collection of all dynamic plans and let $L_{x,\omega} := \{\ell_{\sigma(x,\omega)} : \sigma \in \Sigma\}$ be the collection of all consumption streams so induced by the menu x and the canonical ICP ω. In what follows, $V(x, \sigma)$ is the expected utility from following the dynamic plan σ given the menu x.

Lemma E.14. Let $V \in \Phi^*$ be such that $\varphi^*(V) = \omega$. Then, for all $x \in X$, $V(x, \sigma) = V(\ell_{\sigma(x,\omega)})$ and $V(x) = \max_{\sigma \in \Sigma} V(x, \sigma) = \max_{\ell \in L_{x,\omega}} V(\ell)$.

Proof. For $V \in \Phi^*$ and for any plan σ', an agent with the utility function V is indifferent between following σ' and the consumption stream $\ell_{\sigma'(x,\omega)}$. This is essentially an adaptation of Theorem 9.2 in Stokey, Lucas, and Prescott (1989) where their Equation [7] — which is also known as a no-Ponzi game condition, see Blanchard and Fischer (1989, p 49) — is replaced by the fact that V is equalizing (condition [**E.1**] in Proposition E.13).[26]

(26) Note that Stokey, Lucas, and Prescott (1989) directly work with the optimal plan, but the essential idea is the same — continuation utilities arbitrarily far in the future must contribute arbitrarily little.

To see that there is an optimal plan, notice that x is a compact set of acts, and because there are only finitely many partitions of S, it is possible to find a conserving action at each date after every history. This then gives us a conserving plan. We can now adapt Theorem 9.2 in Stokey, Lucas, and Prescott (1989) where, as above, their Equation [7] is replaced by [**E.1**], to show that σ is indeed an optimal plan. Loosely put, we have just shown that because the plan is conserving and because V is equalizing, the plan must be optimal. This corresponds to the characterization of optimal plans in Theorem 2 of Karatzas and Sudderth (2010). $\square$

E.3. Recursive Representation

We now establish a recursive representation for $\succsim$, thereby proving Theorem 3.

Recall that $\succsim|_L$ has a standard RAA representation $\big((u_s), \delta, \Pi\big)$. That is, there exist functions $V_L^*(\cdot, s) : L \to \mathbb{R}$ such that $V_L^*(\ell, \pi_0) := \sum_s \pi_0(s) V_L^*(\ell, s)$ represents $\succsim|_L$, and

$$V_L^*(\ell, s) := \sum_{s'} \Pi(s, s') \big[u_{s'}(\ell_1(s')) + \delta V_L^*(\ell_2(s'), s')\big]$$

where u_s is nontrivial for some $s \in S$ and $u_s\big(c_s^\dagger\big) = 0$ for all $s \in S$. This implies $V_L^*(\ell^\dagger, s) = 0$ for all s, so that $V_L^*(\ell^\dagger, \pi_0) = 0$. The function V_L^* (recall that V_L^* also denotes the linear extension of V_L^* to $\Delta(L)$) is uniquely determined by the tuple $\big((u_s)_{s \in S}, \delta, \Pi\big)$.

By Proposition E.3 $\succsim$ has a self-generating representation $V \in \Phi^*$ that satisfies $V(\ell^\dagger) = 0$. Now, both $V|_L$ and $V_L^*(\cdot, \pi_0)$ represent $\succsim|_L$ on L. Because $\succsim|_L$ is continuous and satisfies Independence on L, it follows that $V|_L$ and $V_L^*(\cdot, \pi_0)$ are identical up to a positive affine transformation. Given that $V(\ell^\dagger) = V_L^*(\ell^\dagger, \pi_0) = 0$, $V|_L$ and $V_L^*(\cdot, \pi_0)$ only differ by a scaling. Therefore, rescale the collection $(u_s)_{s \in S}$ by a common factor so as to ensure $V|_L = V_L^*(\cdot, \pi_0)$ on L.

Fix ω_0 and observe that by Proposition 3.2, the tuple $\big((u_s)_{s \in S}, \Pi, \delta, \omega_0\big)$ induces a unique value function that satisfies [**Val**]. Notice also that this value function agrees with $V_L^*(\cdot, \pi_0)$ on L. We denote this value function, defined on $X \times \Omega \times S$, by $V^*(\cdot, \omega_0, \pi_0)$.

The next result proves Theorem 3.

Proposition E.15. Let V be a self-generating representation of $\succsim$ such that $\varphi^*(V) = \omega_0$, and suppose $V(\cdot) = V^*(\cdot, \omega_0, \pi_0)$ on L. Then, $V(\cdot) = V^*(\cdot, \omega_0, \pi_0)$ on X.[27]

(27) It follows immediately from Proposition E.15 that in considering dynamic plans, we may restrict attention to stationary plans. This is because we have a recursive formulation with discounting

Proof. In this proof, we refer to objects defined in Appendix E.2. For any x, let $\sigma(x, \omega_0)$ denote the optimal plan for the utility V and let $\sigma^*(x, \omega_0)$ denote the optimal plan for V^*. By Lemma E.14, there exist $\ell_{\sigma(x,\omega_0)}, \ell_{\sigma^*(x,\omega_0)} \in L_{x,\omega_0}$ such that

$$V(x) = V(\ell_{\sigma(x,\omega_0)}) \geq V(\ell_{\sigma^*(x,\omega_0)}) = V^*(\ell_{\sigma^*(x,\omega_0)}) = V^*(x, \omega_0, \pi_0)$$

Reversing the roles of V and V^*, we obtain once again from Lemma E.14 that

$$V^*(x, \omega_0, \pi_0) = V^*(\ell_{\sigma^*(x,\omega_0)}) \geq V^*(\ell_{\sigma(x,\omega_0)}) = V(\ell_{\sigma(x,\omega_0)}) = V(x)$$

In both displays, the second equality obtains because V and V^* agree on L. Combining the two inequalities yields the desired result. $\square$

Suppose V represents $\succsim$ and $V \in \Phi^*$. Then there exists an implementation of V given by $\big((u_s), \mathbb{Q}, (v_s^{(1)}(\cdot, P)), \pi\big)$. For ease of exposition, we shall say that the collection $(v_s^{(1)}(\cdot, P))$ implements V. Then, for all $n \geq 1$, there exists $(v_s^{(n)}) \in \Phi^*$ that implements $v_s^{(n-1)}$ and so on. Notice that each $v_s^{(n)}$ depends on all the past choices of partitions. However, our recursive representation V^* is only indexed by the current state of the canonical ICP, and so is entirely forward looking.

References

Abreu, Dilip, David Pearce, and Ennio Stacchetti. 1990. "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring". *Econometrica* 58 (5): 1041–1063. (Cit. on pp. 4, 24, 44, 45).

Bertsekas, Dimitri P., and Steven Shreve. 2000. *Stochastic Optimal Control: The Discrete-Time Case*. Belmont, MA: Athena Scientific. (Cit. on p. 54).

Blanchard, Olivier Jean, and Stanley Fischer. 1989. *Lectures on Macroeconomics*. Cambridge, Ma: MIT Press. (Cit. on p. 52).

Bolton, Patrick, and Mathias Dewatripont. 1994. "The Firm as Communication Network". *Quarterly Journal of Economics* 109 (4): 809–839. (Cit. on p. 5).

Caplin, Andrew, and Mark Dean. 2015. "Revealed Preference, Rational Inattention, and Costly Information Acquisition". *American Economic Review* 105 (7): 2183–2203. (Cit. on pp. 11, 28).

where all our payoffs are bounded, which obviates the need for non-stationary plans — see, for instance, Proposition 4.4 of Bertsekas and Shreve (2000) or Theorem 1 of Orkin (1974).

Coeurdacier, Nicolas, and Hélène Rey. 2013. "Home Bias in Open Economy Financial Macroeconomics". *Journal of Economic Literature* 51 (1): 63–115. (Cit. on p. 6).

Currie, Janet, and W. Bentley MacLeod. 2017. "Diagnosing Expertise: Human Capital, Decision Making and Performance Among Physicians". *Journal of Labor Economics* 35 (1): 1–43. (Cit. on p. 5).

de Oliveira, Henrique. 2018. "Blackwell's Informativeness Theorem using Diagrams". *Games and Economic Behavior* 109:126–131. (Cit. on pp. 13, 16).

de Oliveira, Henrique, Tommaso Denti, Maximilian Mihm, and M. Kemal Ozbek. 2017. "Rationally Inattentive Preferences and Hidden Information Costs". *Theoretical Economics* 12:621–654. (Cit. on pp. 3, 11, 26, 27).

Dekel, Eddie, Barton L. Lipman, and Aldo Rustichini. 2001. "Representing Preferences with a Unique Subjective State Space". *Econometrica* 69 (4): 891–934. (Cit. on pp. 26, 27).

Dillenberger, David, R. Vijay Krishna, and Philipp Sadowski. 2020. *Information-dependent Utilities and Beliefs*. Tech. rep. Duke University. (Cit. on p. 19).

Dillenberger, David, Juan S. Lleras, Philipp Sadowski, and Norio Takeoka. 2014. "A Theory of Subjective Learning". *Journal of Economic Theory* 153:287–312. (Cit. on pp. 26–28).

Dubins, Lester E., and Leonard J. Savage. 1976. *Inequalities for Stochastic Processes: How to Gamble if You Must.* Mineola, New York: Dover Publications Inc. (Cit. on p. 52).

Dudley, Richard Mansfield. 2002. *Real Analysis and Probability.* 2nd edition. Cambridge University Press. (Cit. on p. 28).

Ellis, Andrew. 2018. "Foundations for Optimal Inattention". *Journal of Economic Theory* 173:56–94. (Cit. on pp. 3, 28).

Ergin, Haluk, and Todd Sarver. 2010. "A Unique Costly Contemplation Representation". *Econometrica* 78 (4): 1285–1339. (Cit. on pp. 11, 26, 27).

Gabaix, Xavier. 2014. "A Sparsity-Based Model of Bounded Rationality". *Quarterly Journal of Economics* 129 (4): 1661–1710. (Cit. on p. 10).

Geanakoplos, John, and Paul R. Milgrom. 1991. "A Theory of Hierarchies Based on Limited Managerial Attention". *Journal of Japanese and International Economies* 5:205–225. (Cit. on pp. 2, 5).

Gilboa, Itzhak, and David Schmeidler. 1989. "Maxmin Expected Utility with a Non-Unique Prior". *Journal of Mathematical Economics* 18:141–153. (Cit. on p. 21).

Greenshtein, Eitan. 1996. "Comparison of Sequential Experiments". *Annals of Statistics* 24 (1): 436–448. (Cit. on pp. 13, 16).

Gul, Faruk, and Wolfgang Pesendorfer. 2006. "Random Expected Utility". *Econometrica* 74 (1): 121–146. (Cit. on p. 18).

– . 2004. "Self-Control and the Theory of Consumption". *Econometrica* 72 (1): 119–158. (Cit. on pp. 8, 27, 29).

Hébert, Benjamin, and Michael Woodford. 2016. *Rational Inattention with Sequential Information Sampling.* Tech. rep. Stanford University GSB. (Cit. on p. 11).

Higashi, Youichiro, Kazuya Hyogo, and Norio Takeoka. 2009. "Subjective random discounting and intertemporal choice". *Journal of Economic Theory* 144 (3): 1015 –1053. (Cit. on p. 14).

Karatzas, Ioannis, and William Sudderth. 2010. "Two Characterizations of Optimality in Dynamic Programming". *Applied Mathematics and Optimization* 61:421–434. (Cit. on p. 53).

Kreps, David M. 1979. "A Representation Theorem for 'Preference for Flexibility'". *Econometrica* 47 (3): 565–577. (Cit. on p. 26).

– . 2012. *Microeconomic Foundations I: Choice and Competitive Markets.* Princeton, NJ: Princeton University Press. (Cit. on p. 33).

Kreps, David M., and Evan L Porteus. 1978. "Temporal Resolution of Uncertainty and Dynamic Choice Theory". *Econometrica* 46 (1): 185–200. (Cit. on p. 8).

Krishna, R. Vijay, and Philipp Sadowski. 2014. "Dynamic Preference for Flexibility". *Econometrica* 82 (2): 655–703. (Cit. on pp. 8, 13, 14, 21, 25–27, 39, 41).

Laffont, Jean-Jacques. 1989. *The Economics of Uncertainty and Information.* The MIT Press. (Cit. on p. 38).

Lian, Chen. 2019. *Consumption with Imperfect Perception of Wealth.* Tech. rep. MIT. (Cit. on p. 10).

Lu, Jay. 2016. "Random Choice and Private Information". *Econometrica* 84 (6): 1983–2027. (Cit. on p. 28).

Maccheroni, Fabio, Massimo Marinacci, and Aldo Rustichini. 2006. "Dynamic Variational Preferences". *Journal of Economic Theory* 128:4–44. (Cit. on p. 27).

Maćkowiak, Bartosz, and Mirko Wiederholdt. 2009. "Optimal Sticky Prices under Rational Inattention". *American Economic Review* 99 (3): 769–803. (Cit. on p. 6).

Magnac, Thierry, and David Thesmar. 2002. "Identifying Dynamic Discrete Decision Processes". *Econometrica* 70 (2): 801–816. (Cit. on p. 17).

Mariotti, Thomas, Martin Meier, and Michele Piccione. 2005. “Hierarchies of Beliefs for Compact Possibility Models”. *Journal of Mathematical Economics* 41 (3): 303–324. (Cit. on p. 32).

Massa, Massimo, and Andrei Simonov. 2006. “Hedging, Familiarity and Portfolio Choice”. *Review of Financial Studies* 19 (2): 633–685. (Cit. on p. 6).

Matějka, Filip, and Alisdair McKay. 2015. “Rational Inattention to Discrete Choices: A New Foundation for the Multinomial Logit Model”. *American Economic Review* 105 (1): 272–298. (Cit. on p. 28).

Mullainathan, Sendhil, and Eldar Shafir. 2013. *Scarcity*. Times Books. (Cit. on p. 4).

Orkin, Michael. 1974. “On Stationary Policies—The General Case”. *The Annals of Statistics* 2 (1): 219–222. (Cit. on pp. 33, 54).

Persico, Nicola. 2000. “Information Acquisition in Auctions”. *Econometrica* 68 (1): 135–148. (Cit. on p. 2).

Piermont, Evan, Norio Takeoka, and Roee Teper. 2016. “Learning the Krepsian State: Exploration Through Consumption”. *Games and Economic Behavior* 100 (69–94). (Cit. on p. 28).

Rust, John. 1994. “Structural Estimation of Markov Decision Processes”. In *Handbook of Econometrics*, ed. by Robert F. Engle and Daniel F. McFadden, vol. IV. Elsevier. (Cit. on pp. 17, 18).

Sims, Christopher A. 2003. “Implications of Rational Inattention”. *Journal of Monetary Economics* 50 (3): 665–690. (Cit. on p. 2).

– . 2011. “Rational Inattention and Monetary Economics”. In *Handbook of Monetary Economics, Volume 3A*, ed. by Benjamin N. Friedman and Mark Woodford. Elsevier. (Cit. on p. 2).

– . 1998. “Stickiness”. *Carnegie-Rochester Conference Series on Public Policy* 49:317–356. (Cit. on p. 2).

Steiner, Jakub, Colin Stewart, and Filip Matějka. 2017. “Rational Inattention Dynamics: Inertia and Delay in DecisionâĂŘMaking”. *Econometrica* 85 (2): 521–553. (Cit. on p. 10).

Stigler, George. 1961. “The Economics of Information”. *Journal of Political Economy* 69 (3): 213–225. (Cit. on p. 2).

Stokey, Nancy L., Robert E. Lucas Jr, and Edward C. Prescott. 1989. *Recursive Methods in Economic Dynamics*. Cambridge, Ma: Harvard University Press. (Cit. on pp. 52, 53).

Woodford, Michael. 2012. *Inattentive Valuation and Reference-Dependent Choice*. Tech. rep. Columbia University. (Cit. on p. 11).