

Subjective Information Choice Processes*

DAVID DILLENBERGER†

R. VIJAY KRISHNA‡

PHILIPP SADOWSKI§

April 23, 2018

Abstract

We propose a class of dynamic choice 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. Additionally, we provide behavioral foundations, ie axioms, for the model.

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 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.

(*) We thank Attila Ambrus, Alexander Bloedel, Andrew Ellis, Itzhak Gilboa, Ed Green, Botond Köszegi, 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>

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 (as is evident from the proliferation of online comparison shopping engines) 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 to outsiders.

In this paper we provide — and fully identify — a class of dynamic models that incorporate *intertemporal* information constraints. Just as with intertemporal budget constraints, intertemporal information 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, these constraints 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. In particular, it applies whether the constraints are cognitive, so that individuals have limited ability to take into account available information, or physical, where the constraint reflects the scarcity of information.

To fix ideas, suppose DM has to manage his portfolio. In each period, there is a set of possible investments he can make, given the monetary value of the portfolio. 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 in order to improve his portfolio choice, DM must also take an unobservable action to acquire information about the true state. His choice of information may affect the feasible set of information strategies in the next period. For instance, it may be that DM is 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 its impact on the feasibility of subsequent information strategies, is not. The following two questions are natural:

(a) Can (unobservable) information constraints be identified from DM's preferences and what

(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).

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 is Theorem 1, which 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. To answer (b), we provide behavioral foundations for the model (Theorem 3), which is challenging because our axioms cannot condition directly on the status of the subjectively controlled process that governs DM's flow of information. We overcome this by introducing a decision-theoretic version of the dynamic programming operator which lets us impose our axioms recursively, and even after histories that have subjective (and hence unobservable) components.

Formally, DM's *Information Choice Process* (ICP) is an entirely subjective control problem, which specifies how future constraints depend on current and past choices of information that are *unobservable* (to the analyst). The payoff-relevant state $s \in S$ changes over time, and we focus on partitional learning in every period.³ The ICP is parametrized by a control state θ , a function $\Gamma(\theta)$ that determines the set of feasible partitions of S , and an operator τ that governs the transition of θ in response to the choice of partition and the realization of $s \in S$. Examples of such ICPS are in Section 2.

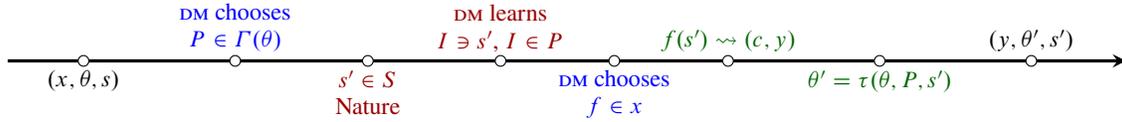


Figure 1: Timeline

The domain of choice consists of sets (or *menus*) of actions, where each action (an *act* on S) results in a state-contingent lottery that yields current consumption and a new menu of acts for the next period. 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 $\pi_s(\cdot | I)$ and then chooses an act f from 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

(2) Recent work on rational inattention 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.

(3) For a discussion of this restriction, see Section 3.3.

$\pi_{s'}$ are determined for next period. The collection of measures $(\pi_s)_{s \in S}$ defines the transition operator for a Markov process on S .

DM's objective is then to maximize the expected utility which consists of state-dependent consumption utilities, (u_s) , and the discounted continuation value:

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

Theorem 1 establishes that $(u_s, \pi_s)_{s \in S}$ and δ are essentially uniquely identified, and that the remaining preference parameter, the ICP, is identified up to the addition or deletion of information strategies that are dominated in an appropriate sense. Identifying the subjectively controlled process (the ICP) from behavior is our main conceptual contribution.

An information plan for DM is a choice of partition conditional on (x, s, θ) . Our approach allows us to infer whether a certain information plan is feasible for DM from observing just one binary comparison. Assume for simplicity that consumption takes values in $[0, 1]$. We construct a menu x (explicitly in Section 4.1) such that DM is indifferent between x and the constant stream of 1 if, and only if, said information plan (or one that is even more informative) is feasible. The menu x resembles a gambling game — as long as DM follows the relevant information plan (or a more informative one), he can make the right bets, which give payoff 1 for sure and keep him in the game — otherwise, with positive probability and in finite time, he will make a bet that has an instantaneous payoff of 0, and after which he gets a payoff of 0 forever, ie, he effectively stops playing. We say that the menu x is *strongly aligned* with the information plan at hand.

Our identification strategy uses such menus to circumscribe the set of all information plans available to DM, which is the behaviorally relevant aspect of his ICP. If DM only faces length T problems, then it suffices to consider ICPS of length T . Such ICPS can be completely identified with a finite amount of data, because there are only finitely many possible length T information plans, and each has a choice problem that is strongly aligned with it.

Apart from establishing that our model does not have free parameters, we see three main benefits of our result. First, it establishes that one snapshot of preferences over continuation problems is sufficient to *forecast* future choice between continuation problems for any menu and after any history of state realizations.⁴ 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 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

(4) That is, full identification, together with the assumption that the model describes how DM *actually* anticipates to choose, allows for the prediction of choice frequencies. This is the case for any menu-choice model that is fully identified.

to fully inform themselves. In order to ameliorate the effects of information constraints, a policy maker must first understand them.

In Section 5 we discuss an axiomatic foundation for our model, with the formal statement of the axioms being deferred to Appendix C. Almost all our axioms are static and establish aspects of standard properties — such as Independence and Temporal Separability, which are central to virtually all existing axiomatic models of dynamic choice — that hold even when behavior depends on contemporaneous unobservable information choice.⁵ The most important innovation behind our representation theorem (Theorem 3) lies not in the content of these axioms, but in the recursive way they are applied: Preferences over dynamic choice problems are required to have the same properties as each of the preferences over continuation problems that together generate them (and thus embody the dynamic programming principle). The main difficulty in this recursive application of our axioms is that continuation preferences can depend not only on the public history, generated by the observable states s , but also on DM’s privately observed information choices, as summarized by the control state θ .⁶ Our recursive formulation hinges on identifying menus where DM has a *unique* subjectively-optimal initial information choice. For those menus, we can perturb all acts with respect to the continuation problem they offer for a fixed state without changing the optimal information. Restrictions on DM’s preferences over those perturbations then indirectly restrict DM’s continuation preferences contingent on a partially private (but fixed) history.

The paper is organized as follows. Section 2 presents examples of ICPS and some of the behavioral patterns they can generate. In Section 3 we introduce the analytical framework, state our utility representation, and describe our notion of comparative informativeness for ICPS. Section 4 establishes our identification result. Section 5 discusses our axioms, and provides the representation theorem. Section 6 surveys related decision-theoretic literature, while other related literature is discussed when relevant. Section 7 extends our identification strategy to a model with direct information costs. The main proofs can be found in the Appendix; additional technical details are in the Supplementary Appendix, Dillenberger, Krishna, and Sadowski (2017).⁷

2. Examples of ICPS and Patterns of Behavior

Recall that an ICP is a tuple $(\Theta, \Gamma, \tau, \theta_0)$, as discussed in the [Introduction](#). ICPS can accommodate *any* dependence of the information constraint on the history of information choices and state realizations. We first describe an ICP that is based on models in the literature on

(5) The only axiom that is not static restricts preferences over consumption streams, which require no future choice so that information constraints are irrelevant.

(6) In principle, past choices of acts might also directly affect continuation preferences. We assume that this is not the case, so that the relevant history consists only of s and θ .

(7) The Supplementary Appendix, Dillenberger, Krishna, and Sadowski (2017), is available online at <http://people.duke.edu/~ps78/documents/ICP-supp.pdf>.

managerial decision making and constrains the accumulation of human capital, or expertise, from experience.⁸ We then provide a few more simple and illustrative examples of ICPS.

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. Allocating time κ_i to source i yields the partition $P_{\theta_i \kappa_i}^i$, where θ_i is a parameter measuring the manager's efficiency in *processing* information from source i , and P_{σ}^i becomes (weakly) finer as σ increases. The manager's information in a given period is the *join* (ie, the coarsest refinement) of all the partitions he chooses. With $\theta = (\theta_1, \dots, \theta_N) \in \mathbb{R}_+^N$ as the information control state, the information constraint is given by

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

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.⁹ 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

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

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 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.

-
- (8) 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.
 - (9) 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.

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).¹⁰ 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 yesterday, then the cost of learning P today is $c(P | Q) = (1 - b)H_\mu(P) + bH_\mu(P | Q)$, where, given a probability μ over S , $H_\mu(P)$ is the entropy of P and $H_\mu(P | Q)$ is the relative entropy of P with respect to Q . Note that $H_\mu(P | P) = 0$ 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.¹¹

Example 2.1 and the second part of Example 2.2 capture, in two alternative ways, the notion of expertise in processing information, that is, a complementarity between information processed in different periods. To see how expertise can help explain a preference pattern that is difficult to explain in the absence of dynamic information constraints, suppose DM has become familiar with a certain set of alternatives and has gained expertise in learning the specific information needed to optimally choose among them. For example, professionals who debate a career change or investors who decide whether or not to enter new markets may find it more difficult to make decisions in the face of new and unfamiliar alternatives, relative to making more routine choices. 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 (see also Section 4.3).¹²

The next two examples can be used to capture *fatigue* in learning, where paying attention in the current period diminishes the ability to pay further attention.

Example 2.3. DM cannot acquire information in two consecutive periods. If he has learned a non-trivial partition of S in the previous period, he cannot afford to learn anything (ie, he can

(10) 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 .

(11) If the prior μ depends on the state (as it does in our model), then entropy, too, will depend on the state. ICP’s can accommodate state dependence.

(12) It is sometimes 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) for a discussion. In that case, 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.

only learn the trivial partition of S) in the current one.¹³ In this case, we may set $\Theta = \{0, 1\}$, $\theta_0 \in \Theta$,

$$\Gamma(\theta) := \begin{cases} \{S\} & \text{if } \theta = 0 \\ \mathcal{P} & \text{if } \theta = 1 \end{cases} \quad \text{and} \quad \tau(P, \theta, s) := \begin{cases} 0 & \text{if } P \neq \{S\}, \theta = 1 \\ 1 & \text{if } P = \{S\} \end{cases}$$

where \mathcal{P} is the collection of all partitions of S .

Example 2.4 (Resource exhaustion). DM is endowed with an initial attention stock K_0 , which he draws down every time he chooses to learn. The amount of attention stock drawn is equal to the cost of choosing a partition. Formally, this corresponds to the ICP in Example 2.2 with parameters $(K_0, \kappa = 0, c, \beta = 1)$.

Conceptually, this type of ICP is reminiscent of the ‘willpower depletion’ model of Ozdenoren, Salant, and Silverman (2012), in which DM is initially endowed with a willpower stock and depletes his willpower whenever he limits his rate of consumption. Consider a simple search problem, where in each period an unemployed worker draws a wage from an iid distribution and needs to decide whether to accept the offer (and work forever at the accepted wage) or to keep searching. Unlike the fixed reservation wage prediction of the standard model, it is easy to show that our model as specified in Example 2.4 will generate a reservation wage that decreases over time, because the expected value of continuing the search decreases as the information constraint tightens due to search-fatigue.

We note that *every* ICP has a finite horizon truncation (see Definition 3.3), 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. (See also the discussion following Proposition 3.4.) The information plans permitted by the finite horizon truncation of an ICP are easily described without any reference to the control state and its transitions, by simply listing the available partitions for every history of information choices and realizations of s , as in the following example.

Example 2.5 (Finite Horizon). The two period truncations of ICPs corresponding to the panels of Figure 2 do not depend on past realizations of s . In the right panel, DM must choose at the outset whether to learn partition P or Q for two successive periods, where P and Q are not ordered in terms of fineness.¹⁴ In the left panel, DM can still commit to either P or Q for both periods, but can alternatively postpone the choice of partition until the second period — at the cost of not learning anything (ie, learning the trivial partition $\{S\}$) in the first period. This option may be valuable if DM does not know at the outset which menu of alternatives he will face in the second period, and hence which information would be most useful.

(13) This example suggests that periods in which individuals pay careful attention are usually followed by periods in which they should rest. In addition to the cognitive interpretation, acquiring information may consume time or physical resources and thus crowd out the completion of other essential tasks; those tasks then have to be performed in consecutive periods, when they, in turn, crowd out further acquisition of information.

(14) That is, it is not the case that every cell in one of them is a union of some cells in the other.

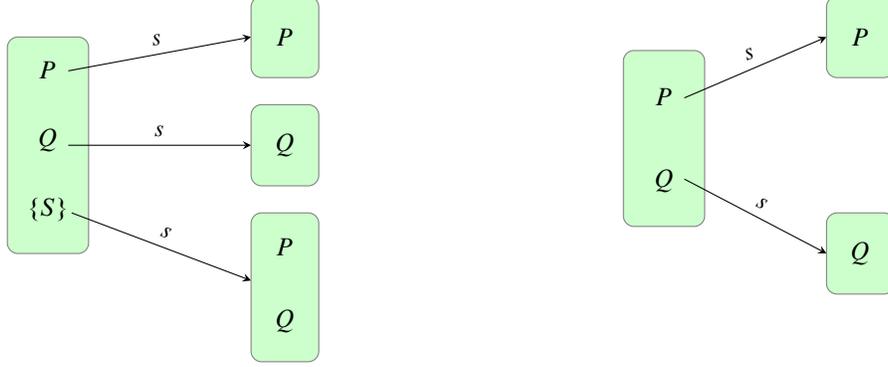


Figure 2: Two-period ICPS

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 $\mathcal{F}(Y)$ the set of acts that map each $s \in S$ to an element of Y , and by $\mathcal{K}(Y)$ 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 $x_1 \in X_1 := \mathcal{K}(\mathcal{F}(\Delta(C)))$. 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 $X_2 := \mathcal{K}(\mathcal{F}(\Delta(C \times X_1)))$, 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 $X_t := \mathcal{K}(\mathcal{F}(\Delta(C \times X_{t-1})))$.

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 $\mathcal{K}(\mathcal{F}(\Delta(C \times X)))$.¹⁵ Note that *both* the current and the continuation menus are now 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. The space L of all consumption streams can be written recursively as $L \simeq \mathcal{F}(\Delta(C \times L))$. 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 ranking \succsim over menus, which is a binary relation on X , and denote its restriction to L by $\succsim|_L$.

The space X of menus, which embodies the descriptive approach of Kreps and Porteus (1978), 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,

(15) Formally, the homeomorphism is written as $X \simeq \mathcal{K}(\mathcal{F}(\Delta(C \times X)))$. See Appendix A.2 for details.

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 choice problem, DM chooses a partition in every period. Let \mathcal{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 $\mathcal{M} = (\Theta, \Gamma, \tau, \theta_0)$, where Θ is a set of *control states*; $\Gamma : \Theta \rightarrow 2^{\mathcal{P}} \setminus \emptyset$ is a set of *feasible partitions* in a given control state θ ; $\tau : \mathcal{P} \times \Theta \times S \rightarrow \Theta$ is a *transition operator* that 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 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-trivial (ie, non-constant), and let $\delta \in (0, 1)$ be a discount factor. Let Π be a fully connected transition operator¹⁶ for a Markov process on S , where $\Pi(s, s') =: \pi_s(s')$ is the probability of transitioning from state s to state s' . Let $0 \notin S$ be an auxiliary state, and denote by π_0 the unique invariant measure of Π .

We consider the following utility representation of \succsim on the space X of menus.

Definition 3.1. A preference \succsim on X has an *ICP representation* $((u_s)_{s \in S}, \delta, \Pi, \mathcal{M})$ if the function $V(\cdot, \theta_0, 0) : X \rightarrow \mathbb{R}$ represents \succsim , where $V : X \times \Theta \times (S \cup \{0\}) \rightarrow \mathbb{R}$ satisfies

$$V(x, \theta, s) =$$

$$[\text{Val}] \quad \max_{P \in \Gamma(\theta)} \sum_{I \in \mathcal{P}} \left[\max_{f \in X} \sum_{s' \in I} \mathbf{E}^{f(s')} \left[u_{s'}(c) + \delta V(y, \tau(P, \theta, s'), s') \right] \pi_s(s' | I) \right] \pi_s(I)$$

In the representation above, for each $s' \in S$, $f(s') \in \Delta(C \times X)$ is a probability measure over $C \times X$ (with the Borel σ -algebra), so that $\mathbf{E}^{f(s')}$ is the expectation with respect to this probability measure.¹⁷

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 ICP representation $((u_s)_{s \in S}, \delta, \Pi, \mathcal{M})$ induces a unique function $V : X \times \Theta \times S \cup \{0\} \rightarrow \mathbb{R}$ that is continuous on X and satisfies [Val]. Moreover, an optimal dynamic information plan exists.

(16) The transition operator Π is *fully connected* if $\Pi(s, s') > 0$ for all $s, s' \in S$.

(17) 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 MDP state θ is controlled by DM and is not observed by the analyst.

A proof is in Appendix A.4.

3.3. Remarks

We now discuss three restrictions that the ICP representation in Definition 3.1 imposes on learning processes, and comment on the behavioral patterns to which these restrictions correspond. Our fourth and final remark points out a limitation of separable models that our approach can address.

First, the realized state of nature, s , is observed by DM at the end of each period (and hence appears as a state variable in the value function [Val]). This rules out an explorative motive for consumption choice (for instance as in a bandit problem), so that the choice of an act has no information value, just as information choice has no direct consumption value; the two interact only through the instrumental value of information.¹⁸ The restriction is implied by the continuity of preferences: If DM does not observe the realized state, he will be willing to pay a premium for acts whose outcomes differ across all states and thus fully reveal the state. Because every open neighborhood in $\mathcal{F}(\Delta(C \times X))$ contains such ‘fully revealing acts’, this would generate violations of continuity that may be implausible in many applications,¹⁹ and that would make the model less tractable.

That being said, even if the state was not observed at the end of each period, but had to be inferred from the outcome of the chosen act, our model would accurately describe behavior on the dense subset of our domain that only includes fully revealing acts; and observing behavior on this subdomain would be sufficient for our identification result. It is then feasible, albeit more cumbersome, to extend our model to the full domain not via continuity (as we currently do), but by keeping track of the effect of the choice of acts on beliefs. Such an extension of the model does not introduce any additional parameters, so that our identification result applies to it as well.

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

(18) To our knowledge, Steiner, Stewart, and Matějka (2017) are the first to 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 cost of a signal at the beginning of a period depends only on its information content (the reduction in entropy given all past signal realizations), but not directly on past information choices, as is possible in the presence of expertise or fatigue in our model. In particular, in the special case of their model where, just like in our model, the state is learned for free at the end of each period, information is valuable only in the period where it is acquired, and the information cost or constraint is history independent and invariant.

(19) For instance, DM might be able to place side bets on fully revealing acts with arbitrarily small stakes, effectively allowing him to learn the state at an arbitrarily small cost at the end of each 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 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 a 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 with 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 and well understood way to achieve it. In particular, we extend the axiomatic literature on subjective learning (Dillenberger et al. (2014) and de Oliveira et al. (2017)) to show that the behavioral implications of deterministic signals (see Section 5.1) are appropriate extensions of known versions of strategic rationality, as in Kreps (1979).

We conclude by pointing out that if the length of a time period in a certain application reflects an aggregation, which is meant to provide an approximation of decisions that are actually made at a higher frequency or rate,²⁰ then, even if a history-independent (or *time-separable*) information constraint is appropriate for the coarser time scale, there may not be a time-separable information constraint for a finer time scale that replicates it. To illustrate, let $S = \{0, 1\} \times \{0, 1\}$ and consider the separable information constraint where DM can learn, in each period, either the first or the second component in S , but not both. Now suppose there is a new data set, where a time period is split in half, and one component realizes every half-period, that is, $S = \{0, 1\}$. In this case, there are only two options to specify a time-separable partitional information constraint: for each half period, DM can either learn the state, or can learn nothing. Neither option captures the original constraint, whereas the non-separable ICP where DM can learn the state in even half-periods only if he did not learn it in odd ones does replicate the original one.²¹ As this example shows, even if information constraints are not separable at a fine enough time scale, it may look as if they are separable at a coarser scale. In contrast to the model with separable constraints, our model is robust to making periods shorter in that case, because ICPS allow for the interdependence of feasible information choices.

(20) For instance, price-setting data at the firm level is often reported monthly or weekly, even though firms can usually change prices more frequently.

(21) The difference between the separable and non-separable environment is not only an artefact of discreteness coming from partitions. Even with noisy channels and an entropy based per-period capacity to process information, as used by Maćkowiak and Wiederholdt (2009), splitting a capacity equally between two half-periods is not equivalent to a model with a full capacity.

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.

We call any ICP that affords only the coarsest partition after t periods a t -period ICP. We first formalize the notion of a finite horizon truncation of an arbitrary ICP.

Definition 3.3. Let $\mathcal{M} = (\Theta, \Gamma, \tau, \theta_0)$ be an ICP. For each $t \geq 1$, its finite horizon truncation, denoted by $\mathcal{M}_{(t)} := (\Theta_{(t)}, \Gamma_{(t)}, \tau_{(t)}, (\theta_0, t))$, is defined as follows: $\Theta_{(t)} := \Theta \times \{0, \dots, t\}$,

$$\Gamma_{(t)}(\theta, j) := \begin{cases} \Gamma(\theta) & \text{if } j \geq 1 \\ \{S\} & \text{if } j = 0 \end{cases} \quad \text{and} \quad \tau_{(t)}(P, (\theta, j), s) := \begin{cases} (\tau(P, \theta, s), j - 1) & \text{if } j > 1 \\ (\theta_0, 0) & \text{if } j = 0, 1 \end{cases}$$

Thus, the t -period ICP $\mathcal{M}_{(t)}$ mimics \mathcal{M} for t periods after which only the coarsest partition is feasible.

A natural way to compare partitions is in terms of fineness, which coincides with Blackwell's comparison of informativeness. To extend this idea to ICPS, first consider two one-period ICPS \mathcal{M} and \mathcal{M}' . Notice that as far as dynamic information plans are concerned, all that matters are the partitions each ICP renders feasible. This suggests the following order on one-period ICPS: \mathcal{M} (*one-period*) *Blackwell dominates* \mathcal{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-period ICPS. \mathcal{M} (*two-period*) *Blackwell dominates* \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) 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 \mathcal{M}' , there is another plan in \mathcal{M} that is more informative in every period and state.

To extend our construction to more than two periods, we note that requirement (ii) above amounts to the one-period continuation ICP $(\Theta, \tau(P, \theta_0, s), \Gamma, \tau)$ being more informative than $(\Theta', \tau'(P', \theta'_0, s), \Gamma', \tau')$. In a similar fashion, we then inductively define an order extending Blackwell dominance to all t -period ICPS, whereby one t -period ICP (*t-period*) *Blackwell dominates* another if for each information plan from the second, there is another plan from the first that is more informative in the first period and, for all $s \in S$, leads to a more informative $(t - 1)$ -period plan starting in the second period.

To illustrate, reconsider Example 2.5. The ICP that corresponds to the constraint in the left panel of Figure 2 (two-period) Blackwell dominates the ICP corresponding to the one on the right panel, but not vice versa. This is because every dynamic information plan available on the right is also feasible on the left, while only the constraint on the left allows the following information 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 .*

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 \mathcal{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 $\mathcal{M}_{(1)}^a$ (the one-period truncation of \mathcal{M}^a) Blackwell dominates $\mathcal{M}_{(1)}^b$ if, and only if, $c^a \leq c^b$ (ie, $c^a(P) \leq c^b(P)$ for all $P \in \mathcal{P}$). Similarly, $\mathcal{M}_{(t)}^a$ (t -period) Blackwell dominates $\mathcal{M}_{(t)}^b$ if, and only if, $c^a \leq c^b$.

Ordering arbitrary ICPS is more delicate, because unlike finite horizon ICPS, 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 3.4. There exists a largest order that satisfies the following: For all $\mathcal{M}, \mathcal{M}^\dagger \in \mathbf{M}$, \mathcal{M} dominates \mathcal{M}^\dagger if for all $P^\dagger \in \Gamma^\dagger(\theta_0^\dagger)$ 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 $(\Theta^\dagger, \Gamma^\dagger, \tau^\dagger, \tau^\dagger(P^\dagger, \theta_0^\dagger, s))$ for all $s \in S$.

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

Proposition 3.4 follows from Propositions A.6 and A.10 in Appendix A.5. In particular, it relies on a metrization of \mathbf{M} which implies that (i) \mathcal{M} dynamically Blackwell dominates \mathcal{M}' if, and only if, $\mathcal{M}_{(t)}$ (t -period) Blackwell dominates $\mathcal{M}'_{(t)}$ for all $t \geq 1$, and (ii) $\mathcal{M}_{(t)}$ converges to \mathcal{M} as $t \rightarrow \infty$.

We note that there are other ways to define dynamic extensions of the static Blackwell order; see, for instance, Greenshtein (1996) and de Oliveira (2016), 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 $((u_s), \delta, \Pi, \mathcal{M})$ 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 \mathcal{M} is unique up to dynamic Blackwell equivalence.²²

The formal proof is in Appendix B. On the subdomain L , V satisfies Independence since it is independent of \mathcal{M} . Indeed, V is then completely characterized by the parameters $((u_s), \delta, \Pi)$. Krishna and Sadowski (2014) 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

(22) In other words, for any additional representation of \succsim with parameters $((u_s^\dagger), \delta^\dagger, \Pi^\dagger, \mathcal{M}^\dagger)$, 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 \mathcal{M} and \mathcal{M}^\dagger dynamically Blackwell dominate each other.

the ICP \mathcal{M} . In Section 4.1 we discuss the main ideas behind the identification strategy for finitely many periods; the extension to infinite horizon is similar, but involves technical issues that we discuss in detail in Appendix B.

Notice that our model is a Markov decision process for DM with state (x, θ, s) , where x and s can be verified by the analyst, while θ cannot. Actions consist of the choice of a partition followed by the choice of an act that depends on the information received. While the choice of act can be verified by the analyst, the choice of partition and subsequent information received cannot. Finally, the analyst knows how x transitions to y when some $f \in x$ is chosen, but the transitions of θ and s are unknown. 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 dynamic environment as does standard 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 $x \succsim \ell$.²³ The comparison in the definition implies that \succsim and \succsim^\dagger have the same ranking over consumption streams in L .²⁴ While any consumption stream requires 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 $((u_s), \delta, \Pi, \mathcal{M})$ and $((u_s^\dagger), \delta^\dagger, \Pi^\dagger, \mathcal{M}^\dagger)$ 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 \mathcal{M} dynamically Blackwell dominates \mathcal{M}^\dagger .*

A proof is in Appendix B. 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 when we sketch the proof of Theorem 1 in Section 4.1. A useful corollary of Theorem 2 is the following characterization of the dynamic Blackwell order: \mathcal{M} dynamically Blackwell dominates \mathcal{M}^\dagger if, and only if, every discounted, expected utility maximizer prefers to have the ICP \mathcal{M} instead of

(23) 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).

(24) 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 .

\mathcal{M}^\dagger regardless of the menu in question.²⁵

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 $u_s = u_{s'}$ for all $s, s' \in S$), and (iii) u is increasing with $u(0) < u(1)$. Rather than providing a general, more abstract intuition, we will base our discussion on the ICPs from Example 2.5, which allow non-trivial information acquisition only in the first two periods; the same ideas extend to any finite horizon truncation, and a limiting argument extends them to arbitrary ICPs. In the sequel, let \mathcal{M} and \mathcal{M}' be ICPs that correspond the left and right panels of Figure 2, respectively. Recall from the discussion in Section 3.4 that \mathcal{M} dynamically Blackwell dominates \mathcal{M}' , but not the other way around.

Suppose the analyst believes DM's preferences are represented by either \mathcal{M} or \mathcal{M}' . How can she verify it is one of them, say \mathcal{M} , and not the other? This is precisely the identification question. We settle it by constructing a menu x such that for any Π and δ , and any $(u_s)_{s \in S}$ as above, $V(x; \mathcal{M}) = V(\mathbf{1}; \mathcal{M}) = V(\mathbf{1})$, and $V(x; \mathcal{M}') < V(\mathbf{1})$, where $\mathbf{1}$ (respectively, $\mathbf{0}$) is the consumption stream that delivers 1 (respectively, 0) in every state at every period,²⁶ and $V(x, \mathcal{M})$ is the value of the menu x under the ICP \mathcal{M} . In that case, we say that the menu x separates \mathcal{M} and \mathcal{M}' . (Holding fixed the ICP \mathcal{M} , we also denote the value of menu x conditional on the choice of partition P by $V(x; P)$.)

It is clear that for x , the optimal initial choice of partition must be $\{S\}$ (ie, the trivial partition), as otherwise $V(x; \mathcal{M}) = V(x; \mathcal{M}')$, because all the other information plans in \mathcal{M} are also available under \mathcal{M}' . We therefore denote that menu, which we now construct explicitly by $x(\{S\}, \mathcal{M})$. For any $J \subset S$, consider the *one-period act* $f_{1,J}$ defined as follows:

$$[4.1] \quad f_{1,J}(s) = \begin{cases} (1, \mathbf{1}) & \text{if } s \in J \\ (0, \mathbf{0}) & \text{if } s \notin J \end{cases}$$

Then, for a partition $R \in \{P, Q\}$, define the *one-period menu* $x_1(R) = \{f_{1,J} : J \in R\}$. Intuitively, $x_1(P)$ requires DM to bet on one of the cells in P . Notice that upon choosing partition $\{S\}$ in the first period, there is a choice between P and Q to be made in the second period. We define the menu $x_2(\{S\}, \mathcal{M})$ as consisting of the act $f_{2,\{S\}}$ (in this case there happens to be only one such act, because the partition $\{S\}$ contains only one cell):

$$f_{2,\{S\}}(s) = (1, \text{Unif}\{x_1(P), x_1(Q)\})$$

(25) This result thus generalizes the seminal characterization of the standard Blackwell order, according to which P is finer than Q if, and only if, every decision maker prefers P to Q regardless of the (static) act in question.

(26) That $V(\mathbf{1}; \mathcal{M}) = V(\mathbf{1})$ follows because the value of consumption stream $\mathbf{1}$ does not depend on information choice, and hence $V(\mathbf{1}; \mathcal{M})$ is independent of \mathcal{M} .

for all $s \in S$, where $\text{Unif}\{y_1, \dots, y_n\}$ is the *uniform lottery* over y_1, \dots, y_n .²⁷ That is, $x_2(\{S\}, \mathcal{M})$ does not require any choice in the initial period, but DM will be asked to bet on either one of the cells in P or on one of the cells in Q , each with positive probability.

To see that $x(\{S\}, \mathcal{M}) := x_2(\{S\}, \mathcal{M})$ separates \mathcal{M} and \mathcal{M}' , notice that with a choice of P or Q in the first period (under \mathcal{M}'), DM is left with the same partition in the second period, while with probability $\frac{1}{2}$ he receives the menu $x_1(R)$, where R is different from his first period information choice. Then $V(x_2(\{S\}, \mathcal{M}); \mathcal{M}') < V(\mathbf{1})$, which establishes the separation claim. Under \mathcal{M} , initially choosing the trivial partition $\{S\}$ is valuable precisely because it allows DM to choose a partition for the second period *after* he knows the second-period menu he is facing, which is drawn randomly from $\{x_1(P), x_1(Q)\}$.

Of course, to fully identify \mathcal{M} , we also need to separate it from other ICPS. As an example, consider the ICP that allows all the information plans that are feasible under \mathcal{M} , except for choosing P twice in a row. A different menu would be needed for that separation. In fact, we can construct three menus, $x(P, \mathcal{M})$, $x(Q, \mathcal{M})$ and $x(\{S\}, \mathcal{M})$, one for each of the initially available partitions in \mathcal{M} , such that facing the ICP \mathcal{M} , and for all $R \in \{P, Q, \{S\}\}$, DM can achieve the same utility from $x(R, \mathcal{M})$ as from the consumption stream $\mathbf{1}$, and *no other* ICP can do strictly better in utility terms. Formally, $V(x(R, \mathcal{M}); \mathcal{M}) = V(\mathbf{1})$, and $V(x(R, \mathcal{M}); \mathcal{M}'') \leq V(\mathbf{1})$ for any ICP \mathcal{M}'' and for $R \in \{P, Q, \{S\}\}$. Further, $V(x(R, \mathcal{M}); \mathcal{M}'') = V(\mathbf{1})$ holds for all $R \in \{P, Q, \{S\}\}$ if, and only if, \mathcal{M}'' dynamically Blackwell dominates \mathcal{M} . We call these requirements *strong alignment* of the collection of menus $\{x(P, \mathcal{M}), x(Q, \mathcal{M}), x(\{S\}, \mathcal{M})\}$ with \mathcal{M} .

Immediately, there is then $R \in \{P, Q, \{S\}\}$ such that $x(R, \mathcal{M})$ separates \mathcal{M} from any ICP \mathcal{M}'' that does not (weakly) dominate it, for example \mathcal{M}' . If, instead, \mathcal{M}'' strictly dominates \mathcal{M} , then there is an analogously constructed collection of ICPS that is strongly aligned with \mathcal{M}'' , so that one of its members provides separation between the two. Thus, the ICP \mathcal{M} is identified up to dynamic Blackwell equivalence.

When constructing $x(P, \mathcal{M})$, note that initially choosing P also determines the partition to be P in the second period. Using the notation above, define for any $J \in P$ the *two-period act* $f_{2,J}$ as:

$$f_{2,J}(s) = \begin{cases} (1, x_1(P)) & \text{if } s \in J \\ (0, \mathbf{0}) & \text{if } s \notin J \end{cases}$$

The *two-period* menu we are after is then $x_2(P, \mathcal{M}) := \{f_{2,J} : J \in P\}$. Intuitively, $x_2(P, \mathcal{M})$ requires DM to bet on one of the cells in P in each of the two periods. By construction, $V(x_2(P, \mathcal{M}); \mathcal{M}) = u(1) + \delta V(x_1(P); P) = u(1) + \delta V(\mathbf{1})$, which implies that no other information plan can get a strictly higher utility. To get this highest utility with the ICP \mathcal{M} , DM must choose P in the first period, as both Q and $\{S\}$ are not finer than P , so that $V(x_1(P), R) < V(x_1(P), P)$ for $R \in \{Q, \{S\}\}$. Since choosing P twice in a row is also feasible

(27) The uniform distribution is not essential; our argument only requires the lottery to have full support on $\{x_1(P), x_1(Q)\}$.

under \mathcal{M}' , the menu $x(P, \mathcal{M}) := x_2(P, \mathcal{M})$ does not separate \mathcal{M} and \mathcal{M}' , but it does separate \mathcal{M} from any ICP where this is not a feasible information strategy. The menu $x(Q, \mathcal{M})$ is constructed analogously to $x(P, \mathcal{M})$.

We now use the above construction to 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 (2016). We again confine our attention in the definition to two-period ICPS.

Definition 4.1. Fix a two-period ICP \mathcal{M} . A (random) sequence of partitions $(P_0, (P_{1,s}))$ is *feasible in \mathcal{M}* if $P_0 \in \Gamma(\theta_0)$ and $P_{1,s} \in \Gamma(\tau(\theta_0, P_0, s))$. Let \mathcal{M}' be another two-period ICP. Then, \mathcal{M} *sequentially Blackwell dominates \mathcal{M}'* if for any sequence of partitions $(Q_0, (Q_{1,s}))$ feasible in \mathcal{M}' , there is another sequence $(P_0, (P_{1,s}))$ feasible in \mathcal{M} such that (i) P_0 is finer than Q_0 , and (ii) $P_{1,s}$ is finer than $Q_{1,s}$ for all $s \in S$.

As above, let \mathcal{M} and \mathcal{M}' be two ICPS that correspond to the left and right panels in Figure 2 respectively. It is easy to see that both \mathcal{M} and \mathcal{M}' sequentially Blackwell dominate each other: \mathcal{M} contains more sequences of partitions than \mathcal{M}' , but any of the additional sequences in \mathcal{M} starts with $\{S\}$ in the first period and is dominated by some sequence feasible in \mathcal{M}' . To understand the discrepancy between the comparative notion in Definition 4.1 and our notion of dynamic Blackwell dominance, recall that an ICP gives DM control over the flow of information. This can be valuable, because it allows him to adjust the information flow to the resolution of uncertainty about the decisions he will face in the future. For example, in the menu $x_2(\{S\}, \mathcal{M})$ above, DM learns whether he faces $x_1(P)$ or $x_1(Q)$ only in the second period, and the ICP \mathcal{M} allows him to postpone the choice of information for the second period until this uncertainty has resolved, while \mathcal{M}' does not. The sequential Blackwell order does not take into account the availability of such dynamic information *plans*, but compares ICPS solely based on the deterministic information *paths* (sequences of partitions) they accommodate, which is inadequate in our dynamic choice setting.²⁸

4.2. Non-Identification of General Markov Decision Processes

There exist negative results in the econometric literature about the identifiability of subjective Markov decision processes. An early and important example appears in Rust (1994). In Rust's model, there is a Markov state space \mathcal{E} , where $\xi \in \mathcal{E}$ fully determines the set of available actions, $D(\xi)$. (For simplicity, we will take this set to be finite.) DM gets instantaneous utility $u(d, \xi)$ from action $d \in D(\xi)$ in state ξ , his discount factor is $\delta \in (0, 1)$, and ξ transitions according to the Markov transition kernel $\mathbf{p}(d\xi' | \xi, d)$, which is controlled by DM's choice of

(28) In other words, the sequential Blackwell order does not condition on the evolution of the menu, as it looks at information plans that can only condition on the evolution of the state s . The dynamic Blackwell order, on the other hand, allows for such dependency by making the dependence a control for DM, and hence our information plans can condition on more than just the evolution of states.

action. This induces the value function (in state ξ when taking the action d)

$$W(\xi, d) := u(d, \xi) + \delta \int \max_{d' \in D(\xi')} W(\xi', d') \mathbf{p}(d\xi' | \xi, d)$$

which in turn induces the choice correspondence $\mathbf{d}(\xi) := \arg \max_{d \in D(\xi)} W(\xi, d)$. The pair (ξ, \mathbf{d}) is observed by the analyst — this is the data in Rust’s model.

Rust shows that identification of state-dependent utilities, the discount factor, and the transition operator is impossible. Indeed, the analyst cannot identify state dependent utilities *even if* the discount factor and the transition operator are known! To show this, Rust constructs a state dependent translation of the utilities that represents the same choice data.²⁹ A second, and perhaps even more fundamental, reason for the failure of identification is that some alternatives may be unavailable or never chosen in state ξ , so that the utility for those alternatives cannot be identified contingent on ξ ; the analyst observes the set $\mathbf{d}(\xi) \subset D(\xi)$, and if $d \notin D(\xi)$ or $d \in D(\xi) \setminus \mathbf{d}(\xi)$, we cannot infer the utility $u(d, \xi)$.

The Markov state in our model is (x, θ, s) .³⁰ In some sense our identification problem is even more ambitious than that in Rust, because we have truly unobservable controls: 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 information choice are both unknown. In order to achieve identification, we leverage restrictions on the Markov decision process that are natural in our context. First, θ and x do not directly enter the utility function, only s does. Second, while the transition operator for the state s is unknown to the analyst, it is not controlled by the decision maker. Finally, the transition of x is controlled by the decision maker through the choice of act f , but the transitions are understood by the analyst (they are specified recursively by the elements of the menu x). Importantly, 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 all possible continuation problems for the same combination of s and θ . This ensures that the second reason for non-identification in Rust mentioned above does not apply in our setting. Moreover, this allows us to identify s -dependent utilities as well as the transition operators for s (by focusing attention on the space L , where information choices do not matter; see Appendix C.2) and θ (see Section 4.1).³¹

(29) Suppose δ and \mathbf{p} are known and consider the utility function $u'(\xi, d) := u(\xi, d) + a(\xi) - \delta \int a(\xi') \mathbf{p}(d\xi' | \xi, d)$, where a is an arbitrary function. It is easy to show (see equation 3.63 in Rust) that the value function is now $W'(\xi, d) := W(\xi, d) + a(\xi)$, which induces the same choice correspondence $\mathbf{d}(\xi)$. Because the transition of ξ is controlled by DM, u and u' are not affinely equivalent, so that the model is not identified.

(30) See the discussion in Section 4 following the statement of Theorem 1.

(31) In our model, $\xi = (x, \theta, s)$ and $u(\xi, \cdot) = u_s(\cdot)$. The transformation in footnote 29 now yields $u'_s := u_s + a(s) - \delta \sum_{s'} \Pi(s, s') a(s')$, which induces another representation, but one that is in the equivalence class identified in Theorem 1, according to which u_s is unique up to the addition of state dependent constants (see footnote 29) that are not controlled by DM.

4.3. Inference from Limited Data

Our identification strategy suggests that useful 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 probabilities used in this randomization are not important, and so can be dictated by the application at hand. Finally, to verify whether DM can follow a particular information plan, the analyst only needs to observe *one* appropriate binary choice — that between the best consumption stream and a choice problem that is uniformly strongly aligned with the plan in question.

To illustrate, suppose there are two periods and consider an investor who must invest in each period a fixed amount in one of two markets, domestic or foreign. On the market where he is active, the investor can either buy (*s*)tocks or (*b*)onds. The relevant state space is $S = \{ss, sb, bs, bb\}$, where, in each period, $ij \in S$ denotes the event that the right decision on the domestic (respectively, foreign) market is to invest in i (respectively, j). Suppose preferences have an ICP representation where the periodic utility payoff from making the right investment choice on the domestic (respectively, foreign) market is 1 (respectively, $1 + \varepsilon$) for each state, while the wrong investment choice in either market yields 0. Suppose $\varepsilon > 0$ so that the foreign market is more profitable. Investing on the domestic market in any period corresponds to choice from the consumption menu³² $D := \{(1, 1, 0, 0), (0, 0, 1, 1)\}$ while investing in the foreign market requires choosing from $F := \{(1 + \varepsilon, 0, 1 + \varepsilon, 0), (0, 1 + \varepsilon, 0, 1 + \varepsilon)\}$.³³ Let DD denote the situation where the investor is active on the domestic market in both periods, FD is the foreign market in the first and the domestic market in the second period, etc.³⁴ Finally, let u_1, u_2 be the consumption (utility) stream that delivers utility u_t in period $t \in \{1, 2\}$.

Suppose the only preference data available to the analyst is that the investor would prefer to be active in the foreign market in both periods, but if for some extraneous reasons (for example foreign markets being closed) he is constrained to invest on the domestic market in either of the two periods, then he strictly prefers to be active in the domestic market in both periods. Formally, $FF \succ DD, DD \succ DF, DD \succ FD$. It is easy to verify that if \succsim admits an ICP representation and if $\varepsilon > 0$ is small enough, then $DD \sim 1, 1$ and $FF \sim 1 + \varepsilon, 1 + \varepsilon$ are implied.³⁵

(32) A consumption menu is a menu that consists of acts without continuation values.

(33) We express outcomes in terms of utilities for convenience, since the goal is to identify the information constraint. As we pointed out before, menu choice is not needed to identify utilities, beliefs, and the discount factor.

(34) These objects can easily be embedded in our domain, by assuming a predetermined continuation stream for the subsequent periods, which we drop from the notation for convenience.

(35) To see this, suppose to the contrary that $1, 1 \succ DD$. Immediately, $P_D P_D$ cannot be feasible (under DM's ICP). But for any partition P that is not finer than P_D , it is easy to verify that $V_1(F; P) > V_1(D; P)$, where V_1 is the value function applied to one period problems. This implies that we cannot have both $DD \succ DF$ and $DD \succ FD$, a contradiction. Hence, $DD \sim 1, 1$.

The following intuition then corresponds exactly to our identification strategy applied to this simple scenario: Let $P_D := \{\{ss, sb\}, \{bs, bb\}\}$ be the coarsest partition in which the acts in menu D are measurable, and analogously $P_F := \{\{ss, bs\}, \{sb, bb\}\}$. $FF \sim 1 + \varepsilon, 1 + \varepsilon$ implies that the investor has an information strategy that allows him to learn at least P_F in both periods, and we say that $P_F P_F$ is feasible (under DM's ICP). Similarly, $DD \sim 1, 1$ implies that $P_D P_D$ is feasible. However, the investor cannot learn (as much or more than) P_D in the first period and P_F in the second, that is $P_D P_F$ is not feasible and, similarly, $P_F P_D$ cannot be feasible as otherwise the representation would imply that $DF \succ DD$, or $FD \succ DD$, respectively. In words, the limited set of observations *implies* that the investor must rely on expertise in order to be able to invest correctly in both periods.

5. Behavioral Characterization

In this section we provide an axiomatic foundation for our model. The main innovation behind our representation theorem lies in the recursive application of our axioms. Therefore, we focus our exposition on this aspect, after describing the axioms, but without burdening the reader with additional notation. Formal statements of all axioms can be found in Appendix C.

In our dynamic setting, the relevant history consists of two disjoint components, namely the *public* history, generated by the observable states s , and the *private* history, generated by DM's unobservable (by the analyst) choices of information. The main difficulty in imposing our axioms is precisely that we want them to hold after histories that contain both public and private components. Our recursive formulation then hinges on identifying those menus where DM has a *unique* subjectively optimal information choice, as for those menus small enough perturbations of the available continuation problems will not change the optimal information, and hence the private history. DM's preferences over continuation problems contingent on the full history are then reflected in his preferences over those perturbations.

All but one of our axioms (namely, Axiom 5) are *static* in the sense that they do not rely on the recursive structure of our domain; they simply restrict preferences on $\mathcal{K}(\mathcal{F}(\Delta(C \times X)))$, ignoring the fact that X is itself again the domain of our preferences (that is, $\mathcal{K}(\mathcal{F}(\Delta(C \times X)))$ is homeomorphic to X). Indeed, our Axioms applied only to initial preferences give rise to a natural, *static* representation, as in display [♣] in Section 5.3. It is then the recursive application of the axioms that exploits the structure of X in a novel way.

Next, suppose $1 + \varepsilon, 1 + \varepsilon \succ FF$. Then $P_F P_F$ is also not feasible. In that case investment must be suboptimal in at least one state in at least one period. Let $\underline{\pi} > 0$ be a lower bound on the state probability across states, periods and state histories, according to Π . Let δ be the discount factor. Then $V(FF) \leq (1 + \varepsilon) + \delta(1 - \underline{\pi})(1 + \varepsilon)$, and hence $V(FF) < V(DD) = 1 + \delta$ as long as $\varepsilon < \frac{\pi}{1 + \delta(1 - \underline{\pi})}$. In that case $DD \succ FF$, a contradiction. Hence, $FF \sim 1 + \varepsilon, 1 + \varepsilon$ for $\varepsilon < \frac{\pi}{1 + \delta(1 - \underline{\pi})}$.

5.1. Versions of Familiar Static Axioms

Axiom 1 summarizes a number of relatively standard properties in the menu-choice literature: the relation \succsim is complete and transitive, Lipschitz continuous, Monotone with respect to set inclusion (bigger menus are better), and satisfies Aversion to Randomization (lower contour sets are convex).

The motivation for our other static axioms (Axioms 2–4) is based on the type of information choice process we envision, where DM is constrained in his choice of partition and takes into account that this choice will also determine a state-dependent utility function over X for next period. We now discuss the extent to which the familiar properties of Temporal Separability, Strategic Rationality and Independence are satisfied when the analyst is not able to observe information choice and hence cannot condition DM’s behavior on it.

Temporal Separability. DM must choose what to learn before picking an act, and this choice may affect the continuation constraint and hence the value of possible continuation problems. We assume, however, that the effect of information choice on the continuation value of an act depends only the marginal distribution over continuation problems the act induces in each state. That is, for any given finite menu x , DM’s optimal learning will not change when substituting act $f \in x$ with g as long as both induce the same marginal distributions over C and X in each state s . *State-Contingent Indifference to Correlation* (Axiom 2) posits that the value of the menu is unchanged under such a substitution.³⁶

Strategic Rationality. Suppose that DM is offered the opportunity to replace the outcome in a certain state of some act in the menu with another. Clearly, no replacement should be worse than a replacement with the worst consumption stream, ℓ_* . This is captured by part (a) of *Indifference to Incentivized Contingent Commitment* (Axiom 3). In general, DM’s attitude towards such replacements may depend on his initial information choice, which is subjective, unobserved, and menu-dependent. Part (b) investigates the conditions under which DM is actually *indifferent* to replacing $f(s)$ with ℓ_* . The natural inference is that DM does not expect to choose act f when the state is s .

Recall that DM must choose a partition of S . Because partitions generate deterministic signals (each state is identified with only one cell of the partition), DM’s choice of partition determines which act he will choose from a given menu, contingent on the state. DM should then be willing to commit to this choice.³⁷ That is, there should be a contingent plan that specifies which act DM will choose for each state, such that he is indifferent between the original menu and one where he is penalized (by receiving the worst consumption stream) whenever

(36) Axiom 2 is closely related to Axiom 5 in Krishna and Sadowski (2014), where other related notions of separability are also mentioned. The important difference is that Axiom 2 requires indifference to correlation in *any* choice problem x , rather than just singletons, because different information may be optimal for different choice problems.

(37) Axiom 3 is conceptually related to the Indifference to State Contingent Commitment Axiom introduced in Dillenberger et al. (2014). Both axioms relate partitional learning to a state contingent notion of strategic rationality, as in Kreps (1979).

his choice does not coincide with that plan.³⁸

Independence. The menus x and y are *concordant* if the same set of initial information choices is optimal for both x and y . In general, DM can potentially tailor his information choice to the menus x and y separately, while for the mixed menu $\frac{1}{2}x + \frac{1}{2}y$ he may need to compromise. However, if x and y are concordant — so that the optimal partitions for x and y are the same — then the same partitions are also optimal for the mixed menu. That is, $\frac{1}{2}x + \frac{1}{2}y$ should also be concordant with x . Since DM can choose the same information across any set of concordant menus, we require \succsim restricted to any such set to satisfy Independence. These requirements are the content of *Concordant Independence* (Axiom 4).

The key is to verify concordance of menus x and y from behavior. Towards this end, recall that $x_1(P)$ is a menu that requires DM to bet on the cells of the partition P in the first period and requires no choice after that, where optimal choice in the first period generates the best consumption stream ℓ^* .³⁹ Thus, if $\frac{1}{2}x + \frac{1}{2}x_1(P) \sim \frac{1}{2}x + \frac{1}{2}\ell^*$, it must be the case that some partition at least as fine as P is optimal for x . We, therefore, call x and y *concordant* if for each $P \in \mathcal{P}$ 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, $\frac{1}{2}y + \frac{1}{2}x_1(P) \sim \frac{1}{2}y + \frac{1}{2}\ell^*$.

5.2. Choice over Consumption Streams

Our last assumption is not static, but imposes assumptions only on $\succsim|_L$, the restriction of \succsim to the set of *consumption streams*, L . The subdomain L is special because it includes no consumption choice to be made in the future, which renders information (choice) inconsequential and hence the distinction between public and private histories irrelevant. Following Krishna and Sadowski (2014), Axiom 5 requires $\succsim|_L$ to satisfy standard versions of (vN-M) Independence, History Independence, and Stationarity.

5.3. A Recursively Formulated Representation Theorem

As noted above, the history of state variables that are relevant for DM consists of the public history of the observable states s , and the private 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.

(38) For example, if $S = \{1, 2\}$, then the menu $\{f, g\}$ must be indifferent to one of the menus $\{f, \ell_*\}$, $\{g, \ell_*\}$, $\{f(1)\ell_*, \ell_*(1)g\}$ or $\{\ell_*(1)f, g(1)\ell_*\}$, where $f'(s)g'$ denotes the act that agrees with f' in state s and with g' otherwise. That is, under partitional learning with only two states, DM either learns the state for sure or learns nothing. The first two cases correspond to DM choosing one of the acts in $\{f, g\}$ unconditionally. In the other two cases, DM expects to learn the state for sure and to choose different acts in different states.

(39) See display [4.1] in Section 4.1 for an example when the prize space is $[0, 1]$ and there are additional conditions placed on utilities.

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 , which implies that these perturbed menus are concordant with 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 [♣] below, this corresponds to one partition generating a higher value than any other available partition).⁴⁰ For such x , if (i) given a state s , DM prefers to *simultaneously* perturb all acts $f \in x$ by mixing the continuation problems they specify in state s with y rather than mixing them with y' , and if (ii) x and the perturbed menus are concordant (verified from behavior as discussed at the end of Section 5.1), 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.

Let Ψ_1 be the space of all binary relations on X that satisfy Axioms 1–5. That is,

$$\Psi_1 := \{ \succsim \text{ on } X : \succsim \text{ satisfies Axioms 1–5} \}$$

Now, define the operator $\mathcal{B} : 2^{\Psi_1} \rightarrow 2^{\Psi_1}$ as follows:

$$[5.1] \quad \mathcal{B}(\Psi) := \left\{ \succsim \in \Psi_1 : \begin{array}{l} \succsim_{(x,s)} \in \Psi \text{ for all } s \in S \text{ and finite } x \in X \\ \text{with a unique optimal partition} \end{array} \right\}$$

for $\Psi \subset \Psi_1$. This allows us to define the set

$$\Psi_2 := \mathcal{B}(\Psi_1)$$

which consists of all preferences that satisfy our (static) axioms with the additional property that continuation preferences (ie, $\succsim_{(x,s)}$ for all finite x with a unique optimal partition) also satisfy our static axioms, by virtue of being in Ψ_1 . It is easy to extend this construction to arbitrary finite horizons by inductively defining $\Psi_{n+1} := \mathcal{B}(\Psi_n)$ for all $n \geq 1$. The construction of Ψ_n for $n \geq 2$ parallels the recursive construction of our domain of choice problems 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 \mathcal{B} in [5.1], which is easily seen to be monotone. It can be shown that $\Psi^* = \lim_{n \rightarrow \infty} \Psi_n = \bigcap_{n \geq 1} \Psi_n$.⁴¹ We say that a *preference relation on X satisfies Axioms 1–5 recursively* if it is in Ψ^* , or equivalently, lies in Ψ_n for all $n \geq 1$.

(40) In terms of \succsim , we focus on x that can be written as a mixture between $x_1(P)$ and some menu x' for which P is optimal and any other optimal partition must be dominated by P , that is, $x = \frac{1}{2}x' + \frac{1}{2}x_1(P) \succsim \frac{1}{2}x' + \frac{1}{2}x_1(Q)$ for all Q , with strict preference if Q is not dominated by P . Because any information choice that is optimal for $x_1(P)$ must dominate P , indeed P must be the unique optimal information choice for x .

(41) This is analogous to the *self-generating set of equilibrium payoffs* in Abreu, Pearce, and Stacchetti (1990), which is also the fixed point of an appropriate operator.

Our representation theorem, Theorem 3, 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) \succsim satisfies Axioms 1–5 recursively.
- (b) There exists an ICP representation of \succsim .

The recursive application of our axioms after histories that account for unobservable information choices is our main innovation. It relies on the operator \mathcal{B} , which is a decision-theoretic version of the dynamic programming operator, so that preferences over menus are *as if* optimal choices of information and consumption will ensue, consistent with the dynamic programming principle. This reduces the characterization of behavior over a dynamic domain to a static characterization of per-period preferences.

As just argued, the recursive application of our axioms requires not only preferences, but also induced preferences over the next period's continuation problems to be in Ψ^* , thereby requiring that preferences over continuation problems two periods ahead again have to be in Ψ^* , and so forth. 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 compatible partition) and violate one of Axioms 1–5.⁴²

Our approach works just as well for finite horizons, as evidenced in the following corollary. Recall that X_n is the space of all n -period dynamic choice problems (see Section 3.1 for a definition), and that $\Psi_n = \mathcal{B}(\Psi_{n-1})$ for all $n > 1$.

Corollary 5.1. *Let $\succsim^{(n)}$ be a binary relation on X_n for some fixed n . Then, the following are equivalent:*

- (a) $\succsim^{(n)} \in \Psi_n$.
- (b) There exists an ICP representation of \succsim over X_n .

Corollary 5.1 follows immediately from Theorem 3. The requirement that $\succsim^{(n)} \in \Psi_n$ is tantamount to requiring that $\succsim^{(n)}$ satisfies Axioms 1–5, and that the continuation preference $\succsim_{(x,s)}^{(n-1)}$ on X_{n-1} satisfies Axioms 1–5 for all finite x with a unique optimal partition (ie, $\succsim_{(x,s)}^{(n-1)} \in \Psi_{n-1}$), and so on.

The proof of Theorem 3 is quite involved. To focus on the most novel parts of our construction, we take as given the following *static* representation of \succsim (see Theorem 2 in the Supplementary Appendix for a proof of its existence), which is the starting point for our proof in Appendix D:

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

(42) We remark that the (recursive) validity of all our axioms except our continuity requirement can be falsified with finite data.

where $\mathbb{Q} \subset \mathcal{P}$ is a set of partitions of S , the measure $\pi(s | I)$ is the probability of s conditional on the event $I \subset S$, and utilities (v_s) over continuation problems depend only on the partition P . 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 . Correspondingly, we need to apply Axioms 1–5 to \succsim to derive it, but do not need to apply the axioms recursively, ie, do not need to impose further restrictions on continuation preferences.

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 D.1). In Lemma D.7 in Appendix D, 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 the recursive application of Axioms 1–5 to show that the representation V of \succsim can be made self-generating.

The remainder of our construction in Appendix D.3 has two main components. First, we construct an ICP $\mathcal{M} = (\Theta, \Gamma, \tau, \theta_0)$ 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 1 and 5 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

$$V_L^*(\ell, s) := \sum_r \Pi(s, r) \mathbf{E}^{\ell(r)} [u_r(c) + \delta V_L^*(\ell', r)]$$

for some tuple $((u_s), \delta, \Pi)$ and where $\ell, \ell' \in L$.⁴³ 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 in the self-generating representation of V , all the utilities 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 \mathcal{M} to find the ICP representation, $((u_s), \delta, \Pi, \mathcal{M})$, 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 .

In Appendix E we discuss how to further strengthen our notion of Separability and impose more structure on continuation preferences (beyond just requiring that they satisfy Axioms 1–5) to characterize the special case of the ICP representation where DM faces the same information constraint each period. This case 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, measured by the expected reduction in entropy. We also confirm

(43) This is the *Recursive Anscombe-Aumann* representation in Krishna and Sadowski (2014). See Appendix C.2 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.

that imposing full Independence implies that information is determined by a trivial choice process: it exogenously arrives over time.

6. Related Literature

We now comment on the menu choice literature that shares some of our basic assumptions.

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 vN - M utility functions over prizes (that is, 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 \mathcal{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 \mathcal{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.⁴⁴

\mathcal{DM} controls his information over time. Thus, his preferences will be interdependent across time, which significantly complicates our analysis, especially because we can no longer appeal to the stationarity assumptions of Krishna and Sadowski (2014). To deal with this complication, we observe that preferences over consumption streams, $\succsim|_L$, should 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, similar to the approach in Gilboa and Schmeidler (1989), where preferences over unambiguous acts (lotteries) are used to calibrate ambiguity averse preferences over all acts. Table 1 summarizes the position of our

(44) To be sure, 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, as we do in Section 7.

model with respect to the papers discussed thus far.

<i>Information</i>	<i>Uncertainty about vN-M taste</i>	<i>Uncertainty about state of world</i>
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)	<i>This paper</i>
Dynamic, choice process	★	<i>This paper</i>

Table 1: *Summary of Literature*

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 Axioms 1–4, then they also satisfy them 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, ie, on public histories. In contrast, in our model of endogenous information choice, past menus affect past information choices which, in turn, can affect the current 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.

There are additional works that 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. 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 Matejka and McKay (2015). To our knowledge, there is no counterpart to our recursive analysis of information choice in the random-choice literature.

7. Direct information Costs

ICPs can generate opportunity costs of information acquisition via tighter future constraints. Our main model allows us to focus *entirely* on the behavioral implications of this new type of

dynamic cost. That said, an alternative way to model limitations on information acquisition is via direct information costs, measured in consumption ‘utils’ (see, for example, Ergin and Sarver (2010), Woodford (2012), Caplin and Dean (2015), de Oliveira et al. (2017), and Hébert and Woodford (2016)).⁴⁵ In this section we (i) provide a recursive model based on *intertemporal* costs rather than ICPS, (ii) argue that the heart of our identification strategy is robust to this change,⁴⁶ and (iii) discuss the benefits of dealing with constraints rather than costs for identification of (or inference about) the relevant parameters.

A *cost structure* \mathcal{C} is a tuple $(\Theta, \theta_0, \tau, \rho)$, where Θ , θ_0 , and τ are as before, and $\rho : \mathcal{P} \times \Theta \rightarrow \mathbb{R}_+$ is a *cost function* that determines the cost of learning a partition as a function of the cost state θ .⁴⁷ As with ICPS, after the choice of P and the realization of s , the transition operator τ determines the continuation control state $\theta' = \tau(P, s, \theta)$. We assume that $\rho(\{S\}, \theta) = 0$ for all θ . With timing unchanged from before, we can show (see Section 1.1 of the Supplementary Appendix) that \mathcal{C} induces a unique value function

$$V(x, \theta, s) = \max_{P \in \mathcal{P}} \left[\sum_{J \in P} \left(\max_{f \in x} \sum_{s' \in J} \pi(s' | J) \mathbf{E}^{f(s')} [u_{s'}(c) + \delta V(y, \tau(P, s, \theta), s')] \right) - \rho(P, \theta) \right]$$

We refer to $((u_s), \delta, \Pi, \mathcal{C})$ as an *Information Cost* representation if $V(\cdot, \theta_0, s_0) : X \rightarrow \mathbb{R}$ represents preferences on X .

In order to compare two cost structures, \mathcal{C} and \mathcal{C}' , we must describe how costly they make it to follow a particular dynamic information plan. Recall that a dynamic information plan prescribes a choice of partition as a function of (x, θ, s) . (See Section 3.2 for a definition of dynamic information plans.)

We say that the cost structure \mathcal{C} *dominates* \mathcal{C}' if for any information plan σ' there is another plan σ such that (i) σ is more informative than σ' at every date and state, and (ii) the expected cost of σ under the cost structure \mathcal{C} is less than that of σ' under the cost structure \mathcal{C}' . \mathcal{C} and \mathcal{C}' are *equivalent* if they dominate each other.

Theorem 4. *Let $((u_s), \delta, \Pi, \mathcal{C})$ be an Information Cost representation for \succsim . Then*

- *The functions $(u_s)_{s \in S}$ are unique up to addition of constants and common scaling,*
- *δ and Π are unique,*
- *If the functions $(u_s)_{s \in S}$ are unbounded,⁴⁸ then \mathcal{C} is unique up to equivalence.*

The intuition for the identification result is similar to that described in Section 4.1 for the case of ICPS with the following qualifiers. First, as is apparent from the theorem and

(45) In static settings, information constraints imply that the amount of information chosen is independent of the scaling of the payoffs involved, which stands in sharp contrast to the stake-dependency under costly information acquisition. Because ICPS can generate opportunity costs of information acquisition, choice may be sensitive to the stakes in a given period, thereby reducing the gap between the two models.

(46) A formal treatment can be found in Section 1 of the Supplementary Appendix.

(47) It is easy to see that an ICP is a cost structure where $\rho(\cdot) \in \{0, \infty\}$.

(48) Unbounded utilities need a non-compact domain. We show in Section 1.1 of the Supplementary Appendix that our approach extends to a domain of dynamic choice problems with a σ -compact consumption space.

immediately intuitive, to identify potentially arbitrarily high information costs, arbitrarily high stakes are needed. This requires utilities to be unbounded, and the analyst to elicit those utilities before eliciting the cost structure. In contrast, identification of the ICP relies only on two outcomes that are strictly ranked in some state (recall that we only require *some* u_s to be non-trivial), which is much less demanding.

Second, the expected cost of a particular information plan, and hence our notion of dominance between cost structures, depends on the discount factor δ and the evolution of the payoff relevant state in S according to Π . In contrast, dynamic Blackwell dominance between ICPs is independent of other preference parameters.

Finally, constructing a choice problem that incentivizes a particular information plan σ requires mixing over continuation problems with exactly the right probabilities to ensure that σ is an optimal information strategy. In contrast, the exact probabilities assigned to different continuation problems is irrelevant in the identification argument for the ICP (in Section 4.1). As a consequence, while our general identification strategy is robust to the consideration of information costs instead of constraints, drawing inference about cost structures in applied contexts will be much more involved than the procedure for inference about constraints that we discuss in Section 4.3.

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. Appendix A.4 proves the existence of the value function V that satisfies [Val]. Appendix A.5 formally defines the dynamic Blackwell order and describes some of its properties.

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 . The following definitions may be found in Chapter 11 of Dudley (2002). For a function $\varphi \in \mathbb{R}^Y$, the supremum norm is $\|\varphi\|_\infty := \sup_y |\varphi(y)|$, and the Lipschitz seminorm is defined by $\|\varphi\|_L := \sup_{y \neq y'} |\varphi(y) - \varphi(y')| / d_Y(y, y')$. This allows us to define the *bounded Lipschitz* norm $\|\varphi\|_{BL} := \|\varphi\|_L + \|\varphi\|_\infty$. Then, $BL(Y) := \{\varphi \in \mathbb{R}^Y : \|\varphi\|_{BL} < \infty\}$ is the space of real-valued, bounded, and Lipschitz functions on Y .

Define d_D on $\Delta(Y)$ as $d_D(\alpha, \beta) := \frac{1}{2} \sup \{ |\int \varphi d\alpha - \int \varphi d\beta| : \|\varphi\|_{BL} \leq 1 \}$. This is the *Dudley metric* $\Delta(Y)$. Theorem 11.3.3 in Dudley (2002) says that for separable Y , d_D induces the topology of weak convergence on $\Delta(Y)$. We note that the factor $\frac{1}{2}$ is not standard. We introduce it to ensure that for all $\alpha, \beta \in \Delta(Y)$, $d_D(\alpha, \beta) \leq 1$.

A.2. Recursive Domain

Let $X_1 := \mathcal{K}(\mathcal{F}(\Delta(C)))$. For acts $f^1, g^1 \in \mathcal{F}(\Delta(C))$, define the metric $d^{(1)}$ on $\mathcal{F}(\Delta(C))$ by $d^{(1)}(f^1, g^1) := \max_s d_D(f^1(s), g^1(s))$. For any $f^1 \in \mathcal{F}(\Delta(C))$ 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). Notice that for all acts f^1 and g^1 , $d^{(1)}(f^1, g^1) \leq 1$.

This allows us to define the Hausdorff metric $d_H^{(1)}$ on X_1 as

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

and because the distance of an act from a set is bounded above by 1, it follows that for all $x_1, y_1 \in X_1$, $d_H^{(1)}(x_1, y_1) \leq 1$. Intuitively, X_1 consists of all one-period Anscombe-Aumann (AA) choice problems.

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

For acts $f^n, g^n \in \mathcal{F}(\Delta(C \times X_{n-1}))$, define the distance between them as $d^{(n)}(f^n, g^n) := \max_s d_D(f^n(s), g^n(s))$. As before, we may now define the Hausdorff metric $d_H^{(n)}$ on X_n as

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

which is also bounded above by 1. 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 $\times_{n=1}^{\infty} X_n$ with the product topology. The Tychonoff metric induces this topology and is given as follows: For $x = (x_1, x_2, \dots), y = (y_1, y_2, \dots) \in \times_{n=1}^{\infty} X_n$,

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

It is easy to see that for all $x, y \in \times_{n=1}^{\infty} X_n$, $d(x, y) \leq 1$. Moreover, and this is easy to verify (because it holds for $d_H^{(n)}$ for each n), $d(\frac{1}{2}x + \frac{1}{2}y, y) = \frac{1}{2}d(x, y)$.

The space of choice problems (menus) X is all members of $\times_{n=1}^{\infty} X_n$ 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} .⁴⁹ The space of menus, X , is our domain for choice, and it follows from standard arguments that X is (linearly) homeomorphic to $\mathcal{K}(\mathcal{F}(\Delta(C \times X)))$. We denote this homeomorphism by $X \simeq \mathcal{K}(\mathcal{F}(\Delta(C \times X)))$. In what follows, we shall abuse notation and use d as a metric both on X and on $\mathcal{K}(\mathcal{F}(\Delta(C \times X)))$. It will be clear from the context precisely which space we are interested in, so there should be no cause for confusion.

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 $n \geq 1$.

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

A.3. Canonical Information Choice Processes

Recall that \mathcal{P} is the space of all partitions of S , where a typical partition is P . The partition P is *finer* than the partition Q if every cell in Q is the union of cells in P . Given a measure μ on S , define the *entropy* of a partition P as $H(P) := -\sum_{J \in P} \mu(J) \log \mu(J)$. Then, we can define a metric d on \mathcal{P} as $d(P, Q) := 2H(P \wedge Q) - H(P) - H(Q)$, where $P \wedge Q$ is the coarsest refinement of P and Q . In Section 5 of the Supplementary Appendix, we show that d is indeed a metric. Thus, (\mathcal{P}, d) is a metric space.

For metric spaces X and Y , we denote by $\mathcal{K}_b(X \times Y)$ 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 $x \neq x'$.

Let $\Omega_1 := \mathcal{K}(\mathcal{P})$, and define recursively for $n > 1$, $\Omega_n := \mathcal{K}_b(\mathcal{P} \times \Omega_{n-1}^S)$. Then, we can set $\Omega' := \prod_{n=1}^{\infty} \Omega_n$. 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 $\psi_1 : \mathcal{P} \times \Omega_1^S \rightarrow \mathcal{P}$ be given by $\psi_1(P, \boldsymbol{\omega}_1) = P$, and define $\Psi_1 : \Omega_2 \rightarrow \Omega_1$ as $\Psi_1(\omega_2) := \{\psi_1(P, \boldsymbol{\omega}_1) : (P, \boldsymbol{\omega}_1) \in \omega_2\}$. Now define recursively, for $n > 1$, $\psi_n : \mathcal{P} \times \Omega_n^S \rightarrow \mathcal{P} \times \Omega_{n-1}^S$ as $\psi_n(P, \boldsymbol{\omega}_n) := (P, (\Psi_{n-1}(\omega_{n,s}))_s)$, and $\Psi_n : \Omega_{n+1} \rightarrow \Omega_n$ 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* is

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

that is, the set of ICPS is the space of all consistent elements of Ω' .

Notice that Ω_1 is a compact metric space when endowed with the Hausdorff metric. Then, inductively, $\mathcal{P} \times \Omega_{n-1}^S$ 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.)

Therefore, 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 5. *The set Ω is homeomorphic to $\mathcal{K}_b(\mathcal{P} \times \Omega^S)$.*

We write the homeomorphism as $\Omega \simeq \mathcal{K}_b(\mathcal{P} \times \Omega^S)$. The theorem is not proved, though it can be in a straightforward way, by adapting the arguments in Mariotti, Meier, and Piccione (2005).

As in the case of one and two periods, all that is relevant (in terms of behavior) for the description of the ICPS \mathcal{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 \mathbf{M} 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 \mathcal{M} and \mathcal{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 $\mathcal{M} = (\Theta, \Gamma, \tau, \theta_0)$ and $\mathcal{M}' = (\Theta', \Gamma', \tau', \theta'_0)$ be two ICPS in \mathbf{M} . A choice of $P \in \Gamma(\theta_0)$ and a realisation of state s results in a new ICP $(\Theta, \Gamma, \tau, \tau(\theta_0, P, s))$. To ease the notational burden, we shall denote this new ICP by $\mathcal{M}(\tau(\theta_0, P, s))$. Further abusing notation, $\mathcal{M}(\theta)$ denotes the ICP $(\Theta, \Gamma, \tau, \theta)$ with initial state θ .

Define the function $D : \mathbf{M} \times \mathbf{M} \rightarrow \mathbb{R}$ as follows:

$$D(\mathcal{M}(\theta_0), \mathcal{M}'(\theta'_0)) := \max \left[d_H(\Gamma(\theta_0), \Gamma'(\theta'_0)) \wedge 1, \frac{1}{2} \max_{P \in \Gamma(\theta_0), s \in S} D(\mathcal{M}(\tau(\theta_0, P, s)), \mathcal{M}'(\tau'(\theta'_0, P, s))) \right]$$

[A.1]

The function D captures the *discrepancy* between the ICPS \mathcal{M} and \mathcal{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 $D \in \mathbf{B}(\mathbf{M} \times \mathbf{M})$ that satisfies equation [A.1].

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

$$TD'(\mathcal{M}(\theta_0), \mathcal{M}'(\theta'_0)) := \max \left[d_H(\Gamma(\theta_0), \Gamma'(\theta'_0)) \wedge 1, \frac{1}{2} \max_{P \in \Gamma(\theta_0), s \in S} D'(\mathcal{M}(\tau(\theta_0, P, s)), \mathcal{M}'(\tau'(\theta'_0, P, s))) \right]$$

for all $D' \in \mathbf{B}(\mathbf{M} \times \mathbf{M})$. It is easy to see that T is monotone in the sense that $D_1 \leq D_2$ implies $TD_1 \leq TD_2$. It also satisfies discounting, ie, $T(D + a) \leq TD + \frac{1}{2}a$ for all $a \geq 0$. This implies that T has a unique fixed point in $\mathbf{B}(\mathbf{M} \times \mathbf{M})$, and this fixed point D satisfies [A.1]. \square

We can now define an isomorphism between ICPS. Two ICPS \mathcal{M} and \mathcal{M}' are *indistinguishable* if $D(\mathcal{M}(\theta_0), \mathcal{M}'(\theta'_0)) = 0$. We now have an easy, recursive characterization of indistinguishability.

Lemma A.2. Let $\mathcal{M}, \mathcal{M}' \in \mathbf{M}$. Then, \mathcal{M} is indistinguishable from \mathcal{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 $(\Theta, \Gamma, \tau, \tau(\theta_0, P, s))$ is indistinguishable from the ICP $(\Theta', \Gamma', \tau', \tau'(\theta'_0, P, s))$.

The proof follows immediately from the definition of the discrepancy function D and so is omitted.

The homeomorphism $\Omega \simeq \mathcal{K}_b(\mathcal{P} \times \Omega^S)$ 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, ω') , where $\omega' = (\omega'_s)_{s \in S}$. To see

that every $\omega \in \Omega$ is indeed an ICP, set $\Gamma^*(\omega) = \{P : (P, \omega') \in \omega\}$ and $\tau^*(\omega, P, s) = \omega'_s$ to obtain the ICP $\mathcal{M}_\omega = (\Omega, \Gamma^*, \tau^*, \omega)$ which is indistinguishable from ω . It is easy to see that for $\omega, \omega' \in \Omega$, $\omega \neq \omega'$ implies $D(\omega, \omega') > 0$.

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

- (a) Every $\mathcal{M} \in \mathbf{M}$ is indistinguishable from a unique $\omega_{\mathcal{M}} \in \Omega$.
- (b) Every $\omega \in \Omega$ induces an $\mathcal{M}_\omega \in \mathbf{M}$ that is indistinguishable from ω .

Proof. We first show that (a) implies (b). Towards this end, let $\mathcal{M} = (\Theta, \Gamma, \tau, \theta_0)$ be an ICP. Recall the definition of the space Ω_n and define the maps $\Phi_n : \Theta \rightarrow \Omega_n$ as follows. Let

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

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 is easy to see that the sequence

$$(\omega_1, \omega_2, \dots, \omega_n, \dots) \in \prod_{n \in \mathbb{N}} \Omega_n$$

is *consistent* in the sense described above. Therefore, there exists $\omega \in \Omega$ such that $\omega = (\omega_1, \omega_2, \dots, \omega_n, \dots)$, ie, the ICP \mathcal{M} corresponds to a *canonical* ICP ω .

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

A.4. Value Function

We now prove Proposition 3.2 for the case of canonical ICPs. The extension to the case of general ICPs is straightforward. 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 $T : \mathbf{C}(X \times \Omega \times (S \cup \{0\})) \rightarrow \mathbf{C}(X \times \Omega \times (S \cup \{0\}))$ as follows:

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

Recall that x is compact. It follows from the Theorem of the Maximum (using standard arguments) that T is well defined. 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 DM who is characterized by $((u_s)_{s \in S}, \Pi, \delta, \omega)$, there exists a unique $V \in \mathbf{C}(X \times \Omega \times (S \cup \{0\}))$ that satisfies the functional equation [Val].

The optimal dynamic information plan is merely the mapping $(x, \omega, s') \mapsto (P, \omega') \in \omega$. Because the set of such (P, ω') is finite, it follows that there is a *conserving plan*.⁵⁰ Given that C is bounded and because of discounting, the conserving plan is actually optimal — see Orkin (1974) or Proposition A6.8 of Kreps (2012). \square

A.5. 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 $\hat{\Omega}_0 := \mathcal{K}_b(\mathcal{P} \times \{\hat{\omega}\})$, and inductively define $\hat{\Omega}_{n+1} := \mathcal{K}_b(\mathcal{P} \times \hat{\Omega}_n)$ 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 \succsim_0 on $\hat{\Omega}_0$ as follows: $\omega_0 \succsim_0 \omega'_0$ if for all $(P', \hat{\omega}) \in \omega'_0$, there exists $(P, \hat{\omega}) \in \omega_0$ such that P is finer than P' . This allows us to define inductively, for all $n \geq 1$, an order \succsim_n on $\hat{\Omega}_n$. For all $\omega_n, \omega'_n \in \hat{\Omega}_n$, $\omega_n \succsim_n \omega'_n$ if for all $(P', \omega'_{n-1}) \in \omega'_n$, there exists $(P, \omega_{n-1}) \in \omega_n$ such that (i) P is finer than P' , and (ii) $\omega_{n-1,s} \succsim_{n-1} \omega'_{n-1,s}$ for all $s \in S$.

It is easy to see that \succsim_n is reflexive and transitive for all n . There is a natural sense in which ω_n extends ω_{n-1} , as we show next.

Lemma A.4. For all $n \geq 0$, \succsim_{n+1} extends \succsim_n , ie, $\succsim_{n+1} |_{\hat{\Omega}_n} = \succsim_n$.

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 $\omega_0 \succsim_0 \omega'_0$. Then, for $(P', \hat{\omega}) \in \omega'_0$, there exists $(P, \hat{\omega}) \in \omega_0$ such that P is finer than P' . Moreover, because \succsim_0 is reflexive, $\hat{\omega} \succsim_0 \hat{\omega}$. But this implies $\omega_0 \succsim_1 \omega'_0$. Conversely, let $\omega_0 \succsim_1 \omega'_0$. Then, for all $(P', \hat{\omega}) \in \omega'_0$, there exists $(P, \hat{\omega}) \in \omega_0$ such that (i) P is finer than P' , and (ii) $\hat{\omega} \succsim_0 \hat{\omega}$ for all $s \in S$. But this implies $\omega_0 \succsim_0 \omega'_0$, which proves that $\succsim_{n+1} |_{\hat{\Omega}_n} = \succsim_n$ when $n = 0$.

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

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

(50) A plan (respectively, an action) at some date and state is conserving if it achieves the supremum in Bellman's equation. See, for instance, Kreps (2012).

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

By definition of $\hat{\Omega}$, there is some n such that $\omega, \omega' \in \hat{\Omega}_n$, and by Lemma A.4, the precise choice of this n is irrelevant. This implies \succsim is well defined. We now show that \succsim has a recursive definition as well.

Proposition A.5. For any $\omega, \omega' \in \hat{\Omega}$, the following are equivalent.

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

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

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

The uniqueness of \succsim on $\hat{\Omega}$ follows immediately from the uniqueness of \succsim_n for all $n \geq 0$. □

We can now prove the existence of a recursive order on Ω . (Notice that $\text{cl}(\hat{\Omega}) = \Omega$, where $\text{cl}(\hat{\Omega})$ is the *closure* of $\hat{\Omega}$.) In particular, for all $\omega, \omega' \in \Omega$, we say that ω *dynamically Blackwell dominates* ω' if for all $(P', \tilde{\omega}') \in \omega'$, there exists $(P, \tilde{\omega}) \in \omega$ 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 A.6. The order \succsim on $\hat{\Omega}$ has a unique continuous extension to Ω , also denoted by \succsim . Moreover, on Ω , \succsim is the unique non-trivial and continuous dynamic Blackwell order.

Proof. Because $\Omega = \text{cl}(\hat{\Omega})$, we simply extend \succsim to Ω by taking its closure, namely $\text{cl}(\succsim)$. Abusing notation, this extension is also denoted by \succsim . It is easy to see that \succsim so defined is continuous and non-trivial. That \succsim is a unique dynamic Blackwell order follows immediately from the facts that $\hat{\Omega}$ is dense in Ω , the continuity of \succsim , and Proposition A.5. □

Let $\text{proj}_n : \Omega \rightarrow \hat{\Omega}_n$ be the natural 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 $\hat{\omega}$, ie, the coarsest partition forever. It is easy to see that 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 A.7. For $\omega, \omega' \in \Omega$, $\omega \succsim \omega'$ if, and only if, for all $n \in \mathbb{N}$, $\omega_n \succsim \omega'_n$.

Proof. The ‘only if’ part is straightforward. The ‘if’ part follows from the continuity of \succsim . □

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

Corollary A.8. For all $\omega, \omega' \in \Omega$ and $m \geq 1$, $\omega_m \succeq \omega'_m$ implies $\omega_n \succeq \omega'_n$ for all $1 \leq n \leq m$.

Proof. Notice that $\omega_m, \omega'_m \in \Omega$. Therefore, by Corollary A.7, it follows that for all $n \geq 1$, $\text{proj}_n \omega_m \succeq \text{proj}_n \omega'_m$. 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 $n \leq m$, $\omega_n \succeq \omega'_n$. \square

Corollary A.9. Let $\omega, \omega' \in \Omega$ be such that $\text{proj}_n(\omega) \not\succeq \text{proj}_n(\omega')$ for some $n \geq 1$, but for all $m < n$, $\text{proj}_m(\omega) \succeq \text{proj}_m(\omega')$. Then, there exists finite sequences $(P_k)_{k=1}^{n-1}$ and $(s_k)_{k=1}^{n-1}$ which induce canonical ICPS $\omega_{(n-k)}^i := \tau^*(\omega_{(n-k+1)}^i, P_k, s_k) \in \hat{\Omega}_{n-k}$ where $P_k \in \Gamma^*(\omega_{n-k+1}^i)$, such that $\Gamma^*(\omega_1^i)$ does not setwise-Blackwell dominate $\Gamma^*(\omega_1^j)$.

Proof. If not, we would have $\omega_n^1 \succeq \omega_n^2$, a contradiction. \square

Let \succeq be a dynamic order on \mathbf{M} defined as follows: For ICPS $\mathcal{M} = (\Theta, \Gamma, \tau, \theta_0)$ and $\mathcal{M}' = (\Theta', \Gamma', \tau', \theta'_0)$,

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

It is easy to see that such a dynamic order exists. Indeed, for any ICP \mathcal{M} , let $\omega_{\mathcal{M}}$ denote the canonical ICP that is indistinguishable from it. (By part (a) of Proposition A.3, there is a unique such $\omega_{\mathcal{M}}$.) The dynamic Blackwell order is induced on \mathbf{M} as follows: \mathcal{M} dynamically Blackwell dominates \mathcal{M}' if, and only if, $\omega_{\mathcal{M}} \succeq \omega_{\mathcal{M}'}$. The dynamic Blackwell order on \mathbf{M} clearly satisfies the condition \star . We now demonstrate that it is the *largest* order that satisfies \star .

Proposition A.10. Let \succeq be a dynamic order on \mathbf{M} that satisfies \star . If $\mathcal{M} \succeq \mathcal{M}'$, then \mathcal{M} dynamically Blackwell dominates \mathcal{M}' .

Proof. We will prove the contrapositive. If \mathcal{M} does not dynamically Blackwell dominate \mathcal{M}' , then $\omega_{\mathcal{M}} \not\succeq \omega_{\mathcal{M}'}$. Corollaries A.7 and A.8 imply that there is a smallest n such that $\omega_{\mathcal{M},n} \not\succeq \omega_{\mathcal{M}',n}$ but that for all $m < n$, $\omega_{\mathcal{M},m} \succeq \omega_{\mathcal{M}',m}$ (where $\omega_{\mathcal{M},n} = \text{proj}_n \omega_{\mathcal{M}}$ as defined in Appendix A.5). From Corollary A.9 it follows that there exists a finite sequence of partitions (P_k) and states (s_k) such that $\Gamma^*(\tau^{*(n)}(\theta_0, (P_k), (s_k)))$ does not setwise Blackwell dominate $\Gamma^*(\tau^{*(n)}(\theta'_0, (P_k), (s_k)))$, 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) . Now recall that \mathcal{M} is indistinguishable from $\omega_{\mathcal{M}}$, and so is \mathcal{M}' from $\omega_{\mathcal{M}'}$. This implies $\Gamma(\tau^{(n)}(\theta_0, (P_k), (s_k)))$ does not setwise Blackwell dominate $\Gamma(\tau^{(n)}(\theta'_0, (P_k), (s_k)))$. Thus, it must necessarily be that $\mathcal{M} \not\succeq \mathcal{M}'$. \square

B. Identification and Behavioral Comparison: Proofs from Section 4

Based on the results and notation from Appendices A.3—A.5, 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 $\omega' \in \Omega^S$ such that $(P, \omega') \in \omega$.

In what follows, ℓ^* and ℓ_* are the best and worst consumption streams, respectively, while c_s^+ and c_s^- denote the best and worst instantaneous consumption (prize) 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 B.1. Let $(P, \omega') \in \omega$. Then, there exists a menu $x(P, \omega')$ recursively defined as $x(P, \omega') = \{f_J : J \in P\}$ with

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

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 $x_{n-1}(P, \omega')$ for each $(P, \omega') \in \omega$, and define, for each cell $J \in P$, the act

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

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

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

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

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

as claimed. □

It is straightforward to verify that

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

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

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

Proof. Fix $(P, \omega) \in \Omega$ and consider the menu $x(P, \omega)$ defined in [★]. As noted above, for all ω' , we have $V(x(P, \omega), (P, \omega), 0) = V(x(P, \omega'), (P, \omega'), 0)$. Moreover, it must be that for all (Q, ω') (even for $Q = P$), we have $V(x(P, \omega), (P, \omega), 0) \geq V(x(P, \omega'), (Q, \omega'), 0)$ and in the case where Q is not finer than P and $Q \neq P$, $V(x(P, \omega'), (P, \omega'), 0) > V(x(P, \omega'), (Q, \omega'), 0)$ by construction of the menu $x(P, \omega')$. (This is straightforward to verify and is a version of Blackwell's theorem on comparison of experiments; see Blackwell (1953) or Theorem 1 on p59 of Laffont (1989).) \square

Lemma B.3. Suppose ω' does not dynamically Blackwell dominate ω . Then, for some $(P, \tilde{\omega}) \in \omega$, $x(P, \tilde{\omega})$ defined in [★] is such that $V(x(P, \tilde{\omega}), \omega, 0) = V(x(P, \tilde{\omega}), (P, \tilde{\omega}), 0) > V(x(P, \tilde{\omega}), \omega', 0)$.

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

From Corollary A.9 it follows that there exist finite sequences of partitions (P_k) and (P'_k) , and states (s_k) such that $\Gamma^*(\tau^{*(n)}(\omega', (P'_k), (s_k)))$ does not setwise Blackwell dominate the set $\Gamma^*(\tau^{*(n)}(\omega, (P_k), (s_k)))$, 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) , $\omega_{n-k}^i = \tau^*(\omega_{n-k+1}^i, P_k, s_k)$ where $P_k \in \Gamma^*(\omega_{n-k+1}^i)$, and $\Gamma^*(\omega_1^1)$.

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

Proof of Theorem 1. It follows from a straightforward extension of the arguments in Krishna and Sadowski (2014) (to the case of a compact prize space) that the collection $((u_s), \Pi, \delta)$ is unique in the sense of the Theorem. Now, define $F_\omega := \{x(P, \tilde{\omega}) : (P, \tilde{\omega}) \in \omega\}$. It follows immediately from Lemma B.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 B.4. Let $\omega, \omega' \in \Omega$. Then, the following are equivalent.

- (a) ω dynamically Blackwell dominates ω' .
- (b) For any $((u_s), \Pi, \delta)$ that induces $\omega \mapsto V(\cdot, \omega, \cdot)$, we have $V(x, \omega, \cdot) \geq V(x, \omega', \cdot)$ for all $x \in X$.

Proof. That (a) implies (b) is easy to see. That (b) implies (a) is merely the contrapositive to Lemma B.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 $\succsim^\dagger|_L = \succsim|_L$. This, and the uniqueness of the Recursive Anscombe-Aumann (RAA) representation (see Appendix C.2 below) together imply that $((u_s), \delta, \Pi) = ((u_s^\dagger), \delta^\dagger, \Pi^\dagger)$ after a suitable (and behaviorally irrelevant) normalization of the state-dependent utilities. Thus, part (b) of Corollary B.4 holds, which establishes the claim.

The ‘if’ part follows immediately from Corollary B.4. \square

C. Axioms

In this section we formally present axioms on the preference \succsim over X ; by Theorem 3 these axioms are necessary and sufficient for an ICP representation as discussed in the text. In addition, Appendix E investigates the implications of further strengthening our notions of Stationarity and Separability, and also shows that imposing Independence implies that information is not determined by a choice process, but instead exogenously arrives over time.

Note that the axioms are numbered in the order of their description in the text, and not sequentially. In particular, we state Axiom 5 immediately after stating Axiom 1.

C.1. Standard Properties

Our first axiom collects basic properties of \succsim that are common in the menu-choice literature.

AXIOM 1 (Basic Properties).

- (a) Order: \succsim is non-trivial, complete, and transitive.
- (b) Continuity: The sets $\{y : y \succsim x\}$ and $\{y : x \succsim y\}$ are closed for each $x \in X$.
- (c) Lipschitz Continuity: There exist $\ell^\#, \ell_\# \in L$ and $N > 0$ such that for all $x, y \in X$ and $t \in (0, 1)$ with $t \geq Nd(x, y)$, we have $(1 - t)x + t\ell^\# \succ (1 - t)y + t\ell_\#$.
- (d) Monotonicity: $x \cup y \succsim x$ for all $x, y \in X$.
- (e) Aversion to Randomization: If $x \sim y$, then $x \succsim \frac{1}{2}x + \frac{1}{2}y$ for all $x, y \in X$.

Items (a)–(d) are standard.⁵¹ Item (e) is familiar from Ergin and Sarver (2010) and de Oliveira et al. (2017) and relaxes Independence in order to accommodate unobserved information choice.

The next axiom captures the special role played by consumption streams, which leave no consumption choice to be made in the future and therefore require no information (that is, all

(51) For a discussion of (c) see Dekel et al. (2007) and for (d) see Kreps (1979).

information alternatives perform equally well). The axiom thus 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 $s \in S$. By Continuity (Axiom 1(b)) and the compactness of L , there exist best and worst consumption streams. As in Appendix C.2, we denote these by ℓ^* and ℓ_* , respectively. For each $I \subset S$, $f \in \mathcal{F}(\Delta(C \times X))$, $(c, y) \in C \times X$, and $\varepsilon \in [0, 1]$, define $f \oplus_{\varepsilon, I}(c, y) \in \mathcal{F}(\Delta(C \times X))$ by

$$(f \oplus_{\varepsilon, I}(c, y))(s) := \begin{cases} (1 - \varepsilon)f(s) + \varepsilon(c, y) & \text{if } s \in I \\ f(s) & \text{otherwise} \end{cases}$$

That is, for any state $s \in I$, the act $f \oplus_{\varepsilon, I}(c, y)$ perturbs the continuation lottery with y .

Fix $c \in C$ and let $\ell_s := \ell_* \oplus_{1, s}(c, \ell) \in L$, and define the induced binary relation \succsim_s on L by $\ell \succsim_s \hat{\ell}$ if $\ell_s \succsim \hat{\ell}_s$. Of course, in principle, \succsim_s depends on the choice of c . However, it will follow from Axiom 5 that the particular choice of c is irrelevant.

AXIOM 5 (Consumption Stream Properties).

- (a) *L*-Independence: For all $x, y \in X$, $t \in (0, 1]$, and $\ell \in L$, $x \succ y$ implies $tx + (1 - t)\ell \succ ty + (1 - t)\ell$.
- (b) *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 $(c, \ell_s) \succsim_{s''}(c, \hat{\ell}_s)$.
- (c) *L*-Stationarity: For all $\ell, \hat{\ell} \in L$ and $c \in C$, $\ell \succsim \hat{\ell}$ if, and only if, $(c, \ell) \succsim (c, \hat{\ell})$.
- (d) *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 5(a) is closely related to the C-Independence axiom in Gilboa and Schmeidler (1989), and is motivated in a similar fashion: 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 (b) through (d) see Krishna and Sadowski (2014).

C.2. Representation on Consumption Streams

Before we proceed, it is useful to record a consequence of assuming Axioms 1 and 5.

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

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

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

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

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

A preference on L has an RAA representation if, and only if, it satisfies Axioms 1 and 5. Corollary 5 of Krishna and Sadowski (2014) establishes this result for the case of finitely many prizes (ie, when C is a finite set), so we cannot directly appeal to their result. Nonetheless, judicious and repeated applications of Corollary 5 of Krishna and Sadowski (2014) allows us to reach the same conclusion for a compact set of prizes.⁵²

It is clear that L is compact, so the continuity of \succsim implies that there exist \succsim -maximal and -minimal elements of L . These we call ℓ^* and ℓ_* . Moreover, given that $\succsim|_L$ has an RAA representation as described above, for each $s \in S$, we let $c_s^+ := \arg \max_{c \in C} u_s(c)$ and $c_s^- := \arg \min_{c \in C} u_s(c)$. Because each u_s is continuous, such c_s^+ and c_s^- must exist. Now, define $f^+ \in \mathcal{F}(\Delta(C))$ to be the act such that $f^+(s) := c_s^+$ — the Dirac measure concentrated at c_s^+ — for all $s \in S$, and similarly, define $f^-(s) := c_s^-$ for all $s \in S$. Then, ℓ^* is the (unique) consumption stream that delivers f^+ at each date and ℓ_* is the (unique) consumption stream that delivers f^- at each date. Observe that the best and worst consumption streams are deterministic, and that for all $\alpha_1 \in \Delta(C)$, $u_s(c_s^-) \leq u_s(\alpha_1) \leq u_s(c_s^+)$. An immediate consequence of this is that for any $c \in C$, $\ell \in L$ and $s \in S$, $(c, \ell^*) \succeq_s (c, \ell) \succeq_s (c, \ell_*)$. Lipschitz Continuity (Axiom 1(c)) implies that $\ell^* \succ \ell_*$ (see Corollary 3.4 in the Supplementary Appendix), so $(c, \ell^*) \succ_s (c, \ell_*)$. We now state our remaining axioms.

C.3. State-Contingent Indifference to Correlation

As discussed in Section 5.1, we assume that DM's value for a menu does not change when substituting act f with g as long as they induce, on each state s , the same marginal distributions over C and X . For any $f \in \mathcal{F}(\Delta(C \times X))$, we denote by $f_1(s)$ and $f_2(s)$ the marginals of $f(s)$ on C and X , respectively.

AXIOM 2 (State-Contingent Indifference to Correlation). For a finite menu x , if $f \in x$ and $g \in \mathcal{F}(\Delta(C \times X))$ are such that $g_1(s) = f_1(s)$ and $g_2(s) = f_2(s)$ for all $s \in S$, then $[(x \setminus \{f\}) \cup \{g\}] \sim x$.⁵³

C.4. Indifference to Incentivized Contingent Commitment

Suppose that, contingent on a sequence of actions and realizations, DM is offered a chance to replace a certain continuation problem with another. DM's attitude towards such replacements may depend on his previous information choices, which are subjective, unobserved, and

(52) The formalities are established in Section 4 of the Supplementary Appendix.

(53) Axiom 2 is closely related to Axiom 5 in Krishna and Sadowski (2014), where other related notions of separability are also mentioned. The important difference is that Axiom 2 requires indifference to correlation in *any* menu x , rather than just singletons, because different information may be optimal for different menus.

menu-dependent. That said, any strategy of choice from a menu gives rise to a consumption stream. Therefore, any continuation problem y should leave DM no worse off than receiving the worst consumption stream, ℓ_* . In particular, because the best consumption stream, ℓ^* , leaves DM strictly better off than ℓ_* in every state, optimal choice from a menu $(1-t)x + t\ell^*$ should give rise to a consumption stream that is also strictly better than ℓ_* .

Formally, let $X^* := \{(1-t)x + t\ell^* : x \in X \text{ is finite, } t \in (0, 1)\}$. For a mapping $e : x \rightarrow (0, 1]$, let $x \oplus_{e,s} y := \{f \oplus_{e(f),s} (c_s^-, y) : f \in x\}$, which perturbs the continuation lottery in state s for any act f in x by giving weight $e(f)$ to (c_s^-, y) . For $x \in X^*$ we then require $x \succ [x \oplus_{e,s} \ell_*]$ and $[x \oplus_{e,s} y] \succeq [x \oplus_{e,s} \ell_*]$ for all $s \in S$ and $y \in X$. This is part (a) of Axiom 3 below.

Part (b) investigates the conditions under which DM is actually *indifferent* to replacing continuation lotteries with the worst consumption stream. In Section 5.1 we suggest a state-contingent notion of strategic rationality, according to which there should be a contingent plan that specifies which act DM will choose for each state, such that he will be indifferent between the original menu and one where he is penalized whenever his choice does not coincide with that plan.

To formalize this state contingent notion of strategic rationality, we define the set of contingent plans \mathcal{E}_x to be the collection of all functions $\xi : S \rightarrow x$. An *Incentivized Contingent Commitment* to $\xi \in \mathcal{E}_x$, is then the set

$$\mathcal{F}(\xi) = \{f \oplus_{1,I^c} (c_s^-, \ell_*) : f \in x \text{ and } I = \{s : f = \xi(s)\}\}$$

which replaces the outcome of f with the worst outcome (c_s^-, ℓ_*) in any state where f should not be chosen according to ξ . Obviously $x \succeq \mathcal{F}(\xi)$ for all $\xi \in \mathcal{E}_x$. However, if for no $s \in S$ is it ever optimal to choose an act outside $\xi(s)$, then $x \sim \mathcal{F}(\xi)$ should hold.

AXIOM 3 (Indifference to Incentivized Contingent Commitment).

- (a) If $x \in X^*$ and $e : x \rightarrow (0, 1]$, then $x \succ [x \oplus_{e,s} \ell_*]$ and $[x \oplus_{e,s} y] \succeq [x \oplus_{e,s} \ell_*]$ for all $s \in S$ and $y \in X$.
- (b) For all $x \in X$, there is $\xi \in \mathcal{E}_x$ such that $x \sim \mathcal{F}(\xi)$.

C.5. Concordant Independence

We envision information constraints where the choice of partition and the actual realization of the payoff-relevant state in the initial period fully determine the available information choices in the subsequent period. We say that x and y are *concordant* if the same initial information choice is optimal for both x and y . As we argue in Section 5.1, if x and y are concordant, then both should be concordant with the convex combination $\frac{1}{2}x + \frac{1}{2}y$. While Independence may be violated when considering menus that lead to different optimal initial information choices, $\succeq|_{X'}$ should satisfy Independence if $X' \subset X$ consists only of concordant menus. That is, any violation of Independence is entirely due to a change in the choice of information. We now introduce our behavioral notion of concordance (Definition C.1 below).

Observe that any one-period menu $z \in \mathcal{K}(L)$ does not permit any choice after the initial period, so that its value depends only on the partition under which it is evaluated. In particular, for $x_1(P) := \{\ell^* \oplus_{1,I^c} (c_s^-, \ell_*) : I \in P\} \in \mathcal{K}(L)$, we have $x_1(P) \sim \ell^*$ if, and only if, $x_1(P)$ is evaluated under a partition that is at least as fine as P . (See also Section 4.1 for a simple version of this menu when the prize space is $C = [0, 1]$ and further assumptions have been made on utilities.) 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 $P \in \mathcal{P}$ 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, $\frac{1}{2}y + \frac{1}{2}x_1(P) \sim \frac{1}{2}y + \frac{1}{2}\ell^*$.⁵⁴

Definition C.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 $P \in \mathcal{P}$.

AXIOM 4 (Concordant Independence). If x and y are concordant, so are x and $\frac{1}{2}x + \frac{1}{2}y$. Furthermore, if $X' \subset X$ consists of pairwise concordant menus, then $\succsim|_{X'}$ satisfies Independence.⁵⁵

D. Existence

We take as a starting point for our proof, the following representation

$$[\mathbf{D.1}] \quad V(x) = \max_{P \in \mathfrak{M}_p^\#} \sum_{J \in P} \left[\max_{f \in x} \sum_s \pi_0(s | J) [u_s(f_1(s)) + v_s(f_2(s), P)] \pi_0(J) \right]$$

where $\mathfrak{M}_p^\#$ is a finite collection of partitions of S , $u_s \in \mathbf{C}(C)$, and $v_s(\cdot, P) \in \mathbf{C}(X)$ for each $s \in S$ and $P \in \mathfrak{M}_p^\#$, with the property that for all $P, P' \in \mathfrak{M}_p^\#, s \in S$, $v_s(\cdot, P)|_L = v_s(\cdot, P')|_L$. Theorem 2 in the Supplementary Appendix establishes that \succsim satisfies Axioms 1–5 if, and only if, it has a representation as in [D.1].

We also take as given the existence of the Recursive Anscombe-Aumann Representation for $\succsim|_L$, described above in Appendix C.2.

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 D.1. We also provide here a behavioral definition for a finite menu x to have a unique optimal partition. In Appendix D.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 D.1.2, we derive some useful properties of consumption streams based on the recursive application of our axioms. Finally, we establish the existence of self-generating representations (by proving Proposition D.5) in Appendix D.1.3.

Given a self-generating representation, Appendix D.2 shows that it is possible to extract the underlying canonical ICP. In Appendix D.3, we note that \succsim has a unique representation on

(54) See Lemma 3.32 in the Supplementary Appendix for a instantiation of this intuition.

(55) If $x, y, z, (1-t)x + tz, (1-t)y + tz \in X', t \in (0, 1)$, and $x \succ y$, then $(1-t)x + tz \succ (1-t)y + tz$.

L . We use this fact, and the canonical ICP extracted from the self-generating representation, to prove Theorem 3.

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

- $u_s \in \mathbf{C}(C)$ for all $s \in S$,
- $\mathbb{Q} \subset \mathcal{P}$,
- $v_s(\cdot, P) \in \text{BL}(X)$ for all $s \in S$ and $P \in \mathbb{Q}$,⁵⁶
- $\pi \in \Delta(S)$,
- $u_s(c_s^\dagger) = v_s(\ell^\dagger, P) = 0$ 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

$$v(x) = \max_{P \in \mathbb{Q}} \sum_{E \in \mathcal{P}} \pi(E) \max_{f \in x} \sum_{s \in S} \pi(s | E) [u_s(f_1(s)) + v_s(f_2(s), P)]$$

In that case, we say that the tuple $((u_s), \mathbb{Q}, (v_s(\cdot, P)), \pi)$ 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 $\mathbf{A} : 2^{\mathbf{C}(X)} \rightarrow 2^{\mathbf{C}(X)}$ as follows:

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

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

Proof. It is easy to see that 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 $\mathbf{A}\Phi \neq \emptyset$.) 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 $\Phi^* \in 2^{\mathbf{C}(X)}$.

(56) The space $\text{BL}(X)$ consists of all bounded Lipschitz functions on X ; see Appendix A.2 for a definition.

To see that $\Phi^* \neq \{\mathbf{0}\}$, ie, Φ^* does not contain only the trivial function $\mathbf{0}$, fix $\mathbb{Q} = \{\{\{s\} : s \in S\}\}$ 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 $((u_s), \mathbb{Q}, \delta V), \pi$, while δV is implemented by $((\delta u_s), \mathbb{Q}, \delta^2 V), \pi$, and so on. Therefore, the set $\Phi_V := \{\delta^n V : n \geq 0\}$ is clearly 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 $((u_s), \mathbb{Q}, (v_s(\cdot, P)), \pi)$ implements v . Then, for all $\lambda \geq 0$, $((\lambda u_s), \mathbb{Q}, (\lambda v_s(\cdot, P)), \pi)$ 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 $((u_s), \mathbb{Q}, (v_s(\cdot, P)), \pi)$ 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. Also, unlike Abreu, Pearce, and Stacchetti (1990), the non-emptiness of Φ^* follows relatively easily, as noted in the proof of Proposition D.1.) In what follows, if \succsim is represented by $V \in \Phi^*$, we shall say that V is a *self-generating representation* of \succsim . In what follows, we show from the representation [D.1] that recursively applying Axioms 1–5 on \succsim implies that it has a self-generating representation.

We begin by behaviorally characterizing when a menu has a unique optimal information choice, by relying on the menu $x_1(P)$ discussed in Section 4.1 and formally defined in the proof of Lemma B.1.

Definition D.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 uniquely optimal information choice.

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 .⁵⁷ Because $\frac{1}{2}x' + \frac{1}{2}x_1(P) \succ \frac{1}{2}x' + \frac{1}{2}x_1(Q)$ if P not finer than Q , learning 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 in the representation in [D.1], let X'_P be defined as follows:

$$X'_P := \{x : V(x) = V(x, P, (v_s(\cdot, P))) \text{ for some } P \in \mathfrak{M}_p^\# \text{ and} \\ V(x) > V(x, Q, (v'_s(\cdot, Q))) \text{ for all } Q \in \mathfrak{M}_p^\# \text{ such that } P \neq Q\}$$

The following corollary records a useful property of the set X'_P .

(57) 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 .

Corollary D.3. Let \succsim have a representation as in [D.1], let $x \in X'_p$ and let Y_x denote the set of menus that are (i) concordant with x , and (ii) have a unique optimal partition. Then, $Y_x = X'_p$.

The corollary follows immediately from the representation in [D.1].

Definition D.4. If for $y, y' \in X, s \in S$, and finite x that has a unique optimal information choice (see Definition D.2), there is $\varepsilon \in (0, 1]$ such that $x \oplus_{\varepsilon, s} y, x \oplus_{\varepsilon, s} y'$, and x are pairwise concordant, then $y \succsim_{(x, s)} y'$ if $[x \oplus_{\varepsilon, s} y] \succsim [x \oplus_{\varepsilon, s} y']$.⁵⁸

With $\succsim_{(x, s)}$ now defined, we can recursively apply our Axioms.

Proposition D.5. Let \succsim be a binary relation on X . Then, the following are equivalent.

- (a) \succsim recursively satisfies Axioms 1–5.
- (b) \succsim has a self-generating representation, ie, there exists a function $V \in \Phi^*$ that represents \succsim .

In Appendix D.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 .

In Appendix D.1.2, we derive some useful properties of consumption streams based on the recursive application of our axioms. Finally, we prove Proposition D.5 in Appendix D.1.3.

D.1.1. Continuation Preferences

Lemma D.6. 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 $d(f_i, x') := \min_j d(f_i, f'_j) < \varepsilon$. 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. \square

Notice that $\mathfrak{M}_p^\#$ in [D.1] is finite and can be taken to be minimal (in the sense that if $\mathfrak{M}_p^\#$ is another set that represents V as in [D.1], then $\mathfrak{M}_p^\# \subset \mathfrak{M}_p^\#$) without affecting the representation. Recall that $X^* := \{(1-t)x + t\ell^* : x \in X \text{ is finite}, t \in (0, 1)\}$.

Lemma D.7. Let \succsim have a representation as in [D.1]. For all $P \in \mathfrak{M}_p^\#$, there exists a finite $x \in X'_p \cap X^*$. In particular, x can be written as $x = \frac{1}{2}x' + \frac{1}{2}x_1(P)$ for some $x' \in X$.

Proof. The finiteness and minimality of $\mathfrak{M}_p^\#$ in [D.1] implies that for any $P \in \mathfrak{M}_p^\#$, 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 immediately from the representation in [D.1] that $x := \frac{1}{2}x' + \frac{1}{2}x_1(P) \in X'_p \cap X^*$, as claimed. \square

Lemma D.8. Let \succsim have a representation as in [D.1]. For all $P \in \mathfrak{M}_p^\#, v_s(y, P) \geq v_s(\ell^*, P)$.

(58) Recall the definition of $x \oplus_{\varepsilon, s} y$ from Appendix C.4. Slightly abusing notation, we write $x \oplus_{\varepsilon, s} y$ if $e(f) = \varepsilon$ for all $f \in x$.

Proof. Suppose instead that $v_s(y, P) < v_s(\ell_*, P)$. Consider $x \in X'_P \cap X^*$ which exists by Lemma D.7. Then, for $\varepsilon > 0$ small enough such that P remains optimal, $[x \oplus_{\varepsilon, s} \ell_*] \succ [x \oplus_{\varepsilon, s} y]$. To see this, suppose $f \oplus_{\varepsilon, s} (c_s^-, y)$ is chosen optimally from the menu $x \oplus_{\varepsilon, s} y$. Then, $v_s(y, P) < v_s(\ell_*, P)$ implies

$$\begin{aligned} & (1 - \varepsilon)[u_s(f_1(s)) + v_s(f_2(s), P)] + \varepsilon [u_s(c_s^-) + v_s(y, P)] \\ & < (1 - \varepsilon)[u_s(f_1(s)) + v_s(f_2(s), P)] + \varepsilon [u_s(c_s^-) + v_s(\ell_*, P)] \end{aligned}$$

This implies $V(x \oplus_{\varepsilon, s} \ell_*) > V(x \oplus_{\varepsilon, s} y)$. But this contradicts part (a) of IICC (Axiom 3), which requires that $[x \oplus_{\varepsilon, s} y] \succeq [x \oplus_{\varepsilon, s} \ell_*]$ for all $y \in X$. \square

Lemma D.9. Let \succeq have a representation as in [D.1]. Fix $P \in \mathfrak{M}_P^\#$. For any finite $x \in X'_P$ and $s \in S$, $\succeq_{(x, s)}$ is independent of the choice of $\varepsilon \in (0, 1)$ for which Definition D.4 applies. Further, $\succeq_{(x, s)}$ is complete on X and is represented by $v_s(\cdot, P)$. Finally, if x' is finite, has a unique optimal partition, and is concordant with x , then $\succeq_{(x, s)} = \succeq_{(x', s)}$.

Proof. Let $x \in X'_P$ be finite, so that $V(x) = V(x, P)$. Fix $s \in S$. Because V in [D.1] is continuous, there is $\varepsilon > 0$ such that P is the unique optimal partition for all $x' \in B(x; \varepsilon)$, and hence all $x', x'' \in B(x; \varepsilon)$ are concordant with each other (see Corollary D.3). By Lemma D.6, $[x \oplus_{\varepsilon, s} y], [x \oplus_{\varepsilon, s} y'] \in B(x; \varepsilon)$ for all $y, y' \in X$, establishing the completeness of $\succeq_{(x, s)}$. Then, $[x \oplus_{\varepsilon, s} y] \succeq [x \oplus_{\varepsilon, s} y']$ if, and only if, $V(x \oplus_{\varepsilon, s} y) \geq V(x \oplus_{\varepsilon, s} y')$. Suppose $f \oplus_{\varepsilon, s} y$ is optimally chosen from $x \oplus_{\varepsilon, s} y$ in state s . Then, it must be that

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

which implies $v_s(y, P) \geq v_s(y', P)$. Conversely, $v_s(y, P) \geq v_s(y', P)$ implies that if $f \oplus_{\varepsilon, s} (c_s^-, y')$ is optimally chosen from $x \oplus_{\varepsilon, s} y'$ in state s , then the inequality displayed above holds, which implies $[x \oplus_{\varepsilon, s} y] \succeq [x \oplus_{\varepsilon, s} y']$. But this is independent of our choice of $\varepsilon > 0$ as long as it maintains concordance.

Finally, if x and x' are concordant and x' has a unique optimal partition, then by Corollary D.3 $x' \in X'_P$. It follows that $\succeq_{(x', s)}$ is also represented by $v_s(\cdot, P)$, and hence $\succeq_{(x, s)} = \succeq_{(x', s)}$, which completes the proof. \square

Lemma D.10. The binary relation $\succeq_{(P, s)}$ on X which is represented by $v_s(\cdot, P)$ satisfies Axioms 1–5 recursively.

Proof. By Lemma D.9, $\succeq_{(P, s)} = \succeq_{(x, s)}$ for some $x \in X'_P$. Since $\succeq_{(x, s)} \in \Psi^*$, it satisfies Axioms 1–5 recursively. \square

D.1.2. Some Properties of Consumption Streams

We now relate preferences on L to those on X .

Let $\tilde{X}_1 := \mathcal{K}(\mathcal{F}(\Delta(C \times \{\ell_*\})))$ be the space of one-period problems that always give ℓ_* at the beginning the second period. Inductively define $\tilde{X}_{n+1} := \mathcal{K}(\mathcal{F}(\Delta(C \times \tilde{X}_n)))$ 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 D.11. The set $\tilde{X} \subset X$ is dense in X .

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

For any $x = (x_1, x_2, \dots) \in X$ 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 D.12. Let \succsim satisfy Axioms 1–5. Then, for any $s \in S$ and $P \in \mathfrak{M}_p^\#$, $\ell^* \succsim_{(P,s)} \ell \succsim_{(P,s)} \ell_*$ for all $\ell \in L$.

Proof. The preference \succsim has a separable and partitional representation as in [D.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 C.2, \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 $Q \in \mathfrak{M}_p^\#$ in the representation [D.1]. But $v_s(\cdot, P)$ represents $\succsim_{(P,s)}$ which implies that $\ell^* \succsim_{(P,s)} \ell \succsim_{(P,s)} \ell_*$ for all $\ell \in L$, $s \in S$. \square

Proposition D.13. Let \succsim satisfy Axioms 1–5 recursively. Then, for all $x \in X$, $\ell^* \succsim x$.

Proof. By the continuity of \succsim and by Lemma D.11, 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 in [D.1] that $V(\tilde{x}_1) \leq V(\ell^*)$ for all $\tilde{x}_1 \in \tilde{X}_1$. Recall that the representation in [D.1] is implied by \succsim satisfying Axioms 1–5.

But $\succsim \in \Psi^*$, so that by Lemma D.10, $\succsim_{(P,s)}$ also satisfies Axioms 1–5 for any $P \in \mathfrak{M}_p^\#$, 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 D.12, 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 satisfies Axioms 1–5 recursively, 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 D.10 $\succsim_{(P,s)}$ satisfies Axioms 1–5 recursively on X , so that we must also have $v_s(\ell^*, P) \geq v_s(\tilde{x}_n, P)$ for all $\tilde{x}_n \in \tilde{X}_n$ (where we have appealed to Lemma D.12 to establish that ℓ^* is the $v_s(\cdot, P)$ -best consumption stream). In particular, this implies that for any lottery $\alpha_2 \in \Delta(\tilde{X}_n)$, $v_s(\ell^*, P) \geq v_s(\alpha_2, P)$.

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

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

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. \square

D.1.3. Proof of Proposition D.5

Proof. To see that (b) implies (a), suppose \succsim has the representation [D.1]. By Proposition 3.30 of the Supplementary Appendix, \succsim satisfies Axioms 1–5. All that remains to establish is that \succsim satisfies these axioms recursively.

Given a representation as in [D.1] that is *also* self-generating, let $x \in X$ be finite and $P \in \mathfrak{M}_P^\#$ be the unique optimal partition for x . By Lemma D.9, $\succsim_{(x,s)}$ is represented by $v_s(\cdot, P)$ on X . Because the representation is self-generating, $\succsim_{(x,s)}$ must satisfy Axioms 1–5 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 satisfies Axioms 1–5 recursively.

To see that (a) implies (b), note that Lemma D.10 has two implications. First, $\succsim_{(P,s)}$ has a separable and partitional representation $v'_s(\cdot, P)$ as in [D.1]. Because $v_s(\cdot, P)$ also represents $\succsim_{(P,s)}$ it follows that $v_s(\cdot, P)$ and $v'_s(\cdot, P)$ are identical up to a monotone transformation. But, by L-Indifference to Timing (Axiom 5(d)), 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 $\succsim_{(P,s)}$ satisfies Axioms 1–5 recursively, it satisfies the hypotheses of Proposition D.13. Together with Lemma D.12 and IICC (Axiom 3), this implies that $\ell^* \succsim_{(P,s)} y \succsim_{(P,s)} \ell_*$ for all $y \in X$. Because $v_s(\cdot, P)$ and $v'_s(\cdot, P)$ both represent $\succsim_{(P,s)}$, they must agree on X because they agree on L . It follows that $v_s(\cdot, P)$ also has a representation as in [D.1], that is, it can be written as

$$v_s(x, P) = \max_{P' \in \mathfrak{M}_P^\#(P)} \sum_{J \in \mathcal{Q}} \pi_0(J) \max_{f \in X} \sum_s \pi_0(s | J) [u'_s(f_1(s)) + v'_s(f_2(s); P')]$$

Then, because $\succsim_{(x,s)}$ satisfies Axioms 1–5 recursively, 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 [D.1], and so on, *ad infinitum*, which demonstrates that $V \in \Phi^*$. \square

D.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 D.14. There is a unique map $\varphi^* : \Phi^* \rightarrow \Omega$ that satisfies for some implementation $((u_s), \mathbb{Q}, (v_s(\cdot, P)), \pi)$ of v , that

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

and is independent of the implementation chosen.

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

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

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

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

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

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

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

and define the map $\varphi^* : \Phi^* \rightarrow \Omega$ as $\varphi^*(v) = (\varphi_1(v), \varphi_2(v), \dots)$, 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. For a self-generating representation, we can find a (not necessarily unique) probability measure π over S^∞ . The formal details are straightforward and hence omitted.

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

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

Let \mathfrak{S}_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 $\sigma_i^{(n)} : \mathfrak{S}_n \rightarrow \mathcal{P} \times \Omega^S$. Similarly, a *dynamic consumption plan* is a sequence $\sigma_c =$

$(\sigma_c^{(1)}, \sigma_c^{(2)}, \dots)$ of mappings where $\sigma_c^{(n)} : \mathfrak{S}_n \rightarrow \mathcal{F}(\Delta(C \times X))$. 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 $\sigma_i^{(n)}(\mathfrak{h}_n) = (P, \omega')$ $\in \omega^{(n-1)}$, $\sigma_c^{(n-1)}(\mathfrak{h}_n)$ is P -measurable, ie, for all $I \in P$ and for all $s, s' \in I$, $\sigma_c^{(n)}(\mathfrak{h}_n)(s) = \sigma_c^{(n)}(\mathfrak{h}_n)(s')$.

Each dynamic plan along with initial states (x, ω_0, π_0) induces a probability measure over $(X \times \Omega \times S)^\infty$ 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 $(x^{(n-1)}, \omega^{(n-1)}, s^{(n-1)})$.⁵⁹

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 D.14 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 D.15. 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

$$[\text{D.2}] \quad \lim_{n \rightarrow \infty} \left\| \mathbf{E}^{\sigma, \pi} v^{(n)}(\cdot, \omega^{(n)}, s^{(n)}, \sigma) \right\|_\infty = 0$$

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 ℓ^\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 $|V(\ell^\dagger \diamond_n x) - V(\ell^\dagger \diamond_n y)| < \varepsilon$.

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

$$\left| \mathbf{E}^{\sigma, \pi} v^{(n)}(x, \omega^{(n)}, s^{(n)}) - \mathbf{E}^{\sigma, \pi} v^{(n)}(y, \omega^{(n)}, s^{(n)}, \sigma) \right| < \varepsilon$$

for all $n \geq N$. Recall that

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

(59) Of course, the choice of plan doesn't affect the evolution of the objective states $(s^{(n)})$.

Moreover, we have

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

which completes the proof. \square

Adapting the terminology of Dubins and Savage (1976), we shall say that a function $V \in \Phi^*$ is *equalizing* if [D.2] holds. (To be precise, if [D.2] 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 $\times_n \mathfrak{S}_n$, 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 D.16. 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 [D.2] in Proposition D.15).⁶⁰

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 (see Footnote 50). We can now adapt Theorem 9.2 in Stokey, Lucas, and Prescott (1989) where, as above, their equation 7 is replaced by [D.2], 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

(60) 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.

D.3. Recursive Representation

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

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

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

where $u_s(c_s^\dagger) = 0$ 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 $((u_s)_{s \in S}, \delta, \Pi)$.

By Proposition D.5 \succsim has a self-generating representation $V \in \Phi^*$ that satisfies $V(\ell^\dagger) = 0$. Now, $V|_L$ and $V_L^*(\cdot, \pi_0)$ both 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 $((u_s)_{s \in S}, \Pi, \delta, \omega_0)$ 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 D.17. 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 .⁶¹

Proof. In this proof, we refer to objects defined in Appendix D.1. 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 D.16, 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 D.16 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

(61) It follows immediately from Proposition D.17 that in considering dynamic plans, we may restrict attention to stationary plans. This is because we have a recursive formulation with discounting 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).

Suppose V represents \succsim and $V \in \Phi^*$. Then, there exists an implementation of V , given by $((u_s), \mathbb{Q}, (v_s^{(1)}(\cdot, P)), \pi)$. 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.

E. Invariant Per-Period Constraint and Fixed Arrival of Information

By Theorem 3, \succsim satisfies Axioms 1–5 recursively if, and only if, it has an ICP representation $((u_s), \delta, \Pi, \mathcal{M})$. We now discuss two special cases of the ICP representation. In the first, DM faces the same information constraint each period. This case 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, measured by the expected reduction in entropy. Recall that $x \in X$ is \succsim -maximal if $x \succsim y$ for all $y \in X$.

AXIOM 6 (Stationary Maximal Choice Problem). $x \in X$ is \succsim -maximal if, and only if, it is $\succsim_{(y,s)}$ -maximal for all $y \in X$ and $s \in S$.

The axiom requires maximal menus to be stable in three ways: Stationarity, because between \succsim and $\succsim_{(y,s)}$ a period has passed; temporal separability, through the comparison of $\succsim_{(y,s)}$ and $\succsim_{(y',s)}$; and State Independence, through the comparison of $\succsim_{(y,s)}$ and $\succsim_{(y,s')}$.

Definition E.1. The ICP $\mathcal{M} = (\Theta, \Gamma, \tau, \theta)$ is an *invariant per-period constraint* if $\Gamma(\theta)$ is constant on Θ (or, equivalently, if Θ is a singleton).

In contrast to a general MIC, an invariant per-period constraint is independent of past information choice, and so does not accommodate any intertemporal trade-offs in processing information.

Proposition E.2. Let \succsim has ICP representation $((u_s), \delta, \Pi, \mathcal{M})$. It satisfies Axiom 6 if, and only if, \mathcal{M} is an invariant per-period constraint.

To see why this must be true, note that ℓ^* is both \succsim -best and $\succsim_{(y,s)}$ -best for all $y \in X$ and $s \in S$. It follows from the argument in Section A.3 that the ICP $\mathcal{M} = (\Theta, \Gamma, \tau, \theta_0)$ is indistinguishable from the ICP $(\Theta, \Gamma, \tau, \tau(\theta_0, P, s))$ for all $P \in \Gamma(\theta_0)$ and $s \in S$. The other direction is immediate.

In the second special case we consider, DM faces a trivial choice between information plans, that is, he can not influence the arrival of information about the state of the world.⁶²

AXIOM 7 (Independence). If $x \succ y$, then $tx + (1-t)z \succ ty + (1-t)z$ for all $x, y, z \in X$ and $t \in (0, 1)$.

(62) This parallels the representation in Krishna and Sadowski (2014), where DM faces a fixed stream of information about his own taste, rather than the state of the world.

Definition E.3. The ICP $\mathcal{M} = (\Theta, \Gamma, \tau, \theta)$ captures *fixed arrival of information* if $\Gamma(\theta)$ is a singleton for all $\theta \in \Theta$.

Proposition E.4. Let \succsim has ICP representation $((u_s), \delta, \Pi, \mathcal{M})$. It satisfies Axiom 7 if, and only if, \mathcal{M} captures fixed arrival of information.

To see why this must be true, suppose instead that $P, P' \in \Gamma(\theta)$ where P and P' are not ranked by fineness for some θ . Then $x_1(P) \sim x_1(P') \sim \ell^* \succ \frac{1}{2}x_1(P) + \frac{1}{2}x_1(P')$, contradicting Independence. This argument easily extends to ICPS that contain *any* two information plans that are not ranked by dynamic Blackwell dominance.

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. 24, 26, 46).
- Bertsekas, Dimitri P., and Steven Shreve. 2000. *Stochastic Optimal Control: The Discrete-Time Case*. Belmont, MA: Athena Scientific. (Cit. on p. 54).
- Blackwell, David. 1953. "Equivalent Comparison of Experiments". *Annals of Mathematical Statistics* 24:265–272. (Cit. on p. 39).
- Blanchard, Olivier Jean, and Stanley Fischer. 1989. *Lectures on Macroeconomics*. Cambridge, Ma: MIT Press. (Cit. on p. 53).
- Bolton, Patrick, and Mathias Dewatripont. 1994. "The Firm as a Communication Network". *Quarterly Journal of Economics* 109 (4): 809–839. (Cit. on p. 6).
- Caplin, Andrew, and Mark Dean. 2015. "Revealed Preference, Rational Inattention, and Costly Information Acquisition". *American Economic Review* 105 (7): 2183–2203. (Cit. on pp. 28, 29).
- Coeurdacier, Nicolas, and H el ene Rey. 2013. "Home Bias in Open Economy Financial Macroeconomics". *Journal of Economic Literature* 51 (1): 63–115. (Cit. on p. 7).
- Currie, Janet, and Bentley MacLeod. 2017. "Diagnosing Expertise: Human Capital, Decision Making, and Performance Among Physicians". *Journal of Labor Economics* 35 (1): 1–43. (Cit. on p. 6).
- 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. 27, 28).
- Dekel, Eddie, Barton L. Lipman, Aldo Rustichini, and Todd Sarver. 2007. "Representing Preferences with a Unique Subjective State Space: A Corrigendum". *Econometrica* 75 (2): 591–600. (Cit. on p. 40).

- Dillenberger, David, R. Vijay Krishna, and Philipp Sadowski. 2017. *Supplement to ‘Subjective Information Choice Processes’*. Tech. rep. Duke University. (Cit. on p. 5).
- 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. 12, 22, 27, 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. 53).
- Dudley, Richard Mansfield. 2002. *Real Analysis and Probability*. 2nd edition. Cambridge University Press. (Cit. on p. 30).
- 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. 27–29, 40).
- Geanakoplos, John, and Paul Milgrom. 1991. “A Theory of Hierarchies Based on Limited Managerial Attention”. *Journal of the Japanese and International Economies* 5:205–225. (Cit. on pp. 2, 6).
- Gilboa, Itzhak, and David Schmeidler. 1989. “Maxmin Expected Utility with a Non-Unique Prior”. *Journal of Mathematical Economics* 18:141–153. (Cit. on pp. 27, 41).
- Greenshtein, Eitan. 1996. “Comparison of Sequential Experiments”. *Annals of Statistics* 24 (1): 436–448. (Cit. on pp. 14, 18).
- Gul, Faruk, and Wolfgang Pesendorfer. 2004. “Self-Control and the Theory of Consumption”. *Econometrica* 72 (1): 119–158. (Cit. on pp. 9, 28, 31).
- Hébert, Benjamin, and Michael Woodford. 2016. *Rational Inattention with Sequential Information Sampling*. Tech. rep. Stanford University GSB. (Cit. on p. 29).
- 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. 15).
- 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 pp. 12, 22, 27, 40).
- . 2012. *Microeconomic Foundations I: Choice and Competitive Markets*. Princeton, NJ: Princeton University Press. (Cit. on p. 35).
- Kreps, David M., and Evan L Porteus. 1978. “Temporal Resolution of Uncertainty and Dynamic Choice Theory”. *Econometrica* 46 (1): 185–200. (Cit. on pp. 9, 10).
- Krishna, R. Vijay, and Philipp Sadowski. 2014. “Dynamic Preference for Flexibility”. *Econometrica* 82 (2): 655–703. (Cit. on pp. 10, 14, 15, 22, 23, 26–28, 39–42, 55).

- Laffont, Jean-Jacques. 1989. *The Economics of Uncertainty and Information*. The MIT Press. (Cit. on p. 39).
- Maccheroni, Fabio, Massimo Marinacci, and Aldo Rustichini. 2006. “Dynamic Variational Preferences”. *Journal of Economic Theory* 128:4–44. (Cit. on p. 28).
- Maćkowiak, Bartosz, and Mirko Wiederholdt. 2009. “Optimal Sticky Prices under Rational Inattention”. *American Economic Review* 99 (3): 769–803. (Cit. on pp. 6, 12).
- 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. 7).
- Matejka, 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. 35, 54).
- Ozdenoren, Emre, Stephen W. Salant, and Dan Silverman. 2012. “Willpower and the Optimal Control of Visceral Urges”. *Journal of the European Economic Association* 10 (2): 342–368. (Cit. on p. 8).
- Persico, Nicola. 2000. “Information Acquisition in Auctions”. *Econometrica* 68 (1): 135–148. (Cit. on p. 2).
- Piermont, Evan, Norio Takeoka, and Roe 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 p. 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 Decision-Making”. *Econometrica* 85:521–553. (Cit. on p. 11).
- 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 p. 53).
- Woodford, Michael. 2012. *Inattentive Valuation and Reference-Dependent Choice*. Tech. rep. Columbia University. (Cit. on p. 29).
- de Oliveira, Henrique. 2016. *Blackwell’s Informativeness Theorem using Category Theory*. Tech. rep. Pennsylvania State University. (Cit. on pp. 14, 18).
- 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, 12, 27–29, 40).