

Subjective Information Choice Processes*

DAVID DILLENBERGER[†] R. VIJAY KRISHNA[‡] PHILIPP SADOWSKI[§]

May 26, 2022 at 13:37 hrs

Abstract

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

KEY WORDS: Dynamic Preferences, Information Choice Process, Dynamic Blackwell Dominance, Rational Inattention, Subjective Markov Decision Process

JEL Classification: D80, D81, D90

(*) An earlier version of this paper has been circulated under the title “Subjective Dynamic Information Constraints”. We thank Attila Ambrus, Alexander Bloedel, Andrew Ellis, Itzhak Gilboa, Ed Green, Botond Kőszegi, Jay Lu, George Mailath, Robert Miller, Pietro Ortoleva, Wolfgang Pesendorfer, and Todd Sarver, as well as the editor, Ran Spiegler, and three anonymous referees for helpful comments. We thank 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>

1. Introduction

In a typical dynamic choice problem, a decision maker (henceforth DM) must choose an action that, contingent on the evolving state of the world, determines a payoff for the current period as well as the collection of actions available in the next period. Faced with such a 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 always be aware of relevant prices at all possible retailers and firms have limited resources they can expend on market analysis. While accounting for such information constraints can significantly change theoretical predictions (see, for instance, Geanakoplos and Milgrom (1991), Stigler (1961), Persico (2000), and the literature on rational inattention pioneered by Sims (1998, 2003)), an inherent difficulty in modeling them, as well as the actual choice of information, is that they are often private and unobservable.

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

To fix ideas, suppose in each period DM can manage his portfolio by choosing from a set of possible investments. Depending on the current state of the economy, each choice of investment results in an instantaneous payoff (eg, a dividend) and a new realization of the monetary value of the portfolio, which determines the continuation investment problem for the next period. Further suppose that to improve his portfolio choice, DM can acquire information about the state in a way that may also determine what information can be acquired in the future. 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 plans, is typically not. A natural question is, Can (unobservable) information constraints be identified from DM’s preferences and if so, what type of data is needed to achieve this identification?

Our main result (Theorem 1) shows that the class of dynamic consumption choice problems we consider is sufficiently rich to identify the entire set of (subjective) parameters governing DM’s preferences. In particular, observable choice is between menus of acts, each of which results in a state-dependent lottery over current consumption and a continuation problem for the next period, and the parameters are (i) a utility function over consumption, (ii) a Markov process that governs the evolution of the state of the world, (iii) the discount factor, and (iv) the Information Choice Process (ICP), which is an information constraint that is identified up to a dynamic extension of Blackwell informativeness.

Formally, our model is a Markov decision process for DM with a state variable (which governs the feasible information choices) that cannot be verified by the analyst; from her point of view, the process is subjective with partially unknown transitions and partially unobservable actions and states. The full identification of this partially subjective Markov decision process in Theorem 1 is our main conceptual contribution.

Understanding ICPs can be useful in a variety of settings. To illustrate this, we now briefly discuss the relevance of expertise and fatigue in organizational and personnel economics.

It is widely accepted that information overload among workers, which is tightly linked to fatigue,¹ emerges as a result of an imbalance between the worker’s processing capacity and his processing requirements. While the worker’s processing capacity (his ICP) is subjective, the processing requirement is determined by his employer. It is thus in the interest of the employer to understand the worker’s ICP so that she can avoid overloading the worker with information.

Organizations typically expect their employees to continue developing expertise (understood here as the ability to make good decisions) either through on-the-job

(1) One explanation for the “social media fatigue”, that is, the reduced interest in social media among US consumers, is that users are experiencing information overload: the amount of information to be processed has become overwhelming for consumers leading them to be more selective about their media exposure and involvement, see Greenfield (2018).

learning or through formal education (see for instance Pastorino 2021). From an organizational perspective, it is then useful to understand the mechanism by which different types of workers can acquire expertise (which is captured by their ICPS). This may determine how well the firm can (re-)match different types of workers to different tasks, choose between on-the-job training and formal education, use its organizational structure to incentivize skill acquisition, or avoid turnover by generating firm specific expertise (as opposed to general human capital).

The paper is organized as follows. Section 2 presents examples of ICPS. Section 3 introduces the analytical framework, states our utility representation, and describes the notion of comparative informativeness for ICPS. Section 4 establishes a duality between the set of ICPS and our choice domain and uses it to illustrate our identification result. Section 5 behaviorally characterizes expertise and fatigue. Section 6 surveys the most related literature. All proofs are in the appendices.

2. Examples of ICPS

An ICP is a subjective controlled process that specifies how future information constraints depend on past choices of information. Formally, given a state space S , from which the payoff-relevant state of the world realizes each period, an ICP is parameterized by an additional subjective *information state* θ , which encodes how today's choice of information affects information constraints in the future; a function $\Gamma(\theta)$ that determines a 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 $s \in S$. A trivial ICP, where $\Gamma(\theta)$ is always a singleton, corresponds to DM facing a particular stream of information, and so our framework subsumes the standard model of dynamic decision making without information choice. Another special case is where $\Gamma(\theta)$ is independent of θ and is non-trivial, so that the set of available signals is constant over time.

The novelty in our theory is that ICPS can accommodate arbitrary history dependence. We now describe a few simple ICPS that feature *expertise* and *fatigue*. Intuitively, expertise is the idea that past information choices facilitate the choice of the same or similar information in the future, while fatigue is the notion that learning more initially reduces the ability to learn and process information later. We will return to these examples in Section 5, where we formally define and behaviorally characterize expertise and fatigue.

Example 2.1. As has been proposed in the literature on managerial decision making, suppose a manager is constrained in the accumulation of human capital, or expertise.² Specifically, the manager faces N different sources of information, where each source i corresponds to a partition P_i of S . Suppose further that the manager has the resources to gain expertise in processing one new source of information every period. This might involve hiring a new expert, purchasing access to new data, or learning how to interpret a particular type of real world data. Once the manager gains that expertise, she can process the now familiar source of information indefinitely. The manager's information in a given period is then the coarsest refinement of all the partitions that correspond to previously chosen information sources. With $\theta = (\theta_1, \dots, \theta_N) \in \{0, 1\}^N$ as the subjective information state, let

$$P_{\theta_i} := \begin{cases} P_i & \text{if } \theta_i = 1 \\ \{S\} & \text{if } \theta_i = 0 \end{cases}$$

The information constraint is then

$$\Gamma(\theta) = \{P : P = P_{\theta_1} \vee \dots \vee P_{\theta_N} \vee P_i \text{ with } i \in \{1, \dots, N\}\}$$

and the state transitions so that the new j th component is given by

$$\tau(P_{\theta_1} \vee \dots \vee P_{\theta_N} \vee P_i, \theta)_j := \begin{cases} 1 & \text{if } j = i \\ \theta_j & \text{otherwise} \end{cases}$$

Example 2.2. DM is subject to fatigue, in that he cannot acquire or process 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 only learn the trivial partition of S) in the current one. This example suggests that periods in which individuals pay careful attention are usually followed by periods in which they should rest.³ In this case, we may set $\Theta = \{0, 1\}$, some $\theta_0 \in \Theta$ as the initial information state,

-
- (2) See, for instance, Geanakoplos and Milgrom (1991). We refer here 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.
- (3) In addition to the cognitive interpretation of fatigue, this type of constraint could also capture advice from an expert who can only be approached infrequently. Alternatively, the acquisition of 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.

and

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

where \mathcal{P} is the collection of all partitions of S . Note that here τ is independent of θ and s .

Our next two examples also capture fatigue and expertise, respectively, and build on the entropy based constraints found in the literature on rational inattention.

Example 2.3. The mental exertion of processing an overload of information in one period results in fatigue that negatively impacts DM's capacity to process information in the subsequent period. Let $c(P)$ measure the amount of information from partition P . For example, $c(P) = H_\mu(P)$, the entropy of P calculated using some probability distribution μ over S (where μ may evolve as information arrives; for instance, if past state realizations become known and states are correlated over time, then μ will depend on those past states). When rested, DM can process an amount K , and she begins each period rested as long as she processed less than $\kappa < K$ in the previous period. Processing more than κ in one period amounts to information overload, and her ability to process information in the next period will be reduced from K by the amount of overload. In that case her per-period information capacity k can serve as the subjective information state. Formally, with attention capacity k , any partition $P \in \Gamma(k) = \{P : c(P) \leq k\}$ can be chosen, whereupon the stock transitions to K if $c(P) \leq \kappa$ and to $K - c(P) + \kappa$ otherwise. This ICP is parametrized by (K, κ, c) . If c is an entropy cost, then the case with $\kappa = K$ corresponds to a typical per period constraint in the literature on rational inattention.

Example 2.4. To capture expertise instead of fatigue in an environment similar to that of the previous example, suppose that DM has capacity K to process *new* information every period. In particular, if partition Q was chosen in the previous period, then the cost of learning P now is $c(P | Q) = H_\mu(P | Q)$, where, given a probability μ over S , $H_\mu(P | Q)$ is the relative entropy of P with respect to Q . Note that $H_\mu(P | P) = 0$. That is, while learning P initially costs $H_\mu(P)$, learning P again in the subsequent period is free. Thus, by changing his choice of information, DM can gain expertise in one area while losing it in another.

The starkness of the examples above is useful in separating the notions of fatigue and expertise, as we will analyze in Section 5. In general, however, ICPs can

accommodate *any* dependence of the information constraint on the history of past information choices and state realizations, and our identification result in Section 4 applies to this very general class of constraints. There are many ways to generalize the last two examples. One way is to just combine them, in which case fatigue and expertise are both present and either may dominate. One could also accommodate partial expertise via the cost function $c(P | Q) = bH_\mu(P) + (1 - b)H_\mu(P | Q)$ for $b \in (0, 1)$.

3. Representation with Information Choice Processes

3.1. Domain

Let S be a finite set from which the payoff-relevant (objective) state realizes each period. For any compact metric space Y , we denote by $\Delta(Y)$ the space of Borel 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 .

There are $T > 1$ periods. 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 Anscombe-Aumann 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 define t -period problems as $X_t := \mathcal{K}(\mathcal{F}(\Delta(C \times X_{t-1})))$ for all $t = 2, \dots, T$. For any $x, y \in X_t$, where $t = 1, \dots, T$, and for any $\alpha \in [0, 1]$, we let $\alpha x + (1 - \alpha)y := \{\alpha f + (1 - \alpha)g : f \in x, g \in y\} \in X_t$.

Our domain X_T consists of T -period *dynamic choice problems* (henceforth, choice problems or *menus*). We analyze DM's preference relation \succeq over X_T .⁴

A consumption stream is a degenerate menu that does not involve choice at any point in time. Let $L_1 := \mathcal{F}(\Delta(C))$, and for each $t > 1$, define $L_t := \mathcal{F}(\Delta(C \times L_{t-1}))$. Thus, each $\ell \in L_t$ is an act that yields a state-dependent lottery over instantaneous consumption and a length- $(t - 1)$ continuation consumption stream (an $\ell' \in L_{t-1}$).

The space X_T subsumes some previously studied domains. For instance, if S is a singleton, X_T reduces to the domain in Kreps and Porteus (1978), which is extended to an infinite horizon in Gul and Pesendorfer (2004). The subspace of consumption

(4) We focus on the finite horizon for expositional clarity. Extending our model and results to the infinite horizon is conceptually straightforward, but formally a bit involved.

streams, where all choices are degenerate, is also a subspace of the (finite horizon version of the) domain in Krishna and Sadowski (2014).

3.2. ICP-Representation

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 *subjective information states*; the mapping $\Gamma : \Theta \rightarrow 2^{\mathcal{P}} \setminus \emptyset$ specifies the set of feasible partitions in any information state θ ; the transition operator $\tau : \mathcal{P} \times \Theta \times S \rightarrow \Theta$ determines the transition of the information state θ , given a particular choice of partition and the realization of an objective state; and θ_0 is the initial information state. Let \mathbf{M} be the space of ICPs. Note that the definition of \mathbf{M} does not take into account the finite horizon T , but it should be clear that for choice from X_T only the first T periods of an ICP are relevant.

In addition, let u be a real-valued, non-constant, and continuous function on C , and let $\delta \in (0, 1)$ be a discount factor.⁵ Let Π be the transition operator for a Markov process on S , where $\Pi(s, s') =: \pi_s(s')$ is the probability of transitioning from state s to state s' , and $\Pi(s, s') > 0$ for all $s, s' \in S$, ie, the process is *fully connected*. It is notationally convenient to let $0 \notin S$ be an auxiliary state, and to denote by π_0 the unique invariant measure of Π .




Figure 1: *Timeline*

Our model suggests the following timing of events and decisions, as illustrated in Figure 1. DM enters a period facing a menu $x \in X_t$, while being equipped with a prior belief π_s over S and an information state θ . He first chooses a partition $P \in \Gamma(\theta)$. For any realization of a cell $I \in P$, which includes the (previously determined) true state, DM updates his beliefs using Bayes' rule to obtain $\pi_s(\cdot | I)$, and then chooses an act $f \in x$. At the end of the period, the true state s' is revealed and DM receives

(5) For simplicity we assume that u is independent of the state s , but all our results can be extended to the case of state-dependent utilities.

the lottery $f(s')$, which determines current consumption c and continuation menu $y \in X_{t-1}$ for the next period. Concurrently, a new $\theta' = \tau(\theta, P, s')$ obtains and a new belief $\pi_{s'}$ is determined for the next period.

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

Definition 3.1. A preference \succsim on X_T has an ICP representation $(u, \delta, \Pi, \mathcal{M})$ if there is a collection of functions $V_t : X_t \times \Theta \times (S \cup \{0\}) \rightarrow \mathbb{R}$ for $t = 1, \dots, T$, such that each V_t satisfies

$$V_t(x_t, \theta, s) =$$

$$[3.1] \quad \max_{P \in \Gamma(\theta)} \sum_{J \in \mathcal{P}} \left[\max_{f \in x_t} \sum_{s' \in J} \mathbf{E}^{f(s')} \left[u(c) + \delta V_{t-1}(x_{t-1}, \tau(P, \theta, s'), s') \right] \pi_s(s' | J) \right] \pi_s(J)$$

with $V_0 = 0$, and $V_T(\cdot, \theta_0, 0)$ represents \succsim .

In the representation above, for each $s' \in S$, $f(s') \in \Delta(C \times X_{t-1})$ is a probability measure over $C \times X_{t-1}$, so that $\mathbf{E}^{f(s')}$ is the expectation over possible realizations (c, x_{t-1}) . Note that t in the representation does not index the time period, but rather the number of periods remaining until T .⁶

A *dynamic information plan* prescribes a choice of $P \in \Gamma(\theta)$ for each tuple (x_t, θ, s) (where s is the realized state in the previous period). Thus, an ICP describes the set of feasible information plans available to DM. Since $\Gamma(\theta)$ is finite, the representation [3.1] implies that an *optimal* dynamic information plan exists.

Before proceeding, we discuss two restrictions that the ICP representation in Definition 3.1 imposes. First, 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, because the same beliefs are used for the evaluation of future consumption acts (those acts that have no continuation values beyond their

(6) 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 describes behavior as the solution to a dynamic programming problem with state variables (x_t, θ, s) , so that implied behavior is dynamically consistent contingent on those state variables. The new aspect is that the state θ is controlled by DM and is not observed by the analyst.

instantaneous payoffs in a certain period), independently of the date of consumption.⁷ The interpretation is that DM does not learn the state in the period prior to the observed choice, and aggregates state-dependent preferences accordingly, using π_0 as his prior belief. One could instead assume that DM does learn the realization of the state in the period prior to his initial choice. Formally, this would simply mean replacing our primitive, \succsim , with a state-dependent family of initial preferences, $\{\succsim_s\}$. Since induced preferences in future periods are already state dependent, aggregated ex ante preferences can be thought of as an expositionally convenient summary of state-dependent preferences starting in the second period.⁸

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

3.3. Comparative Informativeness of ICPs

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

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

(7) Stationarity means that $\mathbf{c} \succsim \mathbf{c}'$ if, and only if for any $c_0 \in C$, we have $(c_0, \mathbf{c}) \succsim (c_0, \mathbf{c}')$, where \mathbf{c} means receiving c in every period (and in every state s) until T .

(8) Formally, if $\{\succsim_s\}$ is a family of preference relations over X_{T-1} , each of which admits a representation as in [3.1], then our identification result applies to each such \succsim_s . To see how \succsim can be used to summarize $\{\succsim_s\}$, note that \succsim induces state-dependent preferences over X_{T-1} as follows: Say that $x \succsim_s x'$ if $\{f_x\} \succsim \{f_{x'}\}$, where the act f_x yields some arbitrary but fixed consumption c in all states s and for all future periods, except in state s where it yields x as a continuation problem. In general, continuation preferences will depend on the unobservable initial choice of information. To sidestep this complication, suppose that $\Gamma(\theta_0) = \{S\}$, so that only the realization of the previous state is learned in the first period. Then $\{\succsim_s\}$ will have a representation as in [3.1], where the initial information state now depends on s .

$P' \in \Gamma'(\theta'_0)$, there exists $P \in \Gamma(\theta_0)$ such that P is finer than P' .

In turn, this suggests a natural extension to two-periods: \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 periods, we note that requirement (ii) above amounts to the continuation ICP $(\Theta, \tau(P, \theta_0, s), \Gamma, \tau)$ one-period Blackwell dominating $(\Theta', \tau'(P, \theta'_0, s), \Gamma', \tau')$. Similarly, we inductively define an order extending Blackwell dominance to t periods, whereby one ICP t -period Blackwell dominates another if for each information plan from the latter, there is another plan from the former that is more informative in the first period and, for any state realization s in the first period, leads to a more informative $(t - 1)$ -period plan starting in the second period.

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




Figure 2: Two-period ICPS, \mathcal{M} and \mathcal{M}'

As another example, consider two of the ICPS introduced in Example 2.3, \mathcal{M}^i for $i = a, b$, which differ only in the costs of acquiring information, that is, they can be parametrized by (K, κ, c^i) . It is easy to see that \mathcal{M}^a one period Blackwell dominates

\mathcal{M}^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}^a t -period Blackwell dominates \mathcal{M}^b if, and only if, $c^a \leq c^b$.

The t -period Blackwell order is reflexive and transitive. We say that \mathcal{M} strictly t -period Blackwell dominates \mathcal{M}' if \mathcal{M} t -period Blackwell dominates \mathcal{M}' , but not vice versa; and that they are t -period Blackwell equivalent if each dominates the other. Two Blackwell equivalent ICPS may differ because dominated information plans may be permitted by one but not the other. In addition, two ICPS that dominate each other may differ in terms of the parameters $(\Theta, \Gamma, \tau, \theta_0)$, but in Appendix A.3 we introduce the space of *canonical* ICPS, which abstracts from this arbitrariness in the parametrization. Blackwell dominance is consistent across periods in the sense that if \mathcal{M} strictly t -period Blackwell dominates \mathcal{M}' , then for $t' > t$, \mathcal{M}' cannot (weakly) t' -period Blackwell dominate \mathcal{M} (see Lemma B.1). We can, thus, define the *dynamic Blackwell order* over the space of ICPS as follows: \mathcal{M} dominates \mathcal{M}' in the dynamic Blackwell order if, for all $t \in \mathbb{N}$, \mathcal{M} t -period Blackwell dominates \mathcal{M}' . It is important to note that \mathcal{M} and \mathcal{M}' may be t -period Blackwell equivalent for all $t \leq T$, even though \mathcal{M} dynamically Blackwell dominates \mathcal{M}' . As a consequence, we will be able to identify the ICP only up to T -period Blackwell equivalence.

Greenshtein (1996) provides another dynamic extension of the static Blackwell order that is further analyzed in de Oliveira (2018). That extension compares (state-dependent) sequences of signals and so does not take into account information choice as we do. That our approach is well suited to our problem is demonstrated by our main identification result, Theorem 1, in the next section.

4. Unique Identification

Our main result states that all the parameters of our model are, essentially, uniquely identified.

Theorem 1. *Let $(u, \delta, \Pi, \mathcal{M})$ be an ICP representation of \succsim . Then, the function u is unique up to a positive affine transformation, δ and Π are unique, and \mathcal{M} is unique up to T -period Blackwell equivalence.⁹*

All the results in this section are proved in Appendix C. On the subdomain L_T , V_T satisfies Independence (that is, for any $\ell, \ell', \hat{\ell} \in L_T$ and any $\alpha \in (0, 1)$, $\ell \succsim \ell'$ if,

(9) In other words, for any additional representation of \succsim with parameters $(u^\dagger, \delta^\dagger, \Pi^\dagger, \mathcal{M}^\dagger)$, it is the case that $\delta^\dagger = \delta$, $\Pi^\dagger = \Pi$, $u^\dagger = au + b$, for some $a > 0$ and $b \in \mathbb{R}$, and \mathcal{M} and \mathcal{M}^\dagger dominate each other in the T -period Blackwell order.

and only if, $\alpha\ell + (1 - \alpha)\hat{\ell} \succeq \alpha\ell' + (1 - \alpha)\hat{\ell}'$ as it is independent of \mathcal{M} , and is thus completely characterized by the parameters (u, δ, Π) . By adapting the arguments in Krishna and Sadowski (2014, Corollary 5), it can be shown that such a representation on L_T is unique up to the addition of constants and a common scaling of u . Our challenge is to identify the ICP \mathcal{M} . In Section 4.1 we discuss our identification strategy.

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

Consider two decision makers with preferences \succeq and \succeq^\dagger , respectively. We say that \succeq has a *greater affinity for dynamic choice* than \succeq^\dagger if for all $x \in X$ and $\ell \in L$, $x \succeq^\dagger \ell$ implies $x \succeq \ell$.¹⁰ The comparison in the definition implies that \succeq and \succeq^\dagger have the same ranking over consumption streams in L_T .¹¹ While consumption streams require no choice of information, a typical choice problem x may allow DM to wait for information to arrive over multiple periods before making a choice. This option should be more valuable the more information plans DM's ICP renders feasible. The uniqueness established in Theorem 1 allows us to formalize this intuition.

Proposition 4.1. Let $(u, \delta, \Pi, \mathcal{M})$ and $(u^\dagger, \delta^\dagger, \Pi^\dagger, \mathcal{M}^\dagger)$ be ICP representations of \succeq and \succeq^\dagger respectively. The preference \succeq has a greater affinity for dynamic choice than \succeq^\dagger if, and only if, $\Pi = \Pi^\dagger$, $\delta = \delta^\dagger$, u and u^\dagger are identical up to a positive affine transformation, and \mathcal{M} T -period Blackwell dominates \mathcal{M}^\dagger .

Proposition 4.1 connects a 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 formalize and use repeatedly in the subsequent sections. A useful corollary of Proposition 4.1 is the following characterization of the dynamic Blackwell order: \mathcal{M} T -period Blackwell dominates \mathcal{M}^\dagger if, and only if, every discounted expected utility maximizer facing an arbitrary T -period menu prefers to have the ICP \mathcal{M} instead of \mathcal{M}^\dagger .¹²

-
- (10) 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).
 - (11) That is, $\ell \succeq \ell'$ if, and only if, $\ell \succeq^\dagger \ell'$ for all $\ell, \ell' \in L$. This follows from Lemma 34 in Appendix F of Krishna and Sadowski (2014) and since both \succeq and \succeq^\dagger satisfy Independence on L_T .
 - (12) This result thus generalizes the seminal characterization of the standard Blackwell order for partitions, according to which P is finer than Q if, and only if, every decision maker prefers P to Q regardless of the (static) choice problem.

4.1. Intuition for Identification of the ICP

We now illustrate the main idea behind the identification of the ICP. To simplify matters, suppose for the rest of this section that consumption is in the set $C = [0, 1]$ and u is strictly increasing, with $u(0) = 0$ and $u(1) = 1$. Rather than providing a general, more abstract intuition, we will base our discussion first on a static example, and then on an ICP that allows non-trivial information acquisition only in the first two periods; the same ideas extend to any finite horizon.

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

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

where, as before, \mathbf{c} means receiving $c \in C$ in every period until T , independently of the state. To test if DM is able to learn some partition that is weakly finer than P , consider the menu $x_1(P) := \{f_{1,J} : J \in P\}$. Contingent on cell $J \in P$, the act $f_{1,J}$ delivers permanent consumption 1 with certainty. Therefore, upon learning a cell in a partition Q that is weakly finer than P , the optimal consumption strategy will guarantee the same consumption, so that $V_1(x_1(P); Q) = 1 = V_1(x_1(P); P)$. Conversely, if Q is not finer than P , then with positive probability DM will learn a cell I which is not a subset of any $J \in P$. In that case, any choice of act from $x_1(P)$ will generate consumption 0 with positive probability, and hence $V_1(x_1(P); Q) < 1 = V_1(x_1(P); P)$.

In order to extend this intuition to two periods, recall the two ICPs, \mathcal{M} and \mathcal{M}' , respectively, in the left and right panels of Figure 2 (Section 3.3). Suppose the analyst believes DM's information constraint is either \mathcal{M} or \mathcal{M}' . How can she verify it is \mathcal{M} , and not \mathcal{M}' ? To answer this, consider the act $f_{2,\{S\}}(s) := (1, \text{Unif}\{x_1(P), x_1(Q)\})$ and the menu $x_2(\{S\}, \mathcal{M}) = \{f_{2,\{S\}}\}$. Note that $x_2(\{S\}, \mathcal{M})$ requires no choice in the first period, but instead offers the bet $\text{Unif}\{x_1(P), x_1(Q)\}$ that provides choice from either $x_1(P)$ or $x_1(Q)$ in the second period. To guarantee consumption of 1, DM must therefore have the option to choose in the second period whether to learn (at least) P or Q , which is not feasible under \mathcal{M}' . Therefore, $V_2(x_2(\{S\}, \mathcal{M}); \mathcal{M}) = 1 > V_2(x_2(\{S\}, \mathcal{M}); \mathcal{M}')$.

In order to fully identify \mathcal{M} , we also need to distinguish it from ICPs other than \mathcal{M}' . For instance, in order to test whether DM can learn the partition P twice in a row, let $f_{2,J}$ be the act that pays $f_{2,J}(s) = (1, x_1(P))$ if $s \in J$ and $(0, \mathbf{0})$ otherwise, and define the two-period choice problem $x_2(P) := \{f_{2,J} : J \in P\}$. An analogous construction for learning Q twice in a row yields $x_2(Q)$. It is easy to see that for

any $y \in \{x_2(P), x_2(Q), x_2(\{S\})\}$, $V_2(y, \mathcal{M}) = 1$. Moreover, for any ICP \mathcal{M}'' , we have that \mathcal{M}'' dynamically Blackwell dominates \mathcal{M} if, and only if, $V_2(y; \mathcal{M}'') = 1$ for all $y \in \{x_2(P), x_2(Q), x_2(\{S\})\}$. In essence, each such y amounts to a betting game, where in each period DM is told a random and history dependent partition and is asked to bet on the correct event in it to receive payoff 1 and stay in the game, rather than exiting the game and receiving 0 indefinitely. Such a betting game generates the same value as a constant stream of 1 if, and only if, DM can learn the relevant partition in each period. We will say that the collection of menus $\{x_2(P), x_2(Q), x_2(\{S\})\}$ is *aligned* with the ICP \mathcal{M} .

Indeed, an analogous construction of an aligned set of menus is possible for *any* ICP, as we demonstrate in Section 4.2 below, where we show that the intuition for identification above is a special case of the general result on aligned menus in Proposition 4.2.

Consider, finally, two ICPs that do not T -period Blackwell dominate each other, and the two corresponding sets of betting games. At least one of the two sets contains a game in which it will be possible to stay until T with certainty under the ICP the set is aligned with (generating the same value as consuming 1 until T) but not under the other ICP (generating a lesser value). In other words, the ICP in our model is identified up to the horizon T and up to T -period Blackwell dominance.

4.2. Aligned Menus

We now describe a duality between ICPs and a class of simple choice problems, which we call *aligned menus*. Aligned menus play a central rôle in our identification technique. To understand the intuition behind such menus, let us consider one-period problems. Fix a partition P . For each cell $J \in P$, define

$$f_{1,J,P}(s) := \begin{cases} c^+ & \text{if } s \in J \\ c^- & \text{otherwise} \end{cases} \quad \text{and} \quad x_1(P) := \{f_{1,J,P} : J \in P\}$$

where c^+ and c^- are, respectively, a u -best and a u -worst consumption choice. Notice that the menu $x_1(P)$ consists only of acts that are binary-valued. Moreover, $f_{1,J,P}$, and hence $x_1(P)$, do not depend on any other features of the utility function u .

The menu $x_1(P)$ represents a collection of bets such that any DM who has access to any partition at least as fine as P will be able to choose the act that gives him the

highest payoff. In other words, any DM with the partition P or finer, when faced with $x_1(P)$, effectively faces no uncertainty at all, as if he were able to make *omniscient* bets where he knows the payoff-relevant events *before* choosing an act.

We can extend this construction, recursively, for ICP $\mathcal{M} = (\Theta, \Gamma, \tau, \theta)$ and partition $P \in \Gamma(\theta)$. To start, let $x_1(P, \mathcal{M}) = x_1(P)$. Let \mathbf{c} denote the consumption stream that provides $c \in C$ in every period until T , independently of the state, and for $P \in \Gamma(\theta_0)$ and $t > 1$ recursively define

$$\Xi_{t-1}(P, \mathcal{M}, s) := \{x_{t-1}(Q, \mathcal{M}) : Q \in \Gamma(\theta'), \theta' = \tau(\theta, P, s)\}$$

and

$$f_{t,J,P}(s) := \begin{cases} \text{Unif}((c^+, \Xi_{t-1}(P, \mathcal{M}, s))) & \text{if } s \in J \\ \mathbf{c}^- & \text{otherwise} \end{cases}$$

where $\text{Unif}(\cdot)$ denotes the uniform distribution over a set. Let

$$x_t(P, \mathcal{M}) = \{f_{t,J,P} : J \in P\} \in X_t.$$

Just as with the one-period menu $x_1(P)$, the menu $x_t(P, \mathcal{M})$ represents a collection of multi-period bets such that any individual who has access to the ICP \mathcal{M} , or any ICP that Blackwell dominates it for the remaining $t - 1$ periods following the choice of P in the first period, again faces no uncertainty at all, and can bet as if he were omniscient.¹³ Let

$$\mathbf{M}_{\mathcal{P}} := \{(P, \mathcal{M}) \in \mathcal{P} \times \mathbf{M} : \mathcal{M} = (\Theta, \Gamma, \tau, \theta_0), P \in \Gamma(\theta_0)\}$$

denote an ICP with an initial feasible choice of partition. Note that $x_T(P, \mathcal{M})$ defines a mapping $x_T : \mathbf{M}_{\mathcal{P}} \rightarrow X_T$, which is an embedding of $\mathbf{M}_{\mathcal{P}}$ in X_T . There is a similar embedding of any t -period ICP with initial choice of partition into X_t .

Define now the mapping $\varphi_T : \mathbf{M} \rightarrow 2^{X_T}$ whereby

$$\varphi_T(\mathcal{M}) := \{x_T(P, \mathcal{M}) : \mathcal{M} = (\Theta, \Gamma, \tau, \theta_0), P \in \Gamma(\theta_0)\}$$

Given an ICP $\mathcal{M} = (\Theta, \Gamma, \tau, \theta_0)$, we say that \mathcal{M} and $\varphi_T(\mathcal{M})$ are *aligned*. For partition $P \in \Gamma(\theta_0)$ we also say that (P, \mathcal{M}) and $x_T(P, \mathcal{M})$ are aligned. Let $V_T(y, \theta_0, 0; \mathcal{M})$

(13) Instead of the uniform lotteries in the construction of $f_{t,J,P}$, the realization of the state $s \in J$ itself could be used for randomization. However, since for three or more states the cardinality of \mathcal{P} exceeds that of S , and hence of $J \in 2^S$, additional randomization in the form of lotteries may be needed for full identification of the ICP.

be the value of menu y under ICP \mathcal{M} as in the representation [3.1]. (For ease of notation, we will often suppress the arguments θ_0 and auxiliary state 0, and write this value as $V_T(y; \mathcal{M})$.) Notice that fixing (u, δ, Π) , each $\mathcal{M} \in \mathbf{M}$ defines a functional on X_T . Conversely, each $y \in X_T$ defines a functional on \mathbf{M} , both via the evaluation functional V_T . The following result formalizes our notion of alignment.

Proposition 4.2. For any $(P, \mathcal{M}), (P', \mathcal{M}') \in \mathbf{M}_{\mathcal{P}}$,

$$V_T(x_T(P, \mathcal{M}); \mathcal{M}) = V_T(\mathbf{c}^+; \cdot) \geq V_T(x_T(P', \mathcal{M}'); \mathcal{M})$$

Furthermore, the inequality is satisfied with equality for all P' with $(P', \mathcal{M}') \in \mathbf{M}_{\mathcal{P}}$ if, and only if, \mathcal{M} T -period Blackwell dominates \mathcal{M}' .

The proposition makes alignment in the sense of duality precise: For each $(P, \mathcal{M}) \in \mathbf{M}_{\mathcal{P}}$ there is a minimal menu (a menu where the payoff in every state is either c^+ or c^-), namely $x_T(P, \mathcal{M})$, that gives DM the highest possible payoff. To see the duality between P and $x_1(P, \mathcal{M}) = x_1(P)$ in the static setting, notice that

- $V_1(x_1(P); P) = V_1(\mathbf{c}^+; \cdot)$,
- $V_1(x_1(P); P) \geq V_1(y_1; \cdot)$, for all $y_1 \in X_1$, and
- $V_1(x_1(P); P) \geq V_1(x_1(P); Q)$ for all partitions, Q , with an equality if, and only if, Q is finer than P .

4.3. Inference from Limited Data

Our identification strategy suggests that inference about the ICP can be made from a small number of observations. In particular, inference benefits from three of its features. First, identification of the ICP is (almost) independent of the other preference parameters, as it only uses the best and worst outcomes c^+ and c^- (1 and 0 in the example discussed in Section 4.1).

Second, while identification of the ICP relies on randomization over continuation problems, the exact probabilities used in this randomization are not important; for example, we could replace the uniform distribution in the proof of Theorem 1 with any distribution with the same support (see Footnote 13).

Finally, Proposition 4.2 implies that it only takes finitely many comparisons to determine whether or not DM's ICP T -period Blackwell dominates a particular ICP \mathcal{M}' , and the number of comparisons required is $|\{(P', \mathcal{M}') : (P', \mathcal{M}') \in \mathbf{M}_{\mathcal{P}}\}|$ (ie, the number of feasible first period partitions available in \mathcal{M}'). Furthermore, to verify

whether DM can follow a *particular information plan*, the analyst need only observe *one* appropriate binary choice — between the best consumption stream and a choice problem that is strongly aligned with the plan in question. For instance, to establish as a lower bound whether DM is able to follow the information plan that chooses $\{S\}$ in the first period and from $\{P, Q\}$ in the second period, as in the ICP \mathcal{M} in Figure 2, it is enough to offer DM a choice between the betting game $x_2(\{S\}, \mathcal{M})$ and the consumption stream **1**. Indifference is observed if, and only if, DM can follow that information plan (or a more informative one). The last observation, in particular, can be useful in many applications. The example of job screening in Section 5.2 serves as an illustration.

5. Characterizing Expertise and Fatigue

In this section, we formally define and behaviorally characterize the notions of expertise and fatigue we described in Sections 1 and 2, making use of the concept of aligned menus from Section 4.2, and in particular Proposition 4.2.

5.1. Behavioral Characterization

Roughly speaking, *fatigue* occurs when the ability to learn decreases with previous choice of information. Given ICP $\mathcal{M} = (\Theta, \Gamma, \tau, \theta_0)$ for which (P_1, s_1, \dots, P_t) is part of an information plan that is undominated, we abuse notation and write $\Gamma(P_1, s_1, \dots, P_t, s_t)$ to denote the ICP that DM faces after the realizations of states (s_1, \dots, s_t) and the choice of partitions (P_1, \dots, P_t) . We write $Q \in \Gamma(P_1, s_1, \dots, P_t, s_t)$ if that continuation ICP allows the choice of partition Q next. We formally define fatigue as follows:

Definition 5.1. ICP \mathcal{M} displays *pure fatigue* if for any t , any sequence of states (s_1, \dots, s_t) , and any two sequences of partitions (P_1, \dots, P_t) and (P'_1, \dots, P'_t) where (i) $P_{t'}$ Blackwell dominates $P'_{t'}$ for all $t' \in \{1, \dots, t\}$; and (ii) $(P_1, s_1, \dots, P_t, s_t)$ and $(P'_1, s_1, \dots, P'_t, s_t)$ are each part of information plans that are undominated in \mathcal{M} , the following holds: If $Q \in \Gamma(P_1, s_1, \dots, P_t, s_t)$ then there is $Q' \in \Gamma(P'_1, s_1, \dots, P'_t, s_t)$ that Blackwell dominates Q .

The ICPs in Examples 2.2 and 2.3 display pure fatigue. In both examples, only the information choice of the previous period matters for the information constraint. For Example 2.2, comparing two histories of information choices as in the definition, (P_1, \dots, P_t) and (P'_1, \dots, P'_t) , there are three relevant cases: (i) P_t and P'_t are not

trivial and both make subsequently learning any non-trivial Q impossible; (ii) both P_t and P'_t are trivial, $P_t = P'_t = S$, and any partition can subsequently be learned; and (iii) P'_t is trivial, but P_t is not, in which case Q can be learned after (P'_1, \dots, P'_t) but not after (P_1, \dots, P_t) . In Example 2.3, if P_t Blackwell dominates P'_t , then it incurs a higher entropy cost and hence leads to a weakly tighter information constraint in the next period.

Turning to *expertise*, we first introduce some notation. For any state s and partition P , denote by P_s the cell in P that contains s . Let $\Upsilon(P) = \{(s, s') \in S \times S : P_s = P'_s\}$, that is, the set of all pair of states that are indistinguishable under P . Intuitively, $\Upsilon(P)$ measures the ignorance of the agent with partition P , and the larger it is, the more ignorant the agent is: If $\Upsilon(P) \subset \Upsilon(P')$, then P' is a coarsening of P and the agent is more ignorant with partition P' than with P . More generally, the difference $\Upsilon(P') - \Upsilon(P)$, which we refer to as the new information of P over P' , collects all the pairs between which P can distinguish while P' cannot, even if the latter is not a coarsening of the former.

Definition 5.2. ICP \mathcal{M} displays *pure expertise* if for any t , any given sequence of partitions and states $(P_1, s_1, \dots, P_t, s_t)$ that is part of an undominated information plan in \mathcal{M} , and any partitions Q and Q' with $\Upsilon(P_i) - \Upsilon(Q) \supset \Upsilon(P_i) - \Upsilon(Q')$ for all $i \leq t$, if $Q \in \Gamma(P_1, s_1, \dots, P_t, s_t)$, then there is $Q'' \in \Gamma(P_1, s_1, \dots, P_t, s_t)$ that Blackwell dominates Q' .

To see the intuition behind the definition, suppose that $t = 2$ and DM needs to choose between learning Q and learning Q' , neither of which dominates the other. Which one will be “easier” to learn depends on how much information those two partitions contain that is not also contained in P_1 , measured by the differences between the corresponding Υ sets; the smaller the difference, the easier it is for DM to exploit the expertise he obtains from previously leaning P_1 . In the extreme case where $Q' = P_1$, DM can automatically afford learning Q' .

The ICPS in Examples 2.1 and 2.4 display pure expertise. In Example 2.1, given the realized sequence of partitions and states $(P_1, s_1, \dots, P_t, s_t)$, the continuation constraint allows choosing one of a fixed number of refinements of the meet of P_1, \dots, P_t . Clearly, if $\Upsilon(P_i) - \Upsilon(Q) \supset \Upsilon(P_i) - \Upsilon(Q')$ for all $i \leq t$, then if Q can be learned, so can Q' . In Example 2.4, a similar reasoning applies, but here only the last partition, P_t matters for the availability of Q versus Q' .

In general, expertise and fatigue may occur simultaneously. For instance, fatigue may always occur, but may be over- or fully- or partially- compensated by expertise

when information is similar between periods; the extent to which expertise and fatigue interact depends on the particular ICP. In the definitions above, ‘pure fatigue’ implies that expertise can never over-compensate fatigue, and ‘pure expertise’ implies that fatigue can not over-compensate expertise. Those definitions, therefore, allow us to focus on the clear cut cases where either expertise or fatigue must be occurring.

We now characterize pure fatigue and pure expertise in terms of behavior. To this end, it is useful to consider menus that can be described via an ICPs they are aligned with. In what follows, we rely on the labels introduced in Definitions 5.1 and 5.2, respectively. Turning first to fatigue, let \mathcal{M}^* be an ICP that allows initially learning the sequence $(P_1, s_1, P_2, \dots, P_t, s_t, Q)$ followed by the coarsest partition $\{S\}$ forever thereafter. Let \mathcal{M}' allow initially learning the sequence (P'_1, s_1, \dots, P'_t) , and construct \mathcal{M}'' from \mathcal{M}' by adding the option to continue the sequence $(P'_1, s_1, \dots, P'_t, s_t)$ with learning Q and then nothing thereafter. Note that this option may be dominated by some option already available under \mathcal{M}' .

Proposition 5.3. Suppose the preference \succeq has an ICP representation. Then, \succeq displays pure fatigue if, and only if, for all \mathcal{M}^* , \mathcal{M}' and \mathcal{M}'' as defined above, $x_T(P_1, \mathcal{M}^*) \sim \mathbf{c}^+$ and $x_T(P'_1, \mathcal{M}'') \not\sim \mathbf{c}^+$ imply that either (i) $x_T(P'_1, \mathcal{M}') \not\sim \mathbf{c}^+$, or (ii) there is \mathcal{M}''' that strictly dominates \mathcal{M}' following the choice of partition P'_1 with $x_T(P'_1, \mathcal{M}''') \sim \mathbf{c}^+$.

The proofs of all the results in this section are in Appendix D. In essence, if (P'_1, \mathcal{M}'') is not a feasible strategy, then neither is (P'_1, \mathcal{M}') , and hence $x_T(P'_1, \mathcal{M}')$ cannot be in the set of menus that is strongly aligned with the true ICP \mathcal{M} .

Turning to pure expertise, let \mathcal{M}^* be as above, but now let \mathcal{M}' allow initially learning the sequence (P_1, s_1, \dots, P_t) , and construct \mathcal{M}'' from \mathcal{M}' by adding the option to continue the sequence $(P_1, s_1, \dots, P_t, s_t)$ with learning Q' and then nothing thereafter.

Proposition 5.4. Suppose the preference \succeq has an ICP representation. Then, \succeq displays pure expertise if, and only if, $x_T(P_1, \mathcal{M}^*) \sim \mathbf{c}^+$ and $x_T(P_1, \mathcal{M}'') \not\sim \mathbf{c}^+$ imply that either (i) $x_T(P_1, \mathcal{M}') \not\sim \mathbf{c}^+$, or (ii) there is \mathcal{M}''' that strictly dominates \mathcal{M}' following the choice of partition P_1 with $x_T(P_1, \mathcal{M}''') \sim \mathbf{c}^+$.

In the context of the proposition, betting on Q is more familiar than betting on Q' after previously betting on $(P_1, s_1, \dots, P_t, s_t)$, and according to the proposition, DM will thus weakly prefer to bet on Q . More generally, expertise 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.¹⁴

5.2. An Example of Job Screening

We now illustrate how expertise and fatigue can influence behavior in a simple two-period example of job choice, and how this can be used by an employer to successfully screen candidates for different kinds of tasks.

Consider a two-dimensional state space, $S := \mathbb{H} \times \mathbb{V}$, with vertical component $\mathbb{V} := \{u, d\}$ and horizontal component $\mathbb{H} := \{\ell, m, r\}$. Suppose again that $u(1) > u(0)$, and simplify notation by writing f_J instead of $f_{1,J,P}$ for the act that pays 1 on event $J \subset S$ and 0 otherwise.

There are three types of individuals who differ in their two-period ICPs.

- I. The first group consists of individuals who are able to concentrate and perform very highly over a short period of time, after which some rest (or play) is needed. In our model, this can be captured by agents being able to learn the state perfectly in any one period, but subsequently needing a period without learning anything, as in Example 2.2. These individuals display pure fatigue according to Definition 5.1.




Figure 3: ICP for Type I individuals

- II. The second group engages in ‘learning by doing’. In our model, this can be captured by agents who can gain the expertise to precisely identify the h component of the state by first coarsely learning the h component in the previous period, in line with Example 2.1. These individuals display pure expertise in the sense of Definition 5.2
- III. The third group describes the generalists: individuals who have broad but limited skills. In our model, this can be captured by agents who can learn a bipartition of

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




Figure 4: ICP for Type II individuals

either \mathbb{V} or \mathbb{H} in any period, independent of the previous information choice. These agents display neither fatigue nor expertise.




Figure 5: ICP for Type III individuals

Note that none of these three ICPS can be ranked in terms of the dynamic Blackwell order.

Consider a firm that wants to hire employees for three different types of jobs. Each job requires making different types of (possibly trivial) decisions in the two periods, and the correctness of decisions is measurable in S . Suppose the firm offers performance wages: it rewards correct decisions by paying 1 and does not reward wrong decisions.

- M.** Managers must make one fully informed decision, which is equally likely to occur in either period 1 or 2, but not both. The performance wage M for this job corresponds to a uniform lottery over two two-period problems¹⁵

$$M = \text{Unif} \{ (\{f_s\}_{s \in S}, 1), (1, \{f_s\}_{s \in S}) \}$$

- E.** Engineers are concerned only with the vertical dimension (\mathbb{V}) and must make a coarse binary decision in period 1 followed by a precise decision in period 2, such that their performance wage amounts to the two-period problem

$$E = (\{f_\ell, f_{\{m, r\}}\}, \{f_\ell, f_m, f_r\})$$

- A.** Administrators are presented each period with a binary decision that may depend on either the \mathbb{V} or the \mathbb{H} dimension, but not on both, such that their performance

(15) Recall that $\mathbf{1}$ represents the constant act that always pays 1 and $\mathbf{1}$ denotes the constant stream that pays 1 in each of the two periods.

wage amounts to the two-period problem

$$A = \left(\text{Unif} \left\{ \{f_u, f_d\}, \{f_\ell, f_{\{m,r\}}\} \right\}, \text{Unif} \left\{ \{f_u, f_d\}, \{f_\ell, f_{\{m,r\}}\} \right\} \right)$$

These performance wages will allow the firm to perfectly screen individuals for the respective positions, as individuals of group [I] select management positions, where they always make the correct decision¹⁶

$$M \sim_{\text{I}} \mathbf{1} >_{\text{I}} rE, A$$

while individuals of group [II] choose to become engineers¹⁷

$$E \sim_{\text{II}} \mathbf{1} >_{\text{II}} M, A$$

and group [III] individuals opt to become administrators

$$A \sim_{\text{III}} \mathbf{1} >_{\text{III}} M, E$$

While screening models where types are ordered (say by the value of a single parameter) have nice properties, solving multi-dimensional screening models where types are not naturally ranked is more difficult, and often does not provide robust insights. While ICPS can be very complicated (high dimensional objects), the example in this section illustrates that there is an analogy between the inequalities that derive lower bounds on the information an individual can process from their valuation of aligned menus, and the collection of incentive compatibility constraints that can be used to screen different ICPS.

6. Related Literature

There are negative results in the econometric literature about the identifiability of subjective Markov decision processes. For instance, Rust (1994) and Magnac and

-
- (16) To see how this behavior aligns with the prediction of Proposition 5.3, let \mathcal{M}^* be the ICP in Figure 3, with $P_1 = \{s\}_{s \in S}$ and $Q = \{S\}$, and let \mathcal{M}' be the ICP in Figure 4 with $P'_1 = \{\ell, \{m, r\}\}$, so that P_1 Blackwell dominated P'_1 . Note that \mathcal{M}'' is again just the ICP in Figure 4, as Q is dominated by $\{\ell, m, r\}$. We have $x_2(P_1, \mathcal{M}^*) = M \sim_{\text{I}} \mathbf{1}$ and $x_2(P'_1, \mathcal{M}'') \not\sim_{\text{I}} \mathbf{1}$, but trivially also $x_2(P'_1, \mathcal{M}') = E \not\sim_{\text{I}} \mathbf{1}$.
- (17) To see how this behavior aligns with the prediction of Proposition 5.4, let \mathcal{M}^* be the ICP in Figure 4, with $P_1 = \{\ell, \{m, r\}\}$ and $Q = \{\ell, m, r\}$, let \mathcal{M}' be the ICP in Figure 5, and let $Q' = \{\{\ell, m\}, r\}$. Then \mathcal{M}'' is the ICP in Figure 5 with the addition that learning Q' is now possible following P_1 . Note that $x_2(P_1, \mathcal{M}^*) = E \sim_{\text{I}} \mathbf{1}$ and $x_2(P_1, \mathcal{M}'') \not\sim_{\text{I}} \mathbf{1}$, but also $x_2(P_1, \mathcal{M}') = A \not\sim_{\text{I}} \mathbf{1}$.

Thesmar (2002) show that in a general Markov decision processes, where utilities depend on the Markov state and choice affects the stochastic evolution of that state, observing (stochastic) choice is insufficient to identify the evolution of shocks or other parameters of the model. The latter proceed to provide structural assumptions that allow identification. More recently, Hu and Shum (2012) provide structural assumptions under which it is possible to identify a Markov process with a discrete set of unobservable state variables from panel choice data under the maintained assumption that all control variables are observable by the analyst, as for instance in the job matching model of Miller (1984) or the real option (optimal stopping) problem in Pakes (1986).

Recall that the Markov state in our model is (x, θ, s) . In some sense, our identification problem is more ambitious than the ones mentioned above, because we have unobservable control variables: θ cannot be observed by the analyst, and the set of available information choices given θ as well as the transition of θ as a function of the information choice are both unknown. Crucially for us, the set of available observable actions in Markov state (x, θ, s) depends only on x , and the analyst can effectively observe choice from all possible continuation problems for the same combination of s and θ . This contrasts with the aforementioned econometric literature where the analyst does not have access to a rich set of tradeoffs at a given state. For instance, even though in Magnac and Thesmar (2002) the evolution of the Markov state is observable by the analyst, the distribution of (iid) taste shocks cannot be identified because DM always has the same set of actions to choose from. Similarly, in Rust (1994), the set of feasible actions is completely determined by the Markov state, and therefore precludes observations from a rich set of continuation problems in a given Markov state. It is this ability to observe choices from a rich set of alternatives at any state that is essential for identification. We note that this feature is orthogonal to the fact that we consider deterministic ex ante menu choice instead of choice frequencies over time. That in principle it may be possible to make inferences about unobserved random variables either by looking at choice frequencies from enough choice sets or by looking at ex ante menu choice was first observed in a static context by Gul and Pesendorfer (2006).

In the decision theoretic literature, de Oliveira et al. (2017) provide an identification result for static information constraints in the face of subjective uncertainty.¹⁸

(18) de Oliveira et al. (2017) permit more general information structures than partitions, and also allow for explicit costs of acquiring information.

A dynamic extension of their model, in which DM faces the same information constraint each period (independent of past information choices), corresponds to an ICP representation where the ICP \mathcal{M} has a degenerate (singleton) subjective state space $\Theta = \{\theta_0\}$. Since this representation is a special case of the ICP representation, our identification result obviously applies. In fact, identifying this representation does not rely on continuation choice problems in periods $t > 1$. To see this, recall that the parameters (u, δ, Π) can be identified from preference over consumption streams $\ell \in L_T$. Further note that, as in de Oliveira et al. (2017), identifying the fixed set of partitions up to Blackwell equivalence requires non-trivial choice only in one period, for instance in $t = 1$.

Krishna and Sadowski (2014) identify the subjective flow of information in a dynamic model in which DM is uncertain about a subjective state, but takes the flow of information as given, rather than facing a constrained choice. 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_T , as does ours. The key difference between the two domains lies in the timing of events: Instead of acts over menus of lotteries as they do, we consider menus of acts over lotteries, which are appropriate for our dynamic extension of de Oliveira et al. (2017). Finally, Piermont, Takeoka, and Teper (2016) identify a different type of dynamic information constraint, where a decision maker learns about his uncertain, but time invariant, consumption taste (only) through consumption.

The decision theoretic literature has also provided testable behavioral foundations for models related to ours. In a static environment, Dillenberger et al. (2014) show that uncertainty about future beliefs about the objective state of the world corresponds to preference for flexibility over menus of acts. de Oliveira et al. (2017) model subjective uncertainty that is not fixed but a hidden choice variable by replacing the Independence axiom in Dillenberger et al. (2014) with Aversion to Randomization. In a dynamic context, Krishna and Sadowski (2014) provide axioms for the model where the flow of information about the subjective state of the world is taken as given by DM. In a companion paper, Dillenberger, Krishna, and Sadowski (2021) build on this literature to provide an axiomatic foundation for a fully recursive infinite horizon extension of our model, where DM controls the flow of information over time. Due to the intertemporal information constraint, preferences are interdependent across time and do not satisfy the stationarity assumptions of Krishna and Sadowski (2014). To deal with this complication, Dillenberger, Krishna, and Sadowski (2021) rely on a

recursive application of their axioms, which can be compared to standard stationarity assumptions: Stationarity requires instantaneous and continuation preferences to be identical, while the recursive axiomatization merely requires them to be of the same class. Importantly, since the Markov state in our model is not directly observable, the axiomatization must be based on inferring this state from preferences. Loosely speaking, this relies on the identification strategy developed in this paper, which can be extended to the infinite horizon.

Appendices

A. Preliminaries

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

A.1. Metrics on Probability Measures

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

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

A.2. Recursive Domain

Let $X_1 := \mathcal{K}(\mathcal{F}(\Delta(C)))$ be the space of menus of acts that pay out lotteries over C . Intuitively, X_1 consists of all one-period Anscombe-Aumann (AA) choice problems. For acts $f^1, g^1 \in \mathcal{F}(\Delta(C))$, define the metric $d^{(1)}$ on $\mathcal{F}(\Delta(C))$ by $d^{(1)}(f^1, g^1) := \max_s d_D(f^1(s), g^1(s)) \leq 1$. 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).

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

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

Now define recursively, for $t > 1$, $X_t := \mathcal{K}(\mathcal{F}(\Delta(C \times X_{t-1})))$. The metric on $C \times X_{t-1}$ is the product metric, that is, $d_{C \times X_{t-1}}((c, x_{t-1}), (c', x'_{t-1})) = \max[d_C(c, c'), d^{(t-1)}(x_{t-1}, x'_{t-1})]$. This induces the Dudley metric on $\Delta(C \times X_{t-1})$.

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

$$d_H^{(t)}(x_t, y_t) := \max \left\{ \max_{f^t \in x_t} d^{(t)}(f^t, y_t), \max_{g^t \in y_t} d^{(t)}(g^t, x_t) \right\}$$

Here, X_t consists of all t -period AA choice problems. The agent faces a menu of acts which pay off in lotteries over consumption and $(t - 1)$ -period AA choice problems that begin the next period. Setting $t = T$ gives us the domain X_T .

A.3. Canonical Information Choice Processes

As mentioned in Section 3.3, in terms of behavior, all that is relevant about the ICP \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 (pseudo-) metrize the set of all ICPs \mathbf{M} , and to provide a canonical space of ICPs.

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 information states, and up to the addition of information states that can never be reached under any information plan, *given the initial state* θ_0 . This definition of indistinguishability is formalized below and leads to the recursive characterization described in Lemma A.2.

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 realization of state s results in a new ICP $(\Theta, \Gamma, \tau, \tau(\theta_0, P, s))$. To simplify notation, we denote this new ICP by $\mathcal{M}(\tau(\theta_0, P, s))$. Further abusing notation, $\mathcal{M}(\theta)$ denotes the ICP $(\Theta, \Gamma, \tau, \theta)$ with initial state θ . Define $D : \mathbf{M} \times \mathbf{M} \rightarrow \mathbb{R}$ as:

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

We endow Θ with the discrete metric, which means that the Hausdorff distance $d_H(A, B) \leq 1$ for all $A, B \subset \Theta$. The function D captures the discrepancy between

\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 $\mathbb{T} : \mathbf{B}(\mathbf{M} \times \mathbf{M}) \rightarrow \mathbf{B}(\mathbf{M} \times \mathbf{M})$ defined as

$$\mathbb{T}D'(\mathcal{M}(\theta_0), \mathcal{M}'(\theta'_0)) := \max \left\{ d_H(\Gamma(\theta_0), \Gamma'(\theta'_0)), \frac{1}{2} \max_{P \in \Gamma(\theta_0), s \in S} D'(\mathcal{M}(\tau(\theta_0, P, s)), \mathcal{M}'(\tau'(\theta'_0, P, s))) \right\}$$

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

We can now define an isomorphism between ICPs. In terms of the discrepancy function D , two ICPs \mathcal{M} and \mathcal{M}' are *indistinguishable* if $D(\mathcal{M}(\theta_0), \mathcal{M}'(\theta'_0)) = 0$. The definition of D immediately implies the following recursive characterization of indistinguishability whose proof is omitted.

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 $\mathcal{M}(\tau(\theta_0, P, s))$ is indistinguishable from the ICP $\mathcal{M}'(\tau'(\theta'_0, P, s))$.

We now construct a set of *canonical* ICPs. Recall that \mathcal{P} is the space of all partitions of S , where a typical partition is P . Then, (\mathcal{P}, d) is a metric space, where d is the discrete metric.

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 $n > 1$, $\Omega_n := \mathcal{K}_b(\mathcal{P} \times \Omega_{n-1}^S)$. Set $\Omega' := \times_{n=1}^{\infty} \Omega_n$. 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 the function (because Ω_n is a space of sets) $\Psi_n : \Omega_{n+1} \rightarrow \Omega_n$ 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 Ω is the set of all consistent elements of Ω' ,

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

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.)

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

Theorem 2. *The set Ω is homeomorphic to $\mathcal{K}_b(\mathcal{P} \times \Omega^S)$.*

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

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 . Also note 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. To show (a), 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 can be verified that the sequence $(\omega_1, \omega_2, \dots, \omega_n, \dots) \in \times_{n \in \mathbb{N}} \Omega_n$ 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 . Observe that if $\omega' \in \Omega$ is indistinguishable from \mathcal{M} , then it must also be indistinguishable from $\omega_{\mathcal{M}}$ (because D is a metric), which proves that $\omega_{\mathcal{M}} = \omega$.

To show (b), 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

B. T-Period Dynamic Blackwell Order

In this section, we construct the T -period dynamic Blackwell order for canonical ICPs. Appendix A.3 exhibits an isomorphism between canonical ICPs and ICPs. It is immediate to verify that the isomorphism induces the T -period dynamic Blackwell order on \mathbf{M} , the space of ICPs, as defined in the text.

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 \succeq_0 on $\hat{\Omega}_0$ as follows: $\omega_0 \succeq_0 \omega'_0$ 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 $n \geq 1$, \succeq_n on $\hat{\Omega}_n$: For all $\omega_n, \omega'_n \in \hat{\Omega}_n$, $\omega_n \succeq_n \omega'_n$ 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) $\omega_{n-1,s} \succeq_{n-1} \omega'_{n-1,s}$ for all $s \in S$.

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

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

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 $\omega_0 \succeq_0 \omega'_0$. 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 \succeq_0 is reflexive, $\hat{\omega} \succeq_0 \hat{\omega}$. But this implies $\omega_0 \succeq_1 \omega'_0$. Conversely, let $\omega_0 \succeq_1 \omega'_0$. 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) $\hat{\omega} \succeq_0 \hat{\omega}$ for all $s \in S$. But this implies $\omega_0 \succeq_0 \omega'_0$, which proves that $\succeq_{n+1} |_{\hat{\Omega}_n} = \succeq_n$ when $n = 0$.

As our inductive hypothesis, we suppose that $\succeq_n |_{\hat{\Omega}_{n-1}} = \succeq_{n-1}$. Let $\omega_n \succeq_n \omega'_n$. 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) $\tilde{\omega}_{n-1,s} \succeq_{n-1} \tilde{\omega}'_{n-1,s}$ for all $s \in S$. But by the induction hypothesis, this is equivalent to $\tilde{\omega}_{n-1,s} \succeq_n \tilde{\omega}'_{n-1,s}$ for all $s \in S$, which implies that $\omega_n \succeq_{n+1} \omega'_n$.

Conversely, let $\omega_n \succeq_{n+1} \omega'_n$. 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) $\tilde{\omega}_{n-1,s} \succeq_n \tilde{\omega}'_{n-1,s}$ for all $s \in S$. However, the induction hypothesis implies $\tilde{\omega}_{n-1,s} \succeq_{n-1} \tilde{\omega}'_{n-1,s}$ for all $s \in S$, proving that $\omega_n \succeq_n \omega'_n$ and therefore $\succeq_{n+1} |_{\hat{\Omega}_n} = \succeq_n$. \square

The T -period dynamic Blackwell order on \mathbf{M} is \succeq_T as defined above, where we set $n = T$.

C. Identification and Behavioral Comparison: Proofs from Section 4

Based on the previous results and notation, we first establish two lemmas that will be used to prove Theorem 1, Proposition 4.1, and Proposition 4.2.

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

Recall that \mathbf{c}^+ and \mathbf{c}^- are the best and worst consumption streams, respectively, that deliver c^+ and c^- in any period and state. For a partition P with generic cell J , define the act

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

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

Now, proceed inductively, and for $t \geq 2$, suppose we have the menu $x_{t-1}(P, \omega')$ for each $(P, \omega') \in \omega$, and define, for each cell $J \in P$, the act

$$f_{t,J}(s) := \begin{cases} (c^+, \text{Unif}_{t-1}(\{x_{t-1}(Q, \tilde{\omega}) : (Q, \tilde{\omega}) \in \omega'_s\})) & \text{if } s \in J \\ \mathbf{c}^- & \text{if } s \notin J \end{cases}$$

Then, given $(P, \omega') \in \omega$, we have the menu $x_t(P, \omega')$ defined as

$$[\star] \quad x_t(P, \omega') := \{f_{t,J} : J \in P\}$$

Note that

$$V_T(\mathbf{c}^+; \omega) = V_T(x_T(P, \omega'); \omega) \geq V_T(x_T(P, \omega'), \tilde{\omega})$$

for all $\tilde{\omega} \in \Omega$, with equality if $\tilde{\omega}$ dynamically Blackwell dominates ω . This amounts to the first part of Proposition 4.2, as well as the "if" direction of the second part. The "only if" part is established by the sequence of Lemmas C.1 and C.2 below.

Lemma C.1. Let $P, Q \in \mathcal{P}$ and suppose Q is not finer than P . Then, for any $\omega, \omega' \in \Omega^S$, $V_T(x_T(P, \omega); (P, \omega)) > V_T(x_T(P, \omega); (Q, \omega'))$.

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

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

Proof. Suppose ω^2 does not t -period dynamically Blackwell dominate ω^1 . Then, there exists a smallest $n \geq 1$ such that for all $m < n$, $\omega^2 \succeq_m \omega^1$, while $\omega^2 \not\succeq_n \omega^1$.

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

Let $(P_1, \tilde{\omega}) \in \omega^1$ be the unique first period choice under ω^1 that makes the sequence (P_k) feasible. Then $x_T(P_1, \tilde{\omega})$ 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 C.1, which completes the proof. \square

Proof of Theorem 1. For the case of finitely many prizes (ie, when C is a finite set), Corollary 5 of Krishna and Sadowski (2014) establishes that the collection (u, Π, δ) is unique in the sense of the Theorem. While we cannot directly appeal to their result, judicious and repeated applications of their corollary allows us to reach the same conclusion for a compact set of prizes. Now, define $F_{T, \omega} := \{x_T(P, \tilde{\omega}) : (P, \tilde{\omega}) \in \omega\}$. It follows immediately from Lemma C.2 that $F_{T, \omega}$ is aligned with ω . \square

This allows us to characterize the T -period Blackwell order in terms of the instrumental value of information.

Corollary C.3. Let $\omega, \omega' \in \Omega$. Then, the following are equivalent.

- (a) ω T -period Blackwell dominates ω' .
- (b) For any (u, Π, δ) that induces $\omega \mapsto V_T(\cdot; \omega)$, $V_T(x; \omega) \geq V_T(x; \omega')$ for all $x \in X_T$.

Proof. That (b) implies (a) is merely the contrapositive to Lemma C.2. To see that (a) implies (b), note first the claim holds when $T = 1$ by Theorem 1 on p59 of Laffont (1989). Intuitively, in a static optimization problem, more information (via a finer partition) is better. Now, pick any information plan from ω' . By assumption, there is an information plan in ω such that the partition chosen in any period is finer than the partition chosen under ω' . But this implies that any consumption strategy for the menu x using the information plan from ω' is also feasible under the information plan from ω . Because information only has instrumental value, it follows that $V_T(x; \omega) \geq V_T(x; \omega')$ for all $x \in X_T$. \square

We are now in a position to prove Proposition 4.1.

Proof of Proposition 4.1. We first show the ‘only if’ part. On L_T , we have $\ell \succeq^\dagger \ell'$ implies $\ell \succeq \ell'$. This implies, by Lemma 34 of Krishna and Sadowski (2014), that $\succeq^\dagger|_{L_T} = \succeq|_{L_T}$. Together with the uniqueness of the Recursive Anscombe-Aumann (RAA) representation (see Corollary 5 in Krishna and Sadowski (2014)), this implies that $(u, \delta, \Pi) = (u^\dagger, \delta^\dagger, \Pi^\dagger)$ after a suitable (and behaviorally irrelevant) normalization of the utilities. Thus, part (b) of Corollary C.3 holds, which establishes the claim.

The ‘if’ part follows immediately from Corollary C.3. \square

We now prove Proposition 4.2.

Proof of Proposition 4.2. That $V_T(x_T(P, \mathcal{M}); \mathcal{M}) = V_T(\mathbf{c}^+; \cdot) \geq V_T(x_T(P', \mathcal{M}'); \mathcal{M})$ follows from the definition of \mathbf{c}^+ and $x_T(P, \mathcal{M})$. All that remains to be shown is that the inequality holds as an equality for all P' with $(P', \mathcal{M}') \in \mathbf{M}_\mathcal{P}$ if, and only if, \mathcal{M} T -period Blackwell dominates \mathcal{M}' .

To see the ‘only if’ part, suppose that \mathcal{M} does *not* T -period Blackwell dominate \mathcal{M}' . Then, by Lemma C.2, it follows that for some P' with $(P', \mathcal{M}') \in \mathbf{M}_\mathcal{P}$, we have $V_T(x_T(P', \mathcal{M}'); \mathcal{M}') = V_t(\mathbf{c}^+; \cdot) > V_T(x_T(P', \mathcal{M}'); \mathcal{M})$, which establishes the contrapositive.

To see the ‘if’ part, suppose \mathcal{M} T -period Blackwell dominates \mathcal{M}' . Then, for every information plan in \mathcal{M}' , and every P' with $(P', \mathcal{M}') \in \mathbf{M}_\mathcal{P}$, there is another information plan in \mathcal{M} with a finer partition after every realization of menu, state,

and date. Following this information plan gives DM the same utility as consuming \mathbf{c}^+ for sure. \square

D. Proofs from Section 5.1

Proof of Proposition 5.3. We first show that the behavioral comparison in the proposition implies fatigue as in Definition 5.1. Given the true underlying ICP \mathcal{M} , take (s_1, \dots, s_t) , and two sequences of partitions (P_1, \dots, P_t) and (P'_1, \dots, P'_t) as in the definition of fatigue, so that $Q \in \Gamma(P_1, s_1, \dots, P_t, s_t)$. Let \mathcal{M}^* be the ICP that allows learning $(P_1, s_1, \dots, P_t, s_t, Q)$ and nothing else. Then $x_T(P_1, \mathcal{M}^*) \sim \mathbf{c}^+$. Let $\mathcal{M}' = \mathcal{M}$. By assumption, \mathcal{M} allows learning (P'_1, s_1, \dots, P'_t) , so that the behavioral comparison applies. Note that $x_T(P'_1, \mathcal{M}') \sim \mathbf{c}^+$ and there is no \mathcal{M}''' that strictly dominates \mathcal{M}' after P'_1 with $x_T(P'_1, \mathcal{M}''') \sim \mathbf{c}^+$. Construct \mathcal{M}'' from \mathcal{M}' by adding the option to continue the sequence $(P'_1, s_1, \dots, P'_t, s_t)$ with learning Q . Now suppose that there is no $Q' \in \Gamma(P'_1, s_1, \dots, P'_t, s_t)$ such that Q' Blackwell dominates Q . Then $x_T(P'_1, \mathcal{M}'') \not\sim \mathbf{c}^+$, a contradiction. Hence, there is $Q' \in \Gamma(P'_1, s_1, \dots, P'_t, s_t)$ that dominates Q .

We now show that fatigue as in Definition 5.1 implies the behavioral comparison in the proposition. For the ICPs in the definition, $x_T(P_1, \mathcal{M}^*) \sim \mathbf{c}^+$ implies that a path that pointwise Blackwell dominates $(P_1, s_1, P_2, \dots, s_t, Q)$ can be learned initially. There are three cases to consider:

- (a) \mathcal{M}' and the true ICP \mathcal{M} dynamically Blackwell dominate each other. Then by fatigue \mathcal{M}'' must also be dominated by \mathcal{M} , and hence $x_T(P'_1, \mathcal{M}'') \sim \mathbf{c}^+$, a contradiction to the assumption that $x_T(P'_1, \mathcal{M}'') \not\sim \mathbf{c}^+$. This rules out the first case.
- (b) \mathcal{M} does not dynamically Blackwell dominate \mathcal{M}' , and hence $x_T(P'_1, \mathcal{M}') \not\sim \mathbf{c}^+$, which is property (i) in the behavioral comparison.
- (c) \mathcal{M} strictly dynamically Blackwell dominates \mathcal{M}' . In that case $x_T(P'_1, \mathcal{M}) \sim \mathbf{c}^+$ which is property (ii) of the behavioral comparison for $\mathcal{M}''' = \mathcal{M}$.

Because the three cases are exhaustive, the proposition is proved. \square

Proof of Proposition 5.4. We first show that the behavioral condition in Proposition 5.4 implies expertise as in Definition 5.2. Given the true ICP \mathcal{M} in the representation, take (s_1, \dots, s_t) , a sequences of partitions (P_1, \dots, P_t) and partitions Q and Q' as in the definition of expertise, so that $Q \in \Gamma(P_1, s_1, \dots, P_t, s_t)$. Let \mathcal{M}^* be an ICP that corresponds to learning $(P_1, s_1, \dots, P_t, s_t, Q)$ and nothing else. Then $x_T(P_1, \mathcal{M}^*) \sim \mathbf{c}^+$. Let $\mathcal{M}' = \mathcal{M}$. By assumption, \mathcal{M} allows learning (P_1, s_1, \dots, P_t) , so that the behavioral comparison applies. Note that $x_T(P_1, \mathcal{M}') \sim \mathbf{c}^+$ and there is no \mathcal{M}''' that strictly

dominates \mathcal{M}' after P_1 with $x_T(P_1, \mathcal{M}''') \sim \mathbf{c}^+$. Construct \mathcal{M}'' from \mathcal{M}' by adding the option to continue the sequence $(P_1, s_1, \dots, P_t, s_t)$ with learning Q' and then nothing thereafter. Now suppose that there is no $Q'' \in \Gamma(P_1, s_1, \dots, P_t, s_t)$ such that Q'' Blackwell dominates Q' . Then $x_T(P_1, \mathcal{M}'') \not\sim \mathbf{c}^+$, a contradiction. Hence, there is $Q'' \in \Gamma(P_1, s_1, \dots, P_t, s_t)$ that dominates Q' .

We now show that expertise in Definition 5.2 implies the behavioral condition in the proposition. For the ICPS in the definition, $x_T(\mathcal{M}^*, P_1) \sim \mathbf{c}^+$ implies that a path that pointwise Blackwell dominates $(P_1, s_1, P_2, \dots, s_t, Q)$ can be learned initially. There are three cases to consider:

- (a) \mathcal{M}' and the true ICP \mathcal{M} dynamically Blackwell dominate each other. Then by expertise \mathcal{M}'' must also be dominated by \mathcal{M} , and hence $x_T(P_1, \mathcal{M}'') \sim \mathbf{c}^+$, a contradiction to the assumption that $x_T(P_1, \mathcal{M}'') \not\sim \mathbf{c}^+$. This rules out the first case.
- (b) \mathcal{M} does not dynamically Blackwell dominate \mathcal{M}' , and hence $x_T(P_1, \mathcal{M}') \not\sim \mathbf{c}^+$, which is property (i) in the behavioral comparison.
- (c) \mathcal{M} strictly dynamically Blackwell dominates \mathcal{M}' . In that case $x_T(P_1, \mathcal{M}) \sim \mathbf{c}^+$ which is property (ii) of the behavioral comparison for $\mathcal{M}''' = \mathcal{M}$.

Because the three cases are exhaustive, the proposition is proved. \square

References

- 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. 21).
- Currie, Janet, and W. Bentley MacLeod. 2017. ‘‘Diagnosing Expertise: Human Capital, Decision Making and Performance Among Physicians’’. *Journal of Labor Economics* 35 (1): 1–43. (Cit. on p. 5).
- de Oliveira, Henrique. 2018. ‘‘Blackwell’s Informativeness Theorem using Diagrams’’. *Games and Economic Behavior* 109:126–131. (Cit. on p. 12).
- 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. 24, 25).
- Dillenberger, David, R. Vijay Krishna, and Philipp Sadowski. 2021. *Information-based Utilities and Beliefs: Foundations and Applications*. Tech. rep. Duke University. (Cit. on p. 25).

- 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 p. 25).
- Dudley, Richard Mansfield. 2002. *Real Analysis and Probability*. 2nd edition. Cambridge University Press. (Cit. on p. 27).
- Geanakoplos, John, and Paul R. Milgrom. 1991. “A Theory of Hierarchies Based on Limited Managerial Attention”. *Journal of Japanese and International Economies* 5:205–225. (Cit. on pp. 2, 5).
- Greenfield, Adam. 2018. *Radical Technologies: The Design of Everyday Life*. UK: Penguin Random House. (Cit. on p. 3).
- Greenshtein, Eitan. 1996. “Comparison of Sequential Experiments”. *Annals of Statistics* 24 (1): 436–448. (Cit. on p. 12).
- Gul, Faruk, and Wolfgang Pesendorfer. 2006. “Random Expected Utility”. *Econometrica* 74 (1): 121–146. (Cit. on p. 24).
- . 2004. “Self-Control and the Theory of Consumption”. *Econometrica* 72 (1): 119–158. (Cit. on p. 7).
- 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. 13).
- Hu, Yingyao, and Matthew Shum. 2012. “Nonparametric Identification of Dynamic Models with Unobserved State Variables”. *Journal of Econometrics* 171:32–44. (Cit. on p. 24).
- Kreps, David M., and Evan L Porteus. 1978. “Temporal Resolution of Uncertainty and Dynamic Choice Theory”. *Econometrica* 46 (1): 185–200. (Cit. on p. 7).
- Krishna, R. Vijay, and Philipp Sadowski. 2014. “Dynamic Preference for Flexibility”. *Econometrica* 82 (2): 655–703. (Cit. on pp. 8, 13, 25, 33, 34).
- Laffont, Jean-Jacques. 1989. *The Economics of Uncertainty and Information*. The MIT Press. (Cit. on pp. 33, 34).
- Magnac, Thierry, and David Thesmar. 2002. “Identifying Dynamic Discrete Decision Processes”. *Econometrica* 70 (2): 801–816. (Cit. on pp. 23, 24).
- 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. 30).
- Massa, Massimo, and Andrei Simonov. 2006. “Hedging, Familiarity and Portfolio Choice”. *Review of Financial Studies* 19 (2): 633–685. (Cit. on p. 21).

- Miller, Robert. 1984. “Job Matching and Occupational Choice”. *Journal of Political Economy* 92 (6): 1086–1120. (Cit. on p. 24).
- Pakes, Ariel. 1986. “Patents as Options: Some Estimates of the Value of Holding European Patent Stocks”. *Econometrica* 54 (4): 755–784. (Cit. on p. 24).
- Pastorino, Elena. 2021. *Careers in Firms: The Role of Learning and Human Capital*. Tech. rep. Stanford University. (Cit. on p. 4).
- 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. 25).
- Rust, John. 1994. “Structural Estimation of Markov Decision Processes”. In *Handbook of Econometrics*, ed. by Robert F. Engle and Daniel F. McFadden, vol. IV. Elsevier. (Cit. on pp. 23, 24).
- Sims, Christopher A. 2003. “Implications of Rational Inattention”. *Journal of Monetary Economics* 50 (3): 665–690. (Cit. on p. 2).
- . 1998. “Stickiness”. *Carnegie-Rochester Conference Series on Public Policy* 49:317–356. (Cit. on p. 2).
- Stigler, George. 1961. “The Economics of Information”. *Journal of Political Economy* 69 (3): 213–225. (Cit. on p. 2).