Comparing Experiments in Discounted Problems *

L. Renou† & X. Venel‡ & M. Whitmeyer§ & C. Williams¶

Latest version: November 11, 2026
First version: June 9, 2023

Abstract

This paper compares statistical experiments in discounted problems, ranging from the simplest ones, where the state is fixed and the flow of information exogenous, to more complex ones, where the decision-maker controls the flow of information or the state changes over time.

KEYWORDS: statistical experiments, repeated problems, discounting, controls, sufficiency, information, value.

JEL CLASSIFICATION: C73, D82.

*We thank audiences at the 2023 Warwick Economic Theory Workshop, 2023 SAET Conference, the 2024 Mediterranean Game Theory Symposium, Collegio Alberto-CEPR workshop, the 2026 Cowles Conference on Economic Theory, the Second Workshop on Game Theory and Information, the University of Toronto, UCL, HEC, QMUL, Institut Henri Poincaré, the University of Southern California, and the University of Iowa for comments and suggestions. LR thanks the hospitality of LUISS. XV is a member of the Gruppo Nazionale per l'Analisi Matematica, la Probabilita è le loro Applicazioni. XV is part of the Italian MIUR PRIN PNRR project "Supply Chain Disruptions, Financial Losses and Their Preventions" (project P2022XT8C8) and the MIUR PRIN 2022 project "Learning in Markets and Society" (project 2022EKNE5K).

†ASU and CEPR, lrenou.econ(at)gmail.com
‡LUISS, xvenel(at)luiss.it
§ASU, mwhitmey(at)asu.edu
¶University of Nebraska-Lincoln, cole.randall.williams(at)gmail.com

1 Introduction

The value of information depends not only on *what* it says, but *when* it is heard. A precise report that arrives late may be less useful than a cruder signal available early. Take, for example, an investor who rebalances a portfolio every month. One information source provides noisy market commentary throughout the quarter; another provides more reliable information only after accounting reports are released. As another example, consider a retailer choosing weekly inventories. One source provides rough demand estimates early in the season; another produces a more accurate market study only near the end. These examples illustrate a central trade-off when comparing information sources in dynamic problems: not only the overall informativeness of a source matters, but also its flow. This paper develops a method to make such comparisons.

We focus primarily on discounted problems, where the decision-maker evaluates sequences of decisions according to the discounting criterion. We start with simple discounted problems, where there is a fixed payoff-relevant state and, at each period, the decision-maker receives a signal, which may depend on all past signals (but not on past decisions), and then makes a decision. We compare the statistical experiments f and g, where experiment f (resp., g) is a collection of kernels $f_t : \Theta \times X^{t-1} \to \Delta(X_t)$ (resp., $g_t : \Theta \times Y^{t-1} \to \Delta(Y_t)$) mapping states and profiles of past signals to probabilities over current signals. Following Blackwell (1951, 1953), we say that the experiment f is more valuable than the experiment g if f induces a higher ex-ante discounted payoff than g, in all decision problems.

We first consider general, but fixed, discount factors. General discount factors are flexible enough to model geometric discounting, hyperbolic discounting, and many other forms of discounting. Let $f^t : \Theta \to \Delta(X^t)$ (resp., $g^t : \Theta \to \Delta(Y^t)$) be the t-period experiment induced by f (resp., g). Our first result, Theorem 1, states that f is more valuable than g in δ-discounted problems if, and only if, the mixture experiment $\sum_t \delta_t f^t$ is Blackwell sufficient for the mixture experiment $\sum_t \delta_t g^t$, which we term δ-*sufficiency*. Informally, the condition says that the experiment f needs to be more informative than g on average, but does not need to be more informative period-per-period. The latter condition is necessary and sufficient when considering *all* discount factors, a topic we explore in Section 2.4. The intuition for δ-sufficiency is simple. In discounted problems, discounted payoffs are equivalently written as *static* expected payoffs with respect to

the discounted probabilities of choosing actions. In turn, the discounted probability of choosing action a can be decomposed as the sum of the probabilities of choosing a at period t, conditional on the signal x^t, multiplied by the probability of x^t. Crucially, since past actions do not influence the evolution of the signals, the probability of x^t in the "as-if static" representation is $\delta_t f^t(x^t|\theta)$ when the state is θ, the probability of x^t in the mixture $\sum_s \delta_s f^s$. As an application of Theorem 1, we compare two experiments, which differ in their (random) arrival time of information.

We then move to more general discounted problems, ranging from relaxing the assumption that the decisions made do not influence the signals received to allowing for history-dependent sets of decisions through relaxing the assumption of a fixed payoff-relevant state, and many others. Economic applications motivate these generalizations. For instance, in a dynamic monopolistic pricing problem, the prices the seller offers over time partly determine what the seller learns about the unknown demand, say because the quantities demanded at each price offered are observed. So, decisions (prices) influence what is learnt (the realized demand). As another instance, in portfolio problems, it is natural to assume that the values of stocks and bonds (the payoff-relevant states) change over time, perhaps due to changes in macroeconomic and political conditions. As yet another instance, in consumption-saving or investment problems, decisions made today influence the possible decisions tomorrow. The generalizations we propose speak to these economic applications, and a few others.

In more general discounting problems, the equivalence with "as-if static" problems breaks down in many ways. For instance, suppose that the decision-maker controls the flow of information through the decisions made. The most natural idea is to let the experiments depend *directly* on the decisions made. This is, however, problematic as changing the decision problem would then change the experiments we compare – what we want is to compare two *fixed* experiments. The solution we adopt is to introduce ancillary variables into the definition of an experiment and to have decisions influence these variables. This formalism allows us to change the decision problem the decision-maker faces, without changing the experiments. We call these variables *informational controls*. For instance, when a firm chooses its pricing and marketing strategies (decisions), it partly controls the fraction and types of consumers made aware of its products (controls) and, ultimately, the realized demand (observations). To compare two *controlled* experiments, we extend them to two larger experiments, where at each period an extended experi-

ment produces a new control and signal, as a function of past controls and signals and the state. We then compare these larger experiments as in the uncontrolled case. We hasten to stress that the extensions are not unique. Intuitively, each decision problem induces a pair of *linked* extensions through its mapping from decisions to controls. As a consequence, we need to do the comparison for *all* possible extensions. We illustrate the comparisons within the class of stationary and bandit problems – Section 3.3.1.

We close the introduction with a brief discussion of the most closely related paper Greenshtein (1996). A more extensive discussion can be found in Section 4. Greenshtein (1996) compares statistical experiments in sequential problems, where the decision-maker's payoff function over sequences of decisions is arbitrary. In particular, payoff functions do not have to be separable across periods, let alone to be discounted sums of instantaneous payoffs. In all other aspects, his model coincides with the model we study in Section 2.5.1, which assumes a fixed state, an exogenous evolution of the signals, and history-dependent choice sets. Since Greenshtein considers a larger class of problems, his ranking of statistical experiments is coarser than ours. In fact, it is even stronger than requiring that f^t is sufficient for g^t at all periods. We also stress that none of the other models we consider are treated in Greenshtein.

Roadmap. We present a large number of novel results, which parallel the richness of the class of dynamic problems. Section 2 groups together all our results, which apply to problems where the decisions made *do not* affect the information the decision-maker obtains, while Section 3 groups the results, where the decisions affect the information the decision-maker receives. We qualify the former problems as *uncontrolled* and the latter as *controlled*. To avoid unnecessary repetitions, Section 3 focuses on the unique challenges that controlled dynamic problems introduce.

Section 2 starts with the simplest problems, where the decision-maker chooses repeatedly from a fixed menu and learns from an exogenous source of information (Sections 2.1 to 2.3). E.g., we can think of a consumer repeatedly purchasing a product of unknown quality and learning from online reviews and word-of-mouth.[1] We then consider inter-temporal dependences in Section 2.5. The first inter-temporal dependence we consider is when the choices the decision-maker makes at a period impact what the decision-maker can do at the next period. E.g., in saving-consumption problems, saving today impacts

[1] If, in addition, the consumer learns from their own experience, the problem is controlled, which is the focus of Section 3.

consumption tomorrow. Another inter-temporal dependence we examine is when the duration of the decision problem is endogenous. To do so, we interpret the discount factor as the probability of the problem lasting one more period, conditional on not having stopped earlier, and assume that the discount factor depends on past choices. E.g., the successful completion of a project is typically random and depends on past efforts. Section 3 parallels Section 2. Section 3.1 starts with the formulation of controlled problems, with an emphasis on the dichotomy between experiments and decision problems. Section 3.2 then revisits the class of problems studied in Section 2.3, while 3.3 extends the analysis to separable problems. Besides the wide range of economic settings we cover, we also discuss more conceptual issues such as dynamic (in)consistency (Sections 2.3 and 3.4), fixed discount vs a class of discount factors (Sections 2.4 and 3.3), and compare our results with more general decision problems (Section 2.5). Applications are in Sections 2.3.1, 3.2.1 and 3.3.1.

2 Discounted Problems: The Uncontrolled Case

This section studies discounted problems in which the flow of information is exogenous: the signals generated by an experiment depend on the state and, possibly past signals, but not on past or current decisions. We call these problems *uncontrolled*. The *controlled* case, in which decisions affect the information the decision-maker receives, is treated in Section 3.

2.1 The Formulation

We first formalize the most basic uncontrolled problem. The decision-maker receives information about an unknown *fixed* state (or parameter) $\theta \in \Theta$ at periods $t = 1, \dots, T \leq +\infty$. The information the decision-maker receives is independent of the decisions made. The discount factor is fixed, and the set of feasible decisions is the same at every period.

In what follows, with the exception of the number of periods, all sets are finite. To avoid trivialities, $|\Theta| > 1$. For any collection of sets $S_1, \dots, S_t$, we write S^t for $S_1 \times \cdots \times S_t$, s_t (resp., s^t) for a generic element of S_t (resp., S^t). We denote by $\Delta(S_t)$ the set of probabilities over S_t and write random vectors in bold.

Experiments. An experiment $f := (f_1, f_2, \dots)$ is a collection of kernels $f_t : \Theta \times X^{t-1} \to \Delta(X_t)$, where X_t is the set of signals at period t and X^{t-1} the set of past signals. We write $f_t(x_t|x^{t-1}, \theta)$ for the probability of the signal x_t given (x^{t-1}, θ) and $f^t(x^t|\theta)$ for $f_1(x_1|\theta) \times \cdots \times f_t(x_t|x^{t-1}, \theta)$. The experiment f models the information the decision-maker receives over time: At each period t, the decision-maker observes the signal x_t. The objective is to compare the experiments f and g, where the experiment $g := (g_1, g_2, \dots)$ is the collection of kernels $g_t : \Theta \times Y^{t-1} \to \Delta(Y_t)$.

Posterior representation. Let p be a prior on Θ. At each period t, the experiment f^t induces a distribution $\mu^t_{p,f} \in \Delta(\Delta(\Theta))$ over posteriors. We define $\mu^t_{p,g}$ analogously.

Decision problems. A decision problem is a tuple (A, u), where A is the non-empty set of decisions, and $u\colon A \times \Theta \to \mathbb{R}$ the payoff function.

Discount factors. A discount factor δ is a probability distribution over $\{1, \dots, T\}$, with δ_t the probability of t.[2] Examples include geometric discounting ($\delta_t = \frac{1-\rho}{1-\rho^T}\rho^{t-1}$ for some $\rho \in (0, 1)$), quasi-hyperbolic discounting, averaging over finitely many periods ($\delta_t = 1/T^*$ for some finite T^*), discounting of a single period ($\delta_t = 1$ for some t), among others.

Discounted problems. A δ-discounted problem is such that the decision-maker values the sequences of decisions $(a_1, \dots, a_T)$ as the discounted sum $\sum_t \delta_t u(a_t, \theta)$.

Strategies. Fix the decision problem (A, u). We write σ (resp., τ) for the strategy adapted to the experiment f (resp., g), that is, σ is the collection of maps $(\sigma_t : A^{t-1} \times X^t \to \Delta(A))_t$ (resp., $(\tau_t : A^{t-1} \times Y^t \to \Delta(A))_t$). We stress that the decision-maker does *not* observe realized payoffs – the only information the decision-maker has are the signals x_t. If the decision-maker were also to observe realized payoffs, we would need to change the definition of experiments to include these additional signals. These experiments would be controlled, which is the focus of Section 3. We denote $\mathbb{P}_{\theta,f,\sigma}$ the probability over histories of signals and decisions, induced by the strategy σ and the experiment f, when the state is θ, and write $\mathbb{E}_{\theta,f,\sigma}$ for the expectation with respect to that probability. We write $\operatorname{marg}_{A_t}\mathbb{P}_{\theta,f,\sigma} \in \Delta(A)$ for the marginal distribution over decisions at period t. We define $\mathbb{P}_{\theta,g,\tau}$, $\mathbb{E}_{\theta,g,\tau}$, and $\operatorname{marg}_{A_t}\mathbb{P}_{\theta,g,\tau}$ similarly.[3]

[2] We can also interpret δ_t as the probability that the decision-maker is active at period t.

[3] Since all sets are finite, the existence and uniqueness of $\mathbb{P}_{\theta,f,\sigma}$ are immediate when $T < \infty$ and follow from the Ionescu-Tulcea extension theorem when $T = +\infty$.

Before proceeding, let us highlight the most important assumptions the above formulation imposes (besides the information being uncontrolled). To start with, the state θ is fixed. In many applications, such as portfolio and insurance problems, the state evolves over time. Similarly, the discount factor is fixed. Again, in applications, it is sometimes convenient to consider a range of discount factors. E.g., it is frequent to restrict attention to sufficiently patient decision-makers. Lastly, the set of feasible decisions is the same at each period, hence independent of past decisions. Yet again, in applications, past decisions may constrain future decisions, e.g., past investments may constrain future investments. We present all these generalizations and a few others in later sections. We now offer a brief review of the seminal results on the comparison of experiments in static problems.

2.2 Static Problems: A Brief Review

This section recalls the static comparison result that underlies our dynamic characterization. We refer the reader to Blackwell and Girshick (1954, Chapter 12), Torgersen (1991), Le Cam (1996) and Strasser (2011, Chapter 9) for extensive reviews.

The comparison of the experiments f and g rests on Definition 1, which states that experiment f is more valuable than experiment g if, regardless of the decision problem and the state, the decision-maker is always better-off with experiment f than with experiment g. This definition, albeit in a slightly different form, first appeared in Blackwell (1951).[4]

Definition 1 *The experiment f is more valuable than the experiment g in static problems if for all decision problems (A, u), for all strategies τ, there exists a strategy σ such that $\mathbb{E}_{\theta,g,\tau} u(\mathbf{a}, \theta) \leq \mathbb{E}_{\theta,f,\sigma} u(\mathbf{a}, \theta)$ for all θ.*

The following theorem summarizes the main results on the comparison of experiments.

Theorem 0 (Blackwell (1951, 1953)) *The following statements are equivalent:*

(1) The experiment f is more valuable than the experiment g in static problems.

[4] Blackwell's original definition compares sets of attainable payoff profiles $\bigcup_\sigma \{u \in \mathbb{R}^{|\Theta|} : u_\theta = \mathbb{E}_{\theta,f,\sigma} u(\mathbf{a}, \theta), \forall \theta\}$ and $\bigcup_\tau \{u \in \mathbb{R}^{|\Theta|} : u_\theta = \mathbb{E}_{\theta,g,\tau} u(\mathbf{a}, \theta), \forall \theta\}$, and requires the inclusion of the latter in the former for all (A, u). Blackwell (1951) attributed the definition to Bohnenblust, Shapley and Sherman (unpublished).

(2) *For all decision problems* (A, u)*, for all strategies* τ*, there exists a strategy* σ *such that* $\text{marg}_A \mathbb{P}_{\theta,g,\tau} = \text{marg}_A \mathbb{P}_{\theta,f,\sigma}$*, for all* θ*.*

(3) *Sufficiency: There exists* $\gamma \colon X \to \Delta(Y)$ *such that*

$$g(y|\theta) = \sum_x \gamma(y|x) f(x|\theta), \text{ for all } (y, \theta).$$

(4) *For every prior* $p \in \Delta(\Theta)$*,* $\mu_{p,f}$ *dominates* $\mu_{p,g}$ *in the convex order.*

The equivalence between (1) and (2) first appeared in Blackwell (1951) as Theorem 2(i), and the equivalence between (1) and (3) as Theorems 5 and 6 in Blackwell (1953) – Blackwell attributed Theorem 5 to Sherman (1951) and Stein (1951). See de Oliveira (2018) for a recent proof. The posterior formulation in (4) already appeared in Blackwell (1951, 1953).

A few remarks are worth making. First, we compare the experiments state-by-state. An alternative is to compare expected payoffs, with respect either to a fully supported prior or every prior. The two approaches are equivalent – see Blackwell (1951). Second, it follows from the sufficiency condition that the decision-maker can emulate $\text{marg}_A \mathbb{P}_{\theta,g,\tau}$ by first sampling a *fictitious* y with probability $\gamma(y|x)$ and then choosing a with probability $\tau(a|y)$. As we will see, the simulation of fictitious profiles of signals will continue to play a central role in our analysis.

2.3 Fixed Discount Factor

This section compares experiments in discounted problems, where the discount factor is fixed (and known). It is the backbone of all our comparisons. We first extend Definition 1 to:

Definition 2 *The experiment* f *is more valuable than the experiment* g *in* δ*-discounted problems if for all decision problems* (A, u)*, for all strategies* τ*, there exists a strategy* σ *such that*

$$\mathbb{E}_{\theta,g,\tau}\left[\sum_t \delta_t u(\mathbf{a}_t, \theta)\right] \leq \mathbb{E}_{\theta,f,\sigma}\left[\sum_t \delta_t u(\mathbf{a}_t, \theta)\right], \text{ for all } \theta.$$

We now show how to derive a characterization similar to the one of Blackwell (1951, 1953), but for discounted problems. The main idea consists in reformulating discounted

problems as static problems and then to apply Theorem 0. More specifically, observe that we can rewrite the decision-maker's expected payoff as follows:

$$\mathbb{E}_{\theta,g,\tau}\left[\sum_t \delta_t u(\mathbf{a}_t,\theta)\right] = \sum_a u(a,\theta)\underbrace{\left[\sum_t \delta_t \text{marg}_{A_t}\mathbb{P}_{\theta,g,\tau}(a)\right]}_{\text{discounted probability of } a}.$$

In words, the decision-maker's expected payoff is the expectation of the static payoff with respect to the discounted probabilities of making decisions. As we shall see, not all discounted problems admit such a reformulation, but our simplest formulation does. It immediately follows from Theorem 0 that f is more valuable than g if for all strategies τ, there exists a strategy σ, which replicates the discounted probabilities of decisions, that is,

$$\sum_t \delta_t \text{marg}_{A_t}\mathbb{P}_{\theta,g,\tau} = \sum_t \delta_t \text{marg}_{A_t}\mathbb{P}_{\theta,f,\sigma}, \text{ for all } \theta.$$

Next, observe that we can rewrite the discounted probability of decision a as follows:

$$\begin{aligned}\sum_t \delta_t \mathbb{P}_{\theta,g,\tau}(\boldsymbol{a}_t = a) &= \sum_{t,y^t} \delta_t \mathbb{P}_{\theta,g,\tau}(y^t)\mathbb{P}_{\theta,g,\tau}(\boldsymbol{a}_t = a|y^t),\\ &= \sum_{t,y^t} \underbrace{\delta_t g^t(y^t|\theta)}_{\text{signal}}\underbrace{\mathbb{P}_{\theta,g,\tau}(\boldsymbol{a}_t = a|y^t)}_{\text{indep. of }\theta \Rightarrow \text{ strategy}}.\end{aligned}$$

To understand the above rewriting, recall that the evolution of the signals does not depend on the decisions made and, consequently, the probability $\mathbb{P}_{\theta,g,\tau}(y^t)$ is independent of the strategy τ, and is equal to $g^t(y^t|\theta)$. In addition, since the decision-maker does not observe θ, the conditional probability $\mathbb{P}_{\theta,g,\tau}(\boldsymbol{a}_t = a|y^t)$ is independent of θ and, thus, we can view it as a behavioral strategy mapping sequences of signals to probabilities over actions. In other words, the static re-formulation states that the experiment g generates the signal "(t,y^t)" with probability $\delta_t g^t(y^t|\theta)$ when the state is θ, the decision-maker observes the realized signals, and makes a decision. (Recall that y^t is an element of Y^t, the set of sequences of signals of length t. Thus, summing over (t,y^t) is equivalent to summing over $\bigcup_{t=1}^T Y^t$.)

Since the same re-formulation also applies to the experiment f, Theorem 0 suggests that for f to be more valuable than g, we need the existence of garblings $\gamma_{t'} : X^{t'} \to \Delta(\cup_{t=1}^T Y^t)$

such that

$$\delta_t g^t(y^t|\theta) = \sum_{t',x^{t'}} \delta_{t'} f^{t'}(x^{t'}|\theta)\gamma_{t'}(y^t|x^{t'}), \text{for all } (t, y^t, \theta).$$

The same static re-formulation also yields a formulation in terms of posteriors. To state that reformulation, we write $\mu^{\delta}_{p,f} := \sum_t \delta_t \mu^t_{p,f}$ for the discounted distribution over posteriors induced by f and define $\mu^{\delta}_{p,g}$ analogously. We stress that the purpose of the static re-formulations is to compare experiments. The underlying problem remains dynamic: information arrives over time and strategies condition on histories of signals.

Theorem 1 *The following statements are equivalent:*

(1) The experiment f is more valuable than the experiment g in δ-discounted problems.

(2) For all decision problems (A, u), for all strategies τ, there exists a strategy σ such that

$$\sum_t \delta_t \text{marg}_{A_t} \mathbb{P}_{\theta,g,\tau} = \sum_t \delta_t \text{marg}_{A_t} \mathbb{P}_{\theta,f,\sigma}, \textit{for all } \theta.$$

(3) δ-sufficiency: For all $t' = 1, \ldots, T$, there exists a garbling $\gamma_{t'} : X^{t'} \to \Delta(\cup_{t=1}^T Y^t)$ such that

$$\delta_t g^t(y^t|\theta) = \sum_{t'=1}^{T} \sum_{x^{t'} \in X^{t'}} \delta_{t'} f^{t'}(x^{t'}|\theta)\gamma_{t'}(y^t|x^{t'}), \textit{for all } (t, y^t, \theta).$$

(4) For every prior $p \in \Delta(\Theta)$, $\mu^{\delta}_{p,f}$ dominates $\mu^{\delta}_{p,g}$ in the convex order.

Intuitively, δ-sufficiency states that the information needed to emulate the experiment g is contained in the experiment f, but f does not need to be more informative than g at all periods, i.e., we do not need g^t to be a garbling of f^t at all t.[5] For instance, f can be less informative than g at earlier periods if it is sufficiently more informative at later periods. To put it simply, all we need is that there is more information in f than g on average, that is, we need the mixture experiment $\sum_t \delta_t f^t$ to be (Blackwell) sufficient for the mixture experiment $\sum_t \delta_t g^t$.

Commitment and consistent planning. The comparisons in Definition 2 are based on the *commitment solution* to the decision problems, that is, as if the decision-maker

[5] We can equivalently define a unique garbling $\gamma : \bigcup_t X^t \to \Delta(\cup_t Y^t)$.

commits ex-ante to optimal strategies. An alternative is to consider the *consistent-planning solution*, that is, equilibria of the intra-personal game between the multiple selves – one self per period – of the decision-maker. Naturally, the two comparisons are identical when the profile of discount factors δ induces time-consistent preferences – a ubiquitous assumption in Economics. More generally, in the uncontrolled problems considered so far, moving from the commitment solution to the consistent-planning solution does not change the ranking of experiments. Indeed, the optimal behavior is myopic since there are no inter-temporal linkages. For a more extensive discussion, we refer the reader to Section 3.4.

2.3.1 Application: Comparing The Arrival Time of News

This section presents our running application: the comparison of experiments which differ in their arrival time of some information. In Section 2.4.2, we revisit this comparison to illustrate additional results.

There is a non-trivial (static) experiment $h\colon \Theta \to \Delta(X)$ and two probability distributions F and G over $\{1,\dots,T\}$, with associated random variables $\boldsymbol{T}_F$ and $\boldsymbol{T}_G$, respectively.[6] We write F_t (resp., G_t) for the value of F (resp., G) at t, which we interpret as the probability that the information arrives by time t. We use the convention $F_0 = G_0 = 0$.

We are interested in comparing two dynamic experiments, f and g, which differ in their (random) arrival time $\boldsymbol{T}_F$ and $\boldsymbol{T}_G$ of the same information, i.e., a signal from h. A formal definition of the experiment f is as follows. At each period t, the set of signals X_t is $X \cup \{\emptyset\}$, where $\emptyset$ is a dummy symbol ("no news"). Conditional on not having observed a signal in X yet, the decision-maker observes the signal $x_t = x$ with probability $\frac{F_t - F_{t-1}}{1-F_{t-1}} h(x|\theta)$ when the state is θ, and the signal $x_t = \emptyset$ with the probability $\frac{1-F_t}{1-F_{t-1}}$. Conditional on having received a signal in X at an earlier period, the signal at period t is $\emptyset$ with probability one. We define the experiment g analogously.

Proposition 1 *The experiment f is δ-sufficient for the experiment g if, and only if,*

$$\mathbb{E}\left[\sum_{t=\boldsymbol{T}_F}^{T} \delta_t\right] \geq \mathbb{E}\left[\sum_{t=\boldsymbol{T}_G}^{T} \delta_t\right] \Leftrightarrow \sum_{t=1}^{T} \delta_t F_t \geq \sum_{t=1}^{T} \delta_t G_t \tag{1}$$

[6] The experiment h is non-trivial if it has a strictly positive information value for some decision problems.

The intuition is simple. Let (A, u) be a decision problem and $p \in \Delta(\Theta)$ a prior. Upon arrival of the signal x at period t, the decision-maker's continuation payoff is

$$\sum_{s=t}^{T} \delta_s \max_a \mathbb{E}_{p,h}[u(a, \boldsymbol{\theta})|x],$$

while the cumulative payoff prior to the arrival of the signal is $\sum_{s=1}^{t-1} \delta_s \max_a \mathbb{E}_{p,h}[u(a, \boldsymbol{\theta})]$. Thus, if we let $V^*_{p,h}(u) \coloneqq \mathbb{E}_{p,h}\Big[\max_a \mathbb{E}_{p,h}[u(a, \boldsymbol{\theta})|\mathbf{x}]\Big] - \max_a \mathbb{E}_{p,h}[u(a, \boldsymbol{\theta})] \geq 0$, we compare two "time" lotteries, where a lottery pays 0 before t, the realized arrival time, and $V^*_{p,h}(u)$ afterwards. (We are translating the payoff streams by $\max_a \mathbb{E}_{p,h}[u(a, \boldsymbol{\theta})]$, the uninformed payoff.) The decision-maker is thus better off with the experiment f than g if, and only if, $\mathbb{E}\left[\sum_{t=\boldsymbol{T}_F}^{T} \delta_t\right] \geq \mathbb{E}\left[\sum_{t=\boldsymbol{T}_G}^{T} \delta_t\right]$.

So far, we have assumed that the two experiments f and g generate a signal from the same static experiment h. Thus, we cannot compare situations, where one source of information delivers faster, but less informative, news than the other. To extend the analysis to such comparisons, we now assume that upon realization of the random time $\boldsymbol{T}_F$ (resp., $\boldsymbol{T}_G$), the experiment f (resp., g) generates a signal from the non-trivial (static) experiment $h_X \colon \Theta \to \Delta(X)$ (resp., $h_Y : \Theta \to \Delta(Y)$). Fix a full-support prior $p \in \Delta(\Theta)$, and define

$$r^* \coloneqq \sup_r \{r \geq 0 \colon V^*_{p,h_X}(u) \geq r V^*_{p,h_Y}(u), \forall u\}.$$

The index r^* quantifies the comparative advantage of experiment h_X over experiment h_Y. In particular, if h_X is (Blackwell) sufficient for h_Y, then $r^* \geq 1$, but it may be strictly greater. In a companion note, we offer a full characterization of r^* – its value is independent of p. For our purpose, all we need to know is that $V^*_{p,h_X}(u) \geq r^* V^*_{p,h_Y}(u)$ binds for some problem (A, u) such that $V^*_{p,h_Y}(u) > 0$. From our prior discussion, we are thus comparing time lotteries of the form $(0, \dots, 0, \dots, 0, \underbrace{V^*_{p,h_X}(u)}_{t_F}, V^*_{p,h_X}(u) \dots)$ and $(0, \dots, 0, \underbrace{V^*_{p,h_Y}(u)}_{t_G}, V^*_{p,h_Y}(u) \dots)$, where t_F and t_G are the arrival times. It immediately follows that:

Corollary 1 *The experiment f is δ-sufficient for the experiment g if, and only if,*

$$\mathbb{E}\left[\sum_{t=\boldsymbol{T}_F}^{T}\delta_t\right]r^* \geq \mathbb{E}\left[\sum_{t=\boldsymbol{T}_G}^{T}\delta_t\right].$$

Informally, if the experiment h_X is less informative than h_Y, hence $r^* < 1$, the information from h_X must arrive sufficiently earlier than the one from h_Y to compensate for its "lack of informativeness." Conversely, if $r^* > 1$, the information from h_X may arrive later than the one from h_Y and still dominate. E.g., if the time horizon is infinite and discounting geometric, i.e., $\delta_t = (1-\rho)\rho^{t-1}$ for some $\rho \in (0,1)$, we need to shift the arrival time $\boldsymbol{T_F}$ by $\left\lceil \frac{\ln(1/r^*)}{\ln \rho} \right\rceil \geq 0$ periods when $r^* \geq 1$ to compensate for the superior informativeness of h_X.[7] The supplementary appendix (B.1) makes this time-shift interpretation precise.

2.3.2 Extension: Time-Varying States

This subsection relaxes the assumption of a fixed state. We show that our analysis generalizes almost immediately to the state evolving over time. Let $p \in \Delta(\Theta_1 \times \cdots \times \Theta_T)$ be the distribution over sequences of states. Without loss of generality, we assume that $\Theta_t = \Theta$ for all t. We assume furthermore that the decision-maker's payoff at period t only depends on the state at period t. (If the payoff was a function of the entire profile of states, then the previous analysis applies verbatim.) We write $\mathbb{P}_{p,f,\sigma}$ for the distribution over states, signals and actions, induced by the distribution of states p, the experiment f and the strategy σ. Similarly, for $\mathbb{P}_{p,g,\tau}$.

As in Section 2.3, we can rewrite the expected payoff as:

$$\mathbb{E}_{p,g,\tau}\left[\sum_t \delta_t u(\boldsymbol{a}_t, \boldsymbol{\theta}_t)\right] = \sum_{a,\theta} u(a,\theta) \underbrace{\left[\sum_t \delta_t \text{marg}_{A_t \times \Theta_t} \mathbb{P}_{p,g,\tau}(a,\theta)\right]}_{\text{discounted probability of } (a,\theta)}.$$

It once again follows immediately that an equivalent statement to f being more valuable than g is that for all decision problems (A,u), for all strategies τ, there exists a strategy

[7] When $r^* \leq 1$, we shift the arrival time $\boldsymbol{T_G}$ by $\left\lceil \frac{\ln r^*}{\ln \rho} \right\rceil \geq 0$ periods.

σ such that

$$\sum_t \delta_t \mathbb{P}_{p,g,\tau}(\boldsymbol{a}_t = a, \boldsymbol{\theta}_t = \theta) = \sum_t \delta_t \mathbb{P}_{p,f,\sigma}(\boldsymbol{a}_t = a, \boldsymbol{\theta}_t = \theta),$$

for all (a, θ). In words, all we need to replicate are the discounted probabilities of states and actions.

Since neither the realizations of the states nor the realizations of the signals depend on the decisions made, the distribution over states and signals is independent of τ. We denote it $\mathbb{P}_{p,g}$. In particular, if the state is fully persistent, i.e., $p(\theta_1, \ldots, \theta_t, \ldots, \theta_T) > 0$ if, and only if, $\theta_t = \theta_{t'}$ for all (t, t'), then

$$\mathbb{P}_{p,g}(\boldsymbol{y}^t = y^t, \boldsymbol{\theta}_t = \theta) = p(\theta, \ldots, \theta, \ldots, \theta) g^t(y^t|\theta).$$

We define $\mathbb{P}_{p,f}$ similarly. In this subsection, $\mu^t_{p,f} \in \Delta(\Delta(\Theta))$ denotes the distribution of the posterior $\mathbb{P}_{p,f}(\boldsymbol{\theta}_t \in \cdot \mid \boldsymbol{x}^t)$ about the current state. We define $\mu^t_{p,g}$ analogously. As before, write $\mu^\delta_{p,f} \coloneqq \sum_t \delta_t \mu^t_{p,f}$ and $\mu^\delta_{p,g} \coloneqq \sum_t \delta_t \mu^t_{p,g}$.

The following theorem is the generalization of Theorem 1 to an environment with evolving states.

Theorem 2 *Let p be the distribution over the sequences of states. The following statements are equivalent.*

1. *The experiment f is more valuable than the experiment g in discounted problems with discount factor δ.*
2. *For all decision problems (A, u), for all strategies τ, there exists a strategy σ such that*
$$\sum_t \delta_t \mathrm{marg}_{A_t \times \Theta_t} \mathbb{P}_{p,g,\tau} = \sum_t \delta_t \mathrm{marg}_{A_t \times \Theta_t} \mathbb{P}_{p,f,\sigma}.$$
3. *δ-sufficiency: There exist garblings $\gamma_{t'} : X^{t'} \to \Delta(\cup_{t=1}^T Y^t)$ such that*
$$\delta_t \mathbb{P}_{p,g}(\boldsymbol{y}^t = y^t, \boldsymbol{\theta}_t = \theta) = \sum_{t', x^{t'}} \delta_{t'} \mathbb{P}_{p,f}(\boldsymbol{x}^{t'} = x^{t'}, \boldsymbol{\theta}_{t'} = \theta) \gamma_{t'}(y^t|x^{t'}), \textit{for all } t, y^t, \theta.$$
4. *$\mu^\delta_{p,f}$ dominates $\mu^\delta_{p,g}$ in the convex order.*

2.4 Families of Discount Factors

In Section 2.3, preferences over action streams $(a_1, \ldots, a_T)$ take the form $\sum_{t=1}^{T} \delta_t u(a_t, \theta)$ and the comparisons are over all static decision problems (A, u), as in Blackwell's work. What is fixed, however, is the inter-temporal aggregation: discounted sums of instantaneous payoffs with a *fixed* discount factor δ. This section relaxes this assumption and compares experiments over both decision problems *and* discount factors.[8]

Given a class $\mathcal{D} \subseteq \Delta(\{1, \ldots, T\})$ of discount factors, we will require the comparison to hold for every $\delta \in \mathcal{D}$. Larger classes $\mathcal{D}$ correspond to larger classes of dynamic preferences. Important classes of discount factors include decreasing and geometric discounting. The latter class is ubiquitous in economics, while the former models decision-makers, who are impatient.

Theorem 1 immediately extends to comparisons over a family $\mathcal{D}$ of discount factors: it suffices to require that the equivalent statements (1)-(4) hold for every $\delta \in \mathcal{D}$. For many families of discount factors, this leads to simple characterizations. The key observation is that δ-sufficiency is a convex property. Consequently, when $\mathcal{D}$ is closed and convex, it is enough to verify the comparison on its extreme points. We begin by establishing this convexity property, then derive characterizations for several important families of discount factors, and finally apply these results to the arrival of information.

2.4.1 Δ-sufficiency and $\Delta_{\downarrow}$-sufficiency

We first argue that the condition of δ-sufficiency satisfies convexity, that is, if experiment f is δ- and $\widehat{\delta}$-sufficient for g, then it is also $\alpha\delta + (1-\alpha)\widehat{\delta}$-sufficient for g, for all $\alpha \in [0, 1]$.

Proposition 2 *The set $\{\delta : f$ is δ-sufficient for $g\}$ is convex.*

Proof If the set is either empty or a singleton, there is nothing to prove. So, assume that the set contains at least two elements, say δ and $\widehat{\delta}$. Since f is δ-sufficient for g, there exist garblings $(\gamma_{t'})_{t'}$ such that

$$\delta_t g^t(y^t|\theta) = \sum_{t'=1}^{T} \sum_{x^{t'} \in X^{t'}} \delta_{t'} f^{t'}(x^{t'}|\theta)\gamma_{t'}(y^t|x^{t'}), \text{ for all } (t, y^t, \theta).$$

[8] It is worth noting that an alternative is to fix a decision problem (A, u) and to do the comparisons over all discount factors. We conduct this exercise in the supplementary appendix (B.2).

Similarly, since f is $\widehat{\delta}$-sufficient for g, there exist garblings $(\widehat{\gamma}_{t'})_{t'}$ such that

$$\widehat{\delta}_t g^t(y^t|\theta) = \sum_{t'=1}^{T} \sum_{x^{t'} \in X^{t'}} \widehat{\delta}_{t'} f^{t'}(x^{t'}|\theta)\widehat{\gamma}_{t'}(y^t|x^{t'}), \text{ for all } (t, y^t, \theta).$$

Let $\alpha \in [0,1]$ and consider the discount factor $\alpha\delta + (1-\alpha)\widehat{\delta}$. It is routine to verify that the garblings

$$\left(\frac{\alpha\delta_t}{\alpha\delta_t + (1-\alpha)\widehat{\delta}_t} \gamma_t + \frac{(1-\alpha)\widehat{\delta}_t}{\alpha\delta_t + (1-\alpha)\widehat{\delta}_t} \widehat{\gamma}_t \right)_t$$

satisfy the required equalities for $\alpha\delta + (1-\alpha)\widehat{\delta}$-sufficiency. If $\alpha\delta_t + (1-\alpha)\widehat{\delta}_t = 0$, the corresponding garbling may be defined arbitrarily. ■

As a consequence of Proposition 2, when the comparison is required to hold for every discount factor in a non-empty, closed, bounded, and convex set D, it is enough to verify the corresponding sufficiency condition at the extreme points of D.[9] In addition, Theorem 1(4) implies that the set of discount factors for which f is δ-sufficient for g is closed. This follows from the boundedness of the continuous convex test functions on $\Delta(\Theta)$ and the preservation of the convex-order inequalities under ℓ^1 limits.

We now present two special families. We begin with the case of *all* discount factors, that is, we want the experiment f to be δ-sufficient for g for all $\delta \in \Delta$, the set of all probability distributions over $\{1, \ldots, T\}$. Since the extreme points of the convex set Δ are the degenerate distributions on each single period, it follows from Proposition 2 that a necessary and sufficient condition is for f^t to be sufficient for g^t for all t.

The same condition also ranks experiments in all separable problems. A separable problem is one where the decision-maker values the sequences of actions $(a_1, \ldots, a_T)$ as $\sum_t u_t(a_t, \theta)$, where when $T = +\infty$, we require $\sum_t \max_{a \in A, \theta \in \Theta} |u_t(a, \theta)| < +\infty$. Discounted problems are special cases of separable problems. We also note that the class of separable problems is equivalent to the class of discounted problems with time-varying payoff functions.[10] The following theorem records these equivalent formulations.

Theorem 3 *The following statements are equivalent:*

[9] We endow Δ with the ℓ^1 norm. Since ℓ^1 satisfies the Radon-Nikodym property, we can write D as the closed convex hull of its extreme points. Compactness is not needed.

[10] Fix the full support discount factor δ. If the decision-maker values the sequence of actions $(a_t)_t$ as $\sum_t \delta_t v_t(a_t, \theta)$, then writing $u_t := \delta_t v_t, \sum_t \delta_t v_t(a_t, \theta) = \sum_t u_t(a_t, \theta)$. The converse is also true if we let $v_t := u_t/\delta_t$.

(1) The experiment f is more valuable than the experiment g in δ-discounted problems, for all $\delta \in \Delta$.

(2) The experiment f is more valuable than the experiment g in all separable problems.

(3) Δ-sufficiency: For all $t = 1, \ldots, T$, f^t is sufficient for g^t.

(4) For every prior $p \in \Delta(\Theta)$ and all $t = 1, \ldots, T$, $\mu^t_{p,f}$ dominates $\mu^t_{p,g}$ in the convex order.

Thus, δ- and Δ-sufficiency are two polar cases, where the first condition assumes a known discount factor, while the second considers all discount factors. In other words, we are considering an even larger class of decision problems, where not only the pair (A, u) may change, but also the discount factor δ. The condition of Δ-sufficiency is stronger than the condition of δ-sufficiency. The theorem also shows that the larger class of separable problems does not induce coarser comparisons than the class of discounted problems, if we consider all discount factors.

Obviously, there are many intermediate cases. As a second family, we consider the set $\Delta_\downarrow$ of decreasing discount factors, a frequently made assumption in economics, as it captures the idea of impatient decision-makers. The set $\Delta_\downarrow$ is convex and its extreme points are the uniform distributions over $\{1, \ldots, t\}$ for all t. For each t, let $\delta^{[t]}$ be the uniform distribution over $\{1, \ldots, t\}$. A direct consequence of Proposition 2 is then:

Corollary 2 *The following statements are equivalent:*

(1) The experiment f is more valuable than the experiment g in δ-discounted problems, for all $\delta \in \Delta_\downarrow$.

(2) $\Delta_\downarrow$-sufficiency: For all $t \geq 1$, $\frac{1}{\min(t,T)} \sum_{t'=1}^{\min(t,T)} f^{t'}$ is sufficient for $\frac{1}{\min(t,T)} \sum_{t'=1}^{\min(t,T)} g^{t'}$.

(3) For every prior $p \in \Delta(\Theta)$ and all $t \geq 1$, $\mu^{\delta^{[\min(t,T)]}}_{p,f}$ dominates $\mu^{\delta^{[\min(t,T)]}}_{p,g}$ in the convex order.

To close this section, we remark that the ranking of experiments may change multiple times as the decision-maker becomes more impatient. With more than two periods, even if f is more valuable than g for very patient and very impatient decision-makers, the converse may hold for moderately patient decision-makers. In other words, comparisons are not monotone with respect to the discount factor. We provide a three-period example in the supplementary appendix (B.3).[11]

[11] Similarly, for a fixed discount factor δ^*, consider the subset Δ_{δ^*} of discount factors that first-order stochastically dominate δ^*, which we interpret as more impatient than δ^* decision-makers. This set is convex and the extreme-points approach yields clean characterizations – see the supplementary appendix B.4.

2.4.2 Application: Comparing the Arrival Time of News

We now return to our main application and derive simple sufficient conditions for Δ-sufficiency and $\Delta_{\downarrow}$-sufficiency. Recall that F and G are the distributions over arrival times and that δ-sufficiency is equivalent to $\sum_{t=1}^{T} \delta_t F_t \geq \sum_{t=1}^{T} \delta_t G_t$. The following result is then immediate:

Corollary 3 *The experiment f is Δ-sufficient for the experiment g if, and only if, G first-order stochastically dominates F.*

Thus, if we are interested in a comparison that holds independently of the discount factor, it suffices that G first-order stochastically dominates F. In most economic applications, discount factors are assumed to be decreasing over time, reflecting the idea that the later the decision-maker is paid, the bigger the payment must be. Now, if $(\delta_t)_t$ is a decreasing sequence, the function $s \mapsto \sum_{t=s}^{T} \delta_t$ is decreasing and convex, so a (weaker) sufficient condition is $\boldsymbol{T}_G$ dominating $\boldsymbol{T}_F$ in the increasing concave order. Since this is equivalent to G second-order stochastically dominating F – see Shaked and Shanthikumar (2007), section 4.A – we have the following:

Corollary 4 *The experiment f is $\Delta_{\downarrow}$-sufficient for the experiment g if, and only if, G second-order stochastically dominates F.*

We stress that all we require is for the discount factor to be decreasing. In particular, this applies to geometric discounting. However, geometric discounting satisfies an even stronger property, namely that the map $t \mapsto \delta_t$ is convex. Thus, a decision-maker would prefer receiving the prize $V^*_{p,h}(u)$ at random time $\boldsymbol{T}_G$ than receiving it for sure at time $\mathbb{E}[\boldsymbol{T}_G]$ – a mean-preserving contraction of $\boldsymbol{T}_G$. Although geometric discounting does not fall within the scope of Proposition 2, it nevertheless admits a sharp characterization. If we want the comparison to hold for geometric discounting, we need $\mathbb{E}[\rho^{\boldsymbol{T}_F}] \geq \mathbb{E}[\rho^{\boldsymbol{T}_G}]$ for all $\rho \in (0,1)$ – see Shaked and Shanthikumar (2007), section 5.A.

With geometric discounting, decision-makers are risk-seeking over the arrival time of prizes, an assumption we do not need for a comparison in terms of mean-preserving spreads. The reason is simple. The information a decision-maker receives at period t remains valuable in later periods (unless the state is independent across periods).

Thus, what matters is the cumulative effect, not the instantaneous effect, of receiving the information.[12]

2.5 History Dependence

In many economic problems, the choices available to a decision-maker at a period depend on past choices. For instance, in saving-consumption problems, past saving decisions constrain current and future consumption. As another instance, irreversibilities in investment have lasting consequences on capital expenditures, particularly when firms face uncertainties about market conditions. As yet another instance, in stopping problems, the decision to stop exploring/gathering information is typically assumed irreversible. This subsection studies two forms of history dependence that do not affect the evolution of information: history-dependent choice sets and history-dependent discount factors. The case in which the information available to the decision-maker depends on past decisions is the object of Section 3.

As already mentioned, discounted problems are instances of sequential problems, a larger class of problems. Greenshtein (1996) considers the larger class of *all* sequential problems, where the decision-maker is constrained to choose sequences of decisions in $\mathcal{A} \subseteq A_1 \times \cdots \times A_t \times \cdots \times A_T$, and his payoff is an arbitrary function of the sequences of decisions made, that is, $u\colon \bigtimes_t A_t \times \Theta \to \mathbb{R}$.[13] Greenshtein (1996, Theorem 1.1a) proves that the experiment f is more valuable than the experiment g for all sequential problems if, and only if, g is a garbling of f, where the garbling is *adapted*. Informally, a garbling is adapted if, to simulate g^t, the decision-maker first draws y_1 with probability $\gamma_1(y_1|x_1)$, next draws y_2 with probability $\gamma_2(y_2|x^2, y^1)$, and then iteratively draws $y_{t'}$ with probability $\gamma_{t'}(y_{t'}|y^{t'-1}, x^{t'})$ for all $t' \leq t$ when the history of simulated and realized signals is (y^{t-1}, x^t).

In other words, unlike what we did, the decision-maker does not simulate new fictitious histories of signals at each period. Instead, at each period, he completes the history of already simulated signals. This is a measurability condition. With the help of an example, Greenshtein (1996) shows that the condition is stronger than Δ-sufficiency.[14]

[12] We refer the reader to DeJarnette et al. (2020) for a study on risk-seeking over time lotteries.

[13] When $T = +\infty$, we require u to be bounded and Borel measurable and $\mathcal{A}$ to be Borel, with respect to the product topology. In constrained problems, strategies must induce a sequence of decisions in $\mathcal{A}$ with probability one in every state.

[14] Greenshtein (1996) also provides conditions linking the two notions of sufficiency.

Since Δ-sufficiency is stronger than δ-sufficiency, the comparisons we have obtained so far are finer than the one of Greenshtein (1996). The results below show that, in discounted problems, sufficiently rich history dependence bring the comparison back to the stronger one of Greenshtein (1996).

2.5.1 History Dependence of the Choice Sets

We assume that the decision-maker is constrained to choose decisions in $\mathcal{A} \subseteq A_1 \times \cdots \times A_t \times \cdots \times A_T$, the set of feasible sequences of decisions. E.g., if $\mathcal{A} = \{(a_1, a_2), (a_1', a_2), (a_1', a_2')\}$, the decision-maker can choose either a_1 or a_1' in the first period, but after choosing a_1 in the first period, the only option in the second period is a_2. A constrained δ-discounted problem is then a triple $(A, \mathcal{A}, u)$, with $\mathcal{A} \subseteq \underbrace{A \times \cdots \times A \times \cdots \times A}_{T \text{ times}}$. A discount factor $\delta \in \Delta$ has full support if $\delta_t > 0$ for every $t \leq T$. We have the following theorem.

Theorem 4 *Fix any full-support discount factor $\bar{\delta} \in \Delta$. The following statements are equivalent:*

1. *The experiment f is more valuable than g in all constrained $\bar{\delta}$-discounted problems.*
2. *For every $\delta \in \Delta$, the experiment f is more valuable than g in all constrained δ-discounted problems.*
3. *The experiment f is more valuable than g in all sequential problems.*
4. *The experiment f is sufficient for g in the sense of Greenshtein (1996).*

The equivalence with the sufficiency condition of Greenshtein (1996) is not entirely surprising as we can think of constrained decision problems as requiring that (i) $U((a_1, \ldots, a_T), \theta) = -\infty$ for all sequences $(a_1, \ldots, a_T) \notin \mathcal{A}$, and (ii) $U((a_1, \ldots, a_T), \theta) = \sum_t \delta_t u(a_t, \theta)$ for all $(a_1, \ldots, a_T) \in \mathcal{A}$. The richness of the class of constrained discounted problems then generates sufficiently many inter-temporal linkages that the comparison becomes equivalent to the one of Greenshtein (1996). In the intermediate case, where $\mathcal{A}$ is a product-space, i.e. $\mathcal{A} = A_1 \times \cdots \times A_t \times \cdots \times A_T$, with $A_t \subseteq A$ for all t, we obtain the equivalence with the class of separable problems, hence Δ-sufficiency. Again, the intuition is simple. In this intermediate case, there are no inter-temporal linkages and we can view the constraints as requiring that (i) $u_t(a, \theta) = -\infty$ if $a \notin A_t$ and (ii) $u_t(a, \theta) = u(a, \theta)$ if $a \in A_t$.

2.5.2 History Dependence of the Discount Factors

Before analyzing the implication of history-dependent discount factors for the comparison of (uncontrolled) experiments, we need to comment on the comparisons we make. So far, we have compared experiments and decision problems, which all have the same duration T. Clearly, it is without loss of generality to assume that the two experiments have the same duration: we can always extend the duration of the shortest experiment to the longest one by adding "dummy" periods, where no additional information is provided. The same is not necessarily true for the decision problems. With history-dependent discount factors, problems with a longer duration than the two experiments matter.

To see this, we consider the class of decision problems, where the decision-maker evaluates sequences $(a_1, \dots, a_t, \dots)$ by $\sum_t \delta_t(a^{t-1})u(a_t, \theta)$, with $\delta_t(a^{t-1})$ the "discount factor" at time t, when the sequence of past choices is a^{t-1}. Models with history-dependent discount factors have antecedents in macroeconomics, e.g., Uzawa (1968) and Schmitt-Grohe and Uribe (2003). We can interpret $\delta_t(a^{t-1})$ as the probability that choice a_t matters at period t when the history of past choices is a^{t-1}. In repeated games with geometric discounting, it is common to interpret the discount factor as the probability that the interaction continues one more period. An analogous interpretation is possible if we let $\delta(a)$ be the probability that the problem continues one more period when the current choice is a.[15]

A history-dependent discounted problem is a triplet (A, u, δ), where $\delta_t \colon A^{t-1} \to [0, 1]$ for all t. It has finite length $\overline{T} < +\infty$ if $\delta_t = 0$ for every $t > \overline{T}$. When the problem does not have finite length, we require $\sum_t \delta_t(a^{t-1})u(a_t, \theta)$ to converge and be bounded uniformly over sequences of decisions and states. We note that, because the discount factor depends on past choices, we cannot change the set of choices A without changing the discount factor. So, when considering all history-dependent discounted problems, we consider all history-dependent discount factors. When $T < +\infty$, we write f^+ (resp., g^+) for the experiment that agrees with f (resp., g) over the first T periods and is uninformative thereafter. When $T = +\infty$, we let $f^+ = f$ and $g^+ = g$.

Theorem 5 *The following statements are equivalent:*

[15] In this case, $\delta_t(a^{t-1}) := \prod_{s \leq t-1} \delta(a_s)$.

1. *The experiment f^+ is more valuable than the experiment g^+ in all history-dependent discounted problems.*
2. *If $T < +\infty$, the experiment f^+ is more valuable than the experiment g^+ in all history-dependent discounted problems of length $T+1$. If $T = +\infty$, the experiment f is more valuable than the experiment g in all history-dependent discounted problems of finite length.*
3. *The experiment f is sufficient for g, in the sense of Greenshtein (1996).*

To get some intuition, assume that $T < +\infty$. The history-dependence of the discount factor induces terms $\delta_T(a^{T-1})u(a_T, \theta)$ in the decision-maker's evaluation, which are instances of general payoffs $\widehat{u}(a_1, \ldots, a_T, \theta)$. For general payoffs, a necessary and sufficient condition is the sufficiency of f for g in the sense of Greenshtein (1996). Now, the class of sequential problems is richer than the class of history-dependent discounted problems of the same duration. What we show is that if we add one more period, we can use the additional period to reward the decision-maker for choosing strategies that would be optimal in a sequential problem of duration T, and punish deviations, otherwise. The dependence of the discount factors on past actions is the leverage we use to reward/punish decisions. When $T = +\infty$, applying the construction at every finite date and the compactness argument in the proof of Theorem 4 implies the existence of an adapted garbling as in Greenshtein (1996). See Proposition 7 for the formal finite-horizon construction.

In summary, the main message is: with sufficiently rich history dependence, comparing experiments in discounted problems is as demanding as comparing experiments in all sequential problems.

3 Discounted Problems: The Controlled Case

So far, the analysis rests on the assumption that the information the decision-maker receives does not depend on the decisions made. While this assumption may be natural in some applications, it is restrictive in many others. For instance, in dynamic pricing problems, the prices a seller chooses impact not only the realized profits, but also what the seller learns about the demand. More generally, the assumption prevents the decision-maker from learning from realized payoffs over time. This section relaxes this assumption.

3.1 Framework

The most natural idea to model controlled information is to incorporate the decisions into the statistical experiments, i.e., to write kernels as $f_t : \Theta \times X^{t-1} \times A^{t-1} \to \Delta(X_t)$. There is, however, a problem with this modeling: the statistical experiments change as we change the decision problems. To extend the approach of Blackwell (1951, 1953), we must be able to vary the decision problems without changing the experiments to be compared.

The solution we adopt parallels the classical distinction between decisions/actions and consequences. For instance, in mechanism design, there is a set of messages (bids, votes, etc) players can choose from and a map from messages to consequences (allocations, elected candidates, etc) – messages are not consequences. A similar distinction is made in decision theory. We follow a similar approach in that we distinguish between decisions and their *informational* consequences. More precisely, we introduce sets of (informational) controls K_t and maps κ_t from actions to controls. The sets of controls are fixed and enter into the definition of statistical experiments, while the maps from actions to controls are part of the decision problem.

We view the controls as the subset of consequences, which directly influence the information the decision-maker receives over time. For instance, in a simple dynamic pricing problem, a seller chooses prices (the decisions), which determine the realized demands, points on the demand curve (the controls) the seller is learning about. In a more realistic pricing problem, a seller chooses not only prices, but also marketing strategies. Different strategies, e.g., a mass campaign in national newspapers or targeted campaigns through social media, reach consumers differently, thus informing some consumers about price changes, while leaving others unaware. In such a problem, the seller only learns about the demand of consumers made aware of the prices offered. Prices and marketing strategies (the decisions) determine, possibly probabilistically, the sample of informed consumers and its demand (the controls), a noisy signal about the true demand (the state). In what follows, we give a formal description of the solution we adopt and derive a characterization for the comparison of statistical experiments.

Experiments. An experiment $f \coloneqq (f_1, f_2, \dots)$ is a collection of kernels $f_t \colon \Theta \times K^{t-1} \times X^{t-1} \to \Delta(X_t)$, where X_t is the set of signals at period t and K_t the set of controls. We

write $f^t(x^t|k^{t-1},\theta)$ for $f_1(x_1|k^0,x^0,\theta)\times\cdots\times f_t(x_t|x^{t-1},k^{t-1},\theta)$.[16] Throughout, we compare the experiments f and g, where the experiment g is the collection of kernels $g_t\colon \Theta\times K^{t-1}\times Y^{t-1}\to\Delta(Y_t)$.

Decision problems. A decision problem is a tuple (A,u,κ), where A is the set of decisions, $u\colon A\times\Theta\to\mathbb{R}$ the payoff function, and $\kappa=(\kappa_t)_t$ with each $\kappa_t\colon A\times K^{t-1}\to\Delta(K_t)$ a control map. The control map summarizes how the decisions influence the information the decision-maker receives over time.

Strategies. Fix the decision problem (A,u,κ). We write σ (resp., τ) for the strategy adapted to the experiment f (resp., g), that is, σ is the collection of maps $(\sigma_t\colon A^{t-1}\times X^t\times K^{t-1}\to\Delta(A))_t$ (resp., $(\tau_t : A^{t-1}\times Y^t\times K^{t-1}\to\Delta(A))_t$). We denote $\mathbb{P}_{\theta,f,\kappa,\sigma}$ the probability over histories of signals, controls and actions, induced by the strategy σ and the experiment f, when the state is θ, and write $\mathbb{E}_{\theta,f,\kappa,\sigma}$ for the expectation with respect to that probability. We define $\mathbb{P}_{\theta,g,\kappa,\tau}$ and $\mathbb{E}_{\theta,g,\kappa,\tau}$ similarly.

We make three important assumptions. First, we assume that the decision-maker observes the realized controls. This assumption allows us to focus on beliefs about the state, which is our main interest. Without this assumption, beliefs about realized controls would also be instrumental. Second, we assume that the control maps do not depend on the state. This assumption is needed to maintain a sharp dichotomy between statistical experiments and decision problems. If the control maps were to depend on the state, the realized controls would also be informative signals about the state – an additional statistical experiment. Third, we assume that only the current choice affects the current control; past choices affect the current control only through the realization of past controls K^{t-1}. We now revisit Sections 2.3 and 2.4, assuming controlled experiments.

3.2 Fixed Discount Factors

Throughout this subsection, if δ has infinite support, we assume that δ has full support. As in the uncontrolled case, f is more valuable than g for a fixed discount factor δ if for all decision problems (A,u,κ), for all strategies τ, there exists a strategy σ, which gives a (weakly) higher payoff in all states. It is immediate to see that f is more valuable than

[16] We adopt the convention: $K_0=\{k_0\}$, $X_0=\{x_0\}$ and $Y_0=\{y_0\}$.

g if, and only if, for all discounted problems (A, u, κ), for all strategies τ, there exists a strategy σ such that

$$\sum_t \delta_t \text{marg}_{A_t} \mathbb{P}_{\theta,g,\kappa,\tau} = \sum_t \delta_t \text{marg}_{A_t} \mathbb{P}_{\theta,f,\kappa,\sigma}, \text{ for all } \theta.$$

Moreover, we can still write the discounted probability of a as:

$$\sum_t \delta_t \mathbb{P}_{\theta,g,\kappa,\tau}(\boldsymbol{a}_t = a) = \sum_{t,y^t,k^{t-1}} \delta_t \mathbb{P}_{\theta,g,\kappa,\tau}(y^t, k^{t-1}) \mathbb{P}_{\theta,g,\kappa,\tau}(\boldsymbol{a}_t = a | y^t, k^{t-1}).$$

Here, $\mathbb{P}_{\theta,g,\kappa,\tau}(y^t, k^{t-1})$ denotes the probability of $\{\boldsymbol{y}^t = y^t, \boldsymbol{k}^{t-1} = k^{t-1}\}$, so $\sum_{y^t,k^{t-1}} \mathbb{P}_{\theta,g,\kappa,\tau}(y^t, k^{t-1}) = 1$ for every t. We use the same convention for all other histories below.

Unlike the uncontrolled case, however, it is no longer true that the probability $\mathbb{P}_{\theta,g,\kappa,\tau}(y^t, k^{t-1})$ is independent of the strategy τ, since the strategy influences the realization of the controls and, consequently, of the signals. Yet, we now argue that a similar idea continues to work.

We start with an informal exposition. Observe that the probability measure $\mathbb{P}_{\theta,g,\kappa,\tau}$ induces kernels $\xi_t^{\boldsymbol{y}} : Y^t \times K^{t-1} \to \Delta(K_t)$, defined by $\xi_t^{\boldsymbol{y}}(k_t | y^t, k^{t-1}) := \mathbb{P}_{\theta,g,\kappa,\tau}(k_t | y^t, k^{t-1})$ whenever $\mathbb{P}_{\theta,g,\kappa,\tau}(y^t, k^{t-1}) > 0$. In turn, this defines an *extended* experiment $\mathbb{P}_{\theta,g,\xi^{\boldsymbol{y}}}$ over profiles of signals and controls, where the probability of (y^t, k^{t-1}) is

$$g_1(y_1|\theta) \times \xi_1^{\boldsymbol{y}}(k_1|y^1, k^0) \times \cdots \times g_{t-1}(y_{t-1}|y^{t-2}, k^{t-2}, \theta) \times \xi_{t-1}^{\boldsymbol{y}}(k_{t-1}|y^{t-1}, k^{t-2}) \times g_t(y_t|y^{t-1}, k^{t-1}, \theta),$$

when the state is θ. (We let K_0 be the singleton $\{k_0\}$.) Analogously, the probability measure $\mathbb{P}_{\theta,f,\kappa,\sigma}$ defines the kernels $\xi_t^{\boldsymbol{x}} : X^t \times K^{t-1} \to \Delta(K_t)$ and the extended experiment $\mathbb{P}_{\theta,f,\xi^{\boldsymbol{x}}}$. The broad idea is to treat the extended experiments $\mathbb{P}_{\theta,g,\xi^{\boldsymbol{y}}}$ and $\mathbb{P}_{\theta,f,\xi^{\boldsymbol{x}}}$ as two *uncontrolled* statistical experiments and to apply the result from Section 2.3.

There are two important caveats, though. First, as the decision problem varies, so does κ and, consequently, the kernels $\xi^{\boldsymbol{y}}$ and $\xi^{\boldsymbol{x}}$. To see this, simply note that $\xi_1^{\boldsymbol{y}}(k_1|y^1, k^0) = \sum_{a_1} \tau_1(a_1|y_1)\kappa_1(k_1|a_1, k^0)$. The solution is to consider a collection of such kernels. Second, and most importantly, the kernels $\xi^{\boldsymbol{y}}$ and $\xi^{\boldsymbol{x}}$ are not independent. As already mentioned, they both depend on the same κ. There is another, more subtle, interdependence. Recall that the aim is to emulate the discounted distribution over actions of any strategy τ

(adapted to g) by a strategy σ (adapted to f), i.e., for all τ, there must exist σ such that

$$\sum_t \delta_t \text{marg}_{A_t} \mathbb{P}_{\theta,g,\kappa,\tau} = \sum_t \delta_t \text{marg}_{A_t} \mathbb{P}_{\theta,f,\kappa,\sigma}.$$

So, suppose, as we do in static problems, that we emulate the distribution over the first-period actions by first sampling a fictitious y_1 with probability $\gamma_1(y_1|x_1)$ when the first-period signal is x_1 and then choosing a_1 with probability $\tau_1(a_1|y_1)$, that is, σ_1 is given by $\sigma_1(a_1|x_1) \coloneqq \sum_{y_1} \tau_1(a_1|y_1)\gamma_1(y_1|x_1)$. It follows that

$$\xi_1^{\boldsymbol{x}}(k_1|x^1,k^0) = \sum_{y_1} \underbrace{\left(\sum_{a_1} \kappa_1(k_1|a_1,k^0)\tau_1(a_1|y^1)\right)}_{=\xi_1^{\boldsymbol{y}}(k_1|y^1,k^0)} \gamma_1(y_1|x_1),$$

that is, $\xi_1^{\boldsymbol{x}}$ is a "garbled" version of $\xi_1^{\boldsymbol{y}}$. Emulating τ with the strategy σ forces the two kernels $\xi^{\boldsymbol{y}}$ and $\xi^{\boldsymbol{x}}$ to be interdependent. For another perspective, suppose that the strategy τ always picks a, while a is associated with relatively "uninformative" controls under f.[17] Then, playing a under σ limits the informativeness of future X-signals. In other words, using information from f at early periods to simulate the information from g constrains the information available at later dates. The characterization needs to capture this dynamic effect.

We turn to a formal exposition. Let $\xi \coloneqq (\xi_t)_t$ be a *test kernel* with $\xi_t : \bigcup_{t'=1}^{T}(Y^{t'} \times K^{t'-1}) \times K^{t-1} \to \Delta(K_t)$ and Ξ the set of all test kernels. We interpret the kernel ξ_t as a map from profiles of *simulated* signals and controls $(y^{t'}, l^{t'-1})$ and *realized* controls k^{t-1} to probabilities over controls at period t. Given ξ, the probability $\mathbb{P}_{\theta,g,\,\xi}$ over profiles of signals and controls is defined by:

$$\begin{aligned}\mathbb{P}_{\theta,g,\,\xi}(y^t, l^{t-1}) =& g_1(y_1|\theta) \times \xi_1(l_1|(y^1,l^0),l^0) \times \cdots \times \\ & g_{t-1}(y_{t-1}|y^{t-2},l^{t-2},\theta) \times \xi_{t-1}(l_{t-1}|(y^{t-1},l^{t-2}),l^{t-2}) \times \\ & g_t(y_t|y^{t-1},l^{t-1},\theta),\end{aligned}$$

for all $(y^t, l^{t-1}) \in Y^t \times K^{t-1}$. The definition of a test kernel may at first be surprising – a reader would perhaps expect it to be defined from $Y^t \times K^{t-1}$ to $\Delta(K_t)$, as we did with $\xi^{\boldsymbol{y}}$. The need to adopt a more complex definition follows from the dual role

[17] That is, for these controls k^{t-1}, the experiments $f_t(\cdot|\cdot,k^{t-1},\theta)$ are not too informative.

garblings have, that is, garblings emulate strategies and, simultaneously, influence the flow of information through the controls. This parallels the dual role decisions have, as determining payoffs and, simultaneously, influencing the flow of information. Formally, as in the uncontrolled case, consider the profile of garblings $\gamma \coloneqq (\gamma_t)_t$, where $\gamma_t : X^t \times K^{t-1} \to \Delta\left(\bigcup_{t'=1}^{T}(Y^{t'} \times K^{t'-1})\right)$ for each $t = 1, \ldots, T$. Given (ξ, γ), we define the probability $\mathbb{P}_{\theta,f,\xi\circ\gamma}$ over signals and controls by:

$$\begin{aligned}\mathbb{P}_{\theta,f,\xi\circ\gamma}(x^t, k^{t-1}) =& f_1(x_1|\theta) \times \left(\sum_{t',y^{t'},l^{t'-1}} \xi_1(k_1|(y^{t'}, l^{t'-1}), k^0)\gamma_1(y^{t'}, l^{t'-1}|x^1, k^0)\right) \times \cdots \times \\ & f_{t-1}(x_{t-1}|x^{t-2}, k^{t-2}, \theta) \times \left(\sum_{t',y^{t'},l^{t'-1}} \xi_{t-1}(k_{t-1}|(y^{t'}, l^{t'-1}), k^{t-2})\gamma_{t-1}(y^{t'}, l^{t'-1}|x^{t-1}, k^{t-2})\right) \times \\ & f_t(x_t|x^{t-1}, k^{t-1}, \theta),\end{aligned}$$

for all (x^t, k^{t-1}). In words, given (x^{t-1}, k^{t-2}), the decision-maker *simulates* a fictitious history $(y^{t'}, l^{t'-1})$ of signals and controls with probability $\gamma_{t-1}(y^{t'}, l^{t'-1}|x^{t-1}, k^{t-2})$, then draws the control k_{t-1} with probability $\xi_{t-1}(k_{t-1}|(y^{t'}, l^{t'-1}), k^{t-2})$, that is, according to the test function ξ_{t-1}, given the simulated history $(y^{t'}, l^{t'-1})$ and the actual realization of past controls k^{t-2}, and then receives the signal x_t with probability $f_t(x_t|x^{t-1}, k^{t-1}, \theta)$, when the state is θ.[18] The probability of transitioning from (x^{t-1}, k^{t-2}) to (x^{t-1}, k^{t-1}) is therefore

$$(\xi_{t-1} \circ \gamma_{t-1})(k_{t-1}|x^{t-1}, k^{t-2}) \coloneqq \sum_{t',y^{t'},l^{t'-1}} \gamma_{t-1}(y^{t'}, l^{t'-1}|x^{t-1}, k^{t-2})\xi_{t-1}(k_{t-1}|(y^{t'}, l^{t'-1}), k^{t-2}).$$

For any collection of test kernels ξ and garblings γ, we thus generate two linked processes $\mathbb{P}_{\theta,g,\,\xi}$ and $\mathbb{P}_{\theta,f,\xi\circ\gamma}$ over signals and controls, which we then treat as *uncontrolled* experiments. We then apply the condition of δ-sufficiency on these two experiments, with the additional consistency requirement that the garblings defining $\mathbb{P}_{\theta,f,\xi\circ\gamma}$ coincide with the garbling defining δ-sufficiency.

Theorem 6 *The following statements are equivalent:*

(1) The experiment f is more valuable than the experiment g in all δ-discounted problems.

[18] We stress that t' may differ from t: At period t, the decision-maker simulates histories of signals and controls $(y^{t'}, l^{t'-1})$ of any length t'.

(2) For all decision problems (A, u, κ)*, for all strategies* τ*, there exists a strategy* σ *such that*

$$\sum_t \delta_t \text{marg}_{A_t} \mathbb{P}_{\theta,g,\kappa,\tau} = \sum_t \delta_t \text{marg}_{A_t} \mathbb{P}_{\theta,f,\kappa,\sigma}, \textit{for all } \theta.$$

(3) δ*-sufficiency: For all* $\xi \in \Xi$*, there exists a profile of kernels* $\gamma = (\gamma_{t'})_{t'=1}^T$*, with* $\gamma_{t'} : X^{t'} \times K^{t'-1} \to \Delta\left(\bigcup_{t=1}^T (Y^t \times K^{t-1})\right)$ *for each* $t' = 1, \ldots, T$*, such that*

$$\delta_t \mathbb{P}_{\theta,g,\xi}(y^t, l^{t-1}) = \sum_{t'=1}^T \sum_{x^{t'} \in X^{t'}, k^{t'-1} \in K^{t'-1}} \delta_{t'} \mathbb{P}_{\theta,f,\xi\circ\gamma}(x^{t'}, k^{t'-1}) \gamma_{t'}(y^t, l^{t-1} | x^{t'}, k^{t'-1}), \textit{for all } (t, y^t, l^{t-1}, \theta).$$

It is instructive to reformulate the third statement as follows: for all ξ, there exist garblings γ and $\widehat{\gamma}$ such that

$$\delta_t \mathbb{P}_{\theta,g,\xi}(y^t, l^{t-1}) = \sum_{t', x^{t'}, k^{t'-1}} \delta_{t'} \mathbb{P}_{\theta,f,\xi\circ\widehat{\gamma}}(x^{t'}, k^{t'-1}) \gamma_{t'}(y^t, l^{t-1} | x^{t'}, k^{t'-1}), \text{ for all } (t, y^t, l^{t-1}, \theta), \quad (2)$$

and $\widehat{\gamma} = \gamma$. Write $\mathbb{P}^t_{\theta,g,\xi}$ for the marginal of $\mathbb{P}_{\theta,g,\xi}$ over $Y^t \times K^{t-1}$ and, analogously, for $\mathbb{P}^t_{\theta,f,\xi\circ\widehat{\gamma}}$. The role of the garbling γ is then to simulate the experiment $\sum_t \delta_t \mathbb{P}^t_{\theta,g,\xi}$ from the experiment $\sum_t \delta_t \mathbb{P}^t_{\theta,f,\xi\circ\widehat{\gamma}}$, as in Blackwell (1951) and Section 2.3.[19] We stress that the garbling γ depends on both the signals and controls because not only signals, but also controls, are informative about the states. As already mentioned, however, using information from one controlled experiment f to simulate another controlled experiment g also has an endogenous effect on the informativeness of f. The role of the garbling $\widehat{\gamma}$ is to capture that effect. Finally, the consistency requirement $\widehat{\gamma} = \gamma$ stipulates that a unique garbling should play these two roles.[20]

3.2.1 Application: Stationary Experiments

As an application of Theorem 6, we consider the class of *stationary experiments*. The experiments f and g are *stationary* if there exist a fixed set of controls K and two static experiments $h_f \colon \Theta \times K \to \Delta(X)$ and $h_g \colon \Theta \times K \to \Delta(Y)$ such that $K_t = K$, $X_t = X$,

[19] Note that if each K_t is a singleton, then $\mathbb{P}^t_{\theta,g,\xi}$ (resp., $\mathbb{P}^t_{\theta,f,\xi\circ\widehat{\gamma}}$) is $g^t(\cdot|\theta)$ (resp., $f^t(\cdot|\theta)$).

[20] We also note that the following algorithm allows one to check for δ-sufficiency: for each $\widehat{\gamma}$, associate all garblings γ, if any, such that Equation (2) holds, a linear programming problem. This defines a convex-valued correspondence from the set of garblings to its power set. If for all ξ, such a correspondence admits a fixed point, then δ-sufficiency holds.

$Y_t = Y$,

$$f_t(x_t|k^{t-1}, x^{t-1}, \theta) = h_f(x_t|k_{t-1}, \theta), \qquad \text{and} \qquad g_t(y_t|k^{t-1}, y^{t-1}, \theta) = h_g(y_t|k_{t-1}, \theta),$$

for all $x_t, y_t, x^{t-1}, y^{t-1}, k^{t-1}, \theta, t$. We fix an arbitrary $k_0 \in K$ so that $f_1(x|\theta) := h_f(x|k_0, \theta)$ and $g_1(y|\theta) := h_g(y|k_0, \theta)$.

Stationary experiments arise naturally in economic applications where the decision-maker repeatedly samples from the same underlying source, as in repeated experimentation and sequential sampling.

For a full-support prior p, recall that $V^*_{p,h}(u)$ is the information value of the static experiment h when (A, u) is the decision problem — see Section 2.3.1. Let h^* be the fully informative experiment and, for each $k \in K$, define

$$R^*_k := \sup_u \frac{V^*_{p,h^*}(u) - V^*_{p,h_f(k,\cdot)}(u)}{V^*_{p,h^*}(u) - V^*_{p,h_g(k,\cdot)}(u)}, \tag{3}$$

where $0/0 = 0$, and a positive number divided by zero is $+\infty$. The number R^*_k is the ratio of the regret of using experiment $h_f(k, \cdot)$ over the regret of using experiment $h_g(k, \cdot)$, with $R^*_k \leq 1$ whenever $h_f(k, \cdot)$ is Blackwell sufficient for $h_g(k, \cdot)$.

Proposition 3 *The experiment f is δ-sufficient for g if*

$$R^*_{k_0} \leq \sum_{t=1}^{T} \delta_t \big(\max\{1, R^*_k\}\big)^{1-t} \qquad \textit{for every } k \in K.^{21} \tag{4}$$

Several remarks are worth making. First, condition (4) requires that $R^*_{k_0} \leq 1$, i.e., the initial experiment $h_f(k_0, \cdot)$ is Blackwell sufficient for $h_g(k_0, \cdot)$. This implies that the n-fold product experiment $h_f^{\otimes n}(k_0, \cdot)$ is Blackwell sufficient for $h_g^{\otimes n}(k_0, \cdot)$. Thus, in problems where the control is k_0 at all periods, experiment f is more valuable than g, regardless of the discount factor.[22] Second, condition (4) is always satisfied whenever $R^*_k \leq 1$ for all $k \in K$. As we shall see later, the point-wise sufficiency of $h_f(k, \cdot)$ for $h_g(k, \cdot)$ guarantees δ-sufficiency for all $\delta \in \Delta$. Third, suppose that $T = +\infty$ and let $\delta_t = (1 - \rho)\rho^{t-1}$, with

[21] By convention, $(\max\{1, R^*_k\})^0 = 1$ and $(\max\{1, R^*_k\})^{1-t} = 0$ when $R^*_k = +\infty$ and $t > 1$.

[22] We refer the reader to Mu et al. (2021), who study when repeated i.i.d. copies of one static experiment are eventually sufficient for repeated i.i.d. copies of another.

$0 < \rho < 1$. Summing the geometric series, condition (4) becomes

$$R^*_{k_0} \leq \frac{1-\rho}{1-\rho/\max\{1, R^*_k\}} \qquad \text{for every } k \in K,$$

where $1/(+\infty) = 0$. Thus, if $R^*_{k_0} \leq 1-\rho$, no additional restrictions are imposed on R^*_k! Intuitively, at the first period, the (static) regret Reg_f under f is less than $R^*_{k_0} \times \text{Reg}_g$. In all subsequent periods, the regret under f is at most Reg_f (since the decision-maker can repeat the first-period behavior at all periods), while the regret under g is at least 0. For the discounted regret to be lower under f, we need $\text{Reg}_f \leq (1-\rho) \times \text{Reg}_g$. Since $R^*_{k_0} \leq 1-\rho$, this is satisfied. If, however, $R^*_{k_0} > 1-\rho$, we need

$$R^*_k \leq \frac{\rho R^*_{k_0}}{R^*_{k_0} + \rho - 1} \qquad \text{for every } k \in K.$$

Subsequent regrets cannot be too large compared to the initial regret.

3.3 Families of Discount Factors

This section parallels Section 2.4, which offered two main insights. The first insight established that the δ-sufficiency of f for all $\delta \in \mathcal{D}$ guarantees that f is more valuable than g in all corresponding δ-discounted problems – this remains true for controlled problems. The second insight showed that the larger class of separable problems induces the same ranking as the class of all discounted problems, i.e., for all $\delta \in \Delta$. Crucially, this second insight does not extend to controlled problems.

More precisely, a separable controlled problem is a collection $(A_t, u_t, \kappa_t)_t$, where A_t is finite, $u_t \colon A_t \times \Theta \to \mathbb{R}$, and $\kappa_t \colon A_t \times K^{t-1} \to \Delta(K_t)$. The decision-maker evaluates the profile $(a_1, \dots, a_t, \dots)$ as $\sum_t u_t(a_t, \theta)$.[23] The following theorem is the analogue of Theorem 3 for controlled problems.

Theorem 7 *Let f and g be two controlled experiments. The following statements are equivalent:*

(1) The experiment f is more valuable than the experiment g in all separable problems.

(2) For all separable problems $(A_t, u_t, \kappa_t)_t$, for all strategies τ, there exists a strategy σ such that

$$\text{marg}_{A_t} \mathbb{P}_{\theta,g,\kappa,\tau} = \text{marg}_{A_t} \mathbb{P}_{\theta,f,\kappa,\sigma}, \text{ for all } (t,\theta).$$

[23] When $T = +\infty$, we require $\sum_t \max_{a \in A_t, \theta \in \Theta} |u_t(a,\theta)| < +\infty$.

(3) $\mathcal{S}$-sufficiency: For all $\xi \in \Xi$, there exist kernels $\gamma_t : X^t \times K^{t-1} \to \Delta(Y^t \times K^{t-1})$ such that

$$\mathbb{P}_{\theta,g,\xi}(y^t, l^{t-1}) = \sum_{x^t,k^{t-1}} \gamma_t(y^t, l^{t-1}|x^t, k^{t-1})\mathbb{P}_{\theta,f,\xi\circ\gamma}(x^t, k^{t-1}), \text{ for all } (y^t, l^{t-1}, t, \theta).$$

In uncontrolled problems, $\mathcal{S}$-sufficiency is equivalent to Δ-sufficiency, the necessary and sufficient condition for f to be more valuable than g in δ-discounted problems for all $\delta \in \Delta$. Theorem 7 states that $\mathcal{S}$-sufficiency remains sufficient for f to be more valuable than g in δ-discounted controlled problems for all $\delta \in \Delta$. As already mentioned, though, the converse is false. In the supplementary appendix (Proposition 9), we provide an example in which the discounted comparison holds for every $\delta \in \Delta$, but f is not more valuable than g in all separable problems.

3.3.1 Applications: Stationary and Bandit Problems

We first revisit the class of stationary experiments. For such experiments, a simple sufficient condition for $\mathcal{S}$-sufficiency is point-wise Blackwell dominance of the per-period experiments.

Proposition 4 *Let f and g be two stationary experiments with controls K. If $h_f(k,\cdot)$ Blackwell dominates $h_g(k,\cdot)$ for each $k \in K$, then, f is $\mathcal{S}$-sufficient for g.*

The proof is a direct application of the constructive methods underlying Theorem 7. Stationarity allows us to construct the garblings $\gamma_t \colon X^t \times K^{t-1} \to \Delta(Y^t \times K^{t-1})$ recursively from the (static) point-wise Blackwell garblings.

The condition is not necessary. For example, if there exists k^* such that $h_f(k^*,\cdot)$ does not Blackwell dominate $h_g(k^*,\cdot)$, but $h_f(k_0,\cdot)$ is fully informative, then f is clearly Δ-sufficient for g. The condition becomes necessary if we require the comparison to hold uniformly over the initial control. Indeed, if $h_f(k,\cdot)$ fails to Blackwell dominate $h_g(k,\cdot)$ for some k, choose $k_0 = k$ and consider decision problems in which the control is held fixed at k. Putting all discount weight on the first period reduces the comparison to the static comparison between $h_f(k,\cdot)$ and $h_g(k,\cdot)$. By Blackwell's theorem, there is then a decision problem in which the decision-maker is strictly better off with g.

BANDIT PROBLEMS. In a classical bandit problem, a decision-maker repeatedly chooses between arms, with the chosen arm influencing both the current payoff and the information generated for future decisions. Bandit problems are a central framework for studying sequential decision-making under uncertainty, with applications including optimal R&D investment (Roberts and Weitzman, 1981), job search and labor market matching (Jovanovic, 1979; Miller, 1984), dynamic pricing under demand uncertainty (Rothschild, 1974; Bergemann and Välimäki, 1996), and the financing of innovation (Bergemann and Hege, 2005).

To model bandit problems as controlled stationary problems, we identify the choice of an arm (the action) with the vector of (expected) state-contingent payoffs it induces, and we treat this vector as an informational control. The experiment then specifies how the rewards/signals depend on the control and the state, and is held fixed across decision/bandit problems. Formally, an action is a vector $u \in \mathbb{R}^{|\Theta|}$, with the convention that the θ-th coordinate $u(\theta)$ is the expected payoff in state θ.[24] If the decision-maker chooses u_t at period t, the realized control is u_t, i.e., $\kappa_t(u_t|u_t, k_{t-1}) = 1$. A *bandit problem* is a finite set $\{u^1, \dots, u^n\}$, with $u^i \in \mathcal{U} \subset \mathbb{R}^{|\Theta|}$ for all i and $\mathcal{U}$ finite.[25] We interpret choosing u^i as choosing arm i.

The *bandit experiment* f is a stationary experiment whose static experiment $h_f(\cdot \mid u, \theta)$ is a finitely supported distribution over rewards $r \in \mathbb{R}$, conditional on the action/control $u \in \mathbb{R}^{|\Theta|}$ and the state $\theta \in \Theta$. We assume that $\mathbb{E}_{h_f}[r \mid u, \theta] = u(\theta)$ for all u and θ. Conditional on the state θ, the reward r_t at period t is a draw from $h_f(\cdot \mid u_t^i, \theta)$ when the decision-maker chooses the arm u_t^i. The decision-maker observes the realized reward r_t at the beginning of period $t+1$. (We assume that the initial control is u^0 and let $h_f(0|u^0, \theta) = 1$ for all θ, i.e., the first-period signal is uninformative.) As in classical bandit problems, the realized payoff is the signal.

Since $\mathbb{E}_{h_f}[r \mid u, \theta] = u(\theta)$, maximizing $\mathbb{E}_{p,h_f,\sigma}[\sum_t \delta_t \mathbf{r}_t]$ with respect to σ is equivalent to maximizing $\mathbb{E}_{p,h_f,\sigma}[\sum_t \delta_t u_t(\boldsymbol{\theta})]$. An alternative interpretation is, therefore, that the realized payoff at period t is $u^i(\theta)$ when the decision-maker chooses arm i and the state is θ. The decision-maker does not observe the realized payoff. Instead, the decision-maker observes the signal r_t. E.g., $r_t = u^i(\theta) + \varepsilon_t$, with ε_t some white noise. This alternative interpretation is mathematically equivalent to our modeling of stationary problems. We

[24] This modeling follows Blackwell (1951).

[25] The finiteness assumption guarantees that the sets of controls and actions are finite, but our arguments do not rely on it.

define the bandit experiment g similarly. We stress that the two experiments are defined for all $u \in \mathbb{R}^{|\Theta|}$. However, when applying Proposition 4 to compare them, we restrict attention to $u \in \mathcal{U}$, the set of actions/controls of the bandit problem.

Proposition 5 *Let f and g be two bandit experiments. If $h_f(u,\cdot)$ Blackwell dominates $h_g(u,\cdot)$ for each $u \in \mathcal{U}$, then f is more valuable than g for all bandit problems with actions in $\mathcal{U}$. Conversely, if there exists an arm $u^* \in \mathbb{R}^{|\Theta|}$ such that $h_f(u^*,\cdot)$ does not Blackwell dominate $h_g(u^*,\cdot)$, then there exist a finite set $\mathcal{U}^*$ with $u^* \in \mathcal{U}^*$ and a bandit problem with actions in $\mathcal{U}^*$ for which the decision-maker is strictly better off under g than f.*

Two concluding remarks are worth making. First, we note that if two arms satisfy $u^1 = u^2$, the current modeling assumes the same distribution over rewards. Yet, if we interpret u^i as the mean reward, two arms can have the same means, but different variances. We can easily accommodate this by adding an extra label, say for the variance, to each arm. Second, a special case is to assume that the experiments depend only on the θ-th coordinate of u when the state is θ. We illustrate this case with a simple family of bandit problems: the *scalar payoff-noise family*. Assume that $h_f(\cdot|u,\theta)$ takes the value $r = \frac{u(\theta)}{\alpha_f(u(\theta))}$ (resp. $r = 0$) with probability $\alpha_f(u(\theta))$ (resp., $1 - \alpha_f(u(\theta))$), where $\alpha_f : \mathbb{R} \to (0,1]$, $\alpha_f(0) = 1$. Clearly, $u(\theta)$ is the mean of h_f conditional on (u,θ). We define h_g similarly.

Now, if the map $w \mapsto \rho_f(w) := w/\alpha_f(w)$ is injective, observing $r \neq 0$ identifies the mean $u^i(\theta)$ of the chosen arm u^i.[26] Moreover, if $\alpha_f(w) \geq \alpha_g(w)$ for all w, then the garbling

$$\gamma(0|0) = 1, \gamma(0|r) = 1 - \frac{\alpha_g(\rho_f^{-1}(r))}{\alpha_f(\rho_f^{-1}(r))}, \gamma\left(\frac{\rho_f^{-1}(r)}{\alpha_g(\rho_f^{-1}(r))}\middle| r\right) = \frac{\alpha_g(\rho_f^{-1}(r))}{\alpha_f(\rho_f^{-1}(r))}, r \in \{\rho_f(u(\theta)) : \theta \in \Theta\} \setminus \{0\}$$

guarantees that $h_f(u,\cdot)$ Blackwell dominates $h_g(u,\cdot)$. We thus have:

Corollary 5 *Assume that ρ_f is injective. In the scalar payoff-noise family, f is Δ-sufficient for g if $\alpha_f \geq \alpha_g$.*

If $T \geq 2$, the condition is also necessary. To see this, assume that $\alpha_f(w) < \alpha_g(w)$ for some $w \neq 0$. Consider the arm $u^* = (w, 0, \ldots, 0)$. For $h_g(\cdot|u^*,\cdot)$ to be a garbling of $h_f(\cdot|u^*,\cdot)$, we need $\gamma(0|0) = 1$ since $h_g(0|u^*,\theta) = h_f(0|u^*,\theta) = 1$ for all $\theta \neq \theta_1$. Since $\alpha_f(w) < \alpha_g(w)$,

[26] Examples include $\alpha(w) = c \in (0,1]$ (constant), $\alpha(w) = e^{-\lambda|w|}$ (exponential), and $\alpha(w) = 1/(1+\lambda|w|)$ (hyperbolic) for $\lambda > 0$ and all $w \neq 0$.

it follows that $h_g(\cdot|u^*,\cdot)$ cannot be a garbling of $h_f(\cdot|u^*,\cdot)$. Necessity then follows from Proposition 5.

3.4 Time (In)consistency

We now return to the issue of time (in)-consistency. For most of the paper, we compare experiments assuming that the decision-maker *commits* to strategies, that is, we compare

$$C_f(A,u,\kappa,p,\delta) \coloneqq \sup_{\sigma} \mathbb{E}_{p,f,\kappa,\sigma}\left[\sum_{t=1}^{T}\delta_t u(\boldsymbol{a_t},\boldsymbol{\theta})\right] \text{ and } C_g(A,u,\kappa,p,\delta) \coloneqq \sup_{\tau} \mathbb{E}_{p,g,\kappa,\tau}\left[\sum_{t=1}^{T}\delta_t u(\boldsymbol{a_t},\boldsymbol{\theta})\right].$$

for all problems (A,u,κ,p).[27] We call it the commitment solution.

An alternative is to consider an intra-personal game between the multiple selves – one self per period – of the decision-maker. We call it the consistent planning solution. Formally, for each self-s, we let $\delta^{(s)} \in \Delta(\{s,\dots,T\})$ be the discount factor and require self-s choosing $\sigma_s : A^{s-1}\times X^s \times K^{s-1} \to \Delta(A)$ to maximize

$$\mathbb{E}_{p,f,\kappa,(\sigma_s,\sigma_{-s})}\left[\sum_{t=s}^{T}\delta_t^{(s)}u(\boldsymbol{a_t},\boldsymbol{\theta})\middle| a^{s-1},x^s,k^{s-1}\right],$$

for each (a^{s-1},x^s,k^{s-1}), where σ_{-s} is the strategy profile of selves other than self-s. The strategy profile σ is thus a subgame-perfect Nash equilibrium (henceforth, SPNE) of the intra-personal game. We would then compare

$$V_f(A,u,\kappa,p,\boldsymbol{\delta}) \coloneqq \sup_{\sigma\in SPNE} \mathbb{E}_{p,f,\kappa,\sigma}\left[\sum_{t=1}^{T}\delta_t^{(1)}u(\boldsymbol{a_t},\boldsymbol{\theta})\right] \text{ and}$$
$$V_g(A,u,\kappa,p,\boldsymbol{\delta}) \coloneqq \sup_{\tau\in SPNE} \mathbb{E}_{p,g,\kappa,\tau}\left[\sum_{t=1}^{T}\delta_t^{(1)}u(\boldsymbol{a_t},\boldsymbol{\theta})\right],$$

for all problems (A,u,κ,p), where $\boldsymbol{\delta}$ is the profile of discount factors – one for each self.

Naturally, the two comparisons are identical when the profile of discount factors $\boldsymbol{\delta}$ induces time-consistent preferences – a ubiquitous assumption in Economics. More generally, with the simplest uncontrolled problems considered in Section 2.3, the two comparisons are also identical. Yet, as the following example demonstrates, this is no longer

[27] We write $p \in \Delta(\Theta)$ for a fully-supported prior.

true with controlled problems and, more generally, with problems featuring non-trivial inter-temporal leakages.

EXAMPLE. There are three periods $t \in \{1,2,3\}$, two states $\theta \in \{0,1\}$, and two controls $k \in \{k_{\text{no info}}, k_{\text{info}}\}$. The prior probability that the state $\theta = 1$ is $3/4$. The discounting factors are $\delta^{(1)} = (0, 1/2, 1/2)$ for self-1, $\delta^{(2)} = (0,1)$ for self-2, and $\delta^{(3)} = (1)$ for self-3.

We compare the "uninformative" experiment f with the "partially informative" experiment g. The initial control is $k_{\text{no info}}$. When the control at period $t-1$ is k_{info} (resp., $k_{\text{no info}}$), the experiment g generates the signal $y_t \in \{0,1\}$ at period t, where y_t matches the state with probability $q \in (3/4, 1)$ (resp., $1/2$).[28]

We now construct a decision problem, where the consistent-planning solution induces a strictly higher value when the experiment is uninformative. At each of the three periods, the decision-maker has three actions: a_0, a_1, and e. The action e has payoff zero, regardless of the state. The action a_i has payoff 1 if the state is i and zero, otherwise. We interpret a_i as guessing that the state is i. The control map is as follows. At the first period, all actions induce the control $k_{\text{no info}}$. At periods 2 and 3, the action e induces the control k_{info}, while the actions a_0 and a_1 induce the control $k_{\text{no info}}$.

When the experiment is f, the maximum consistent-planning value is $3/4$: the experiment is uninformative, a_1 is the best guess at every information set, and there is an SPNE in which every self chooses a_1. This is the payoff from self-1's viewpoint. When the experiment is g, self-3 chooses the best guess at each information set. Since the initial control is $k_{\text{no info}}$ and the first-period control is fixed at $k_{\text{no info}}$, self-2 has no useful information when choosing at the second period. Self-2 chooses e since self-2 only cares about self-3's payoff and e generates information for self-3. The ex-ante expected payoff from self-1's viewpoint is therefore $(1/2) \times 0 + (1/2) \times q$, which is strictly less than $3/4$. The decision-maker is *strictly* better off with the uninformative experiment f than the partially informative experiment g! Intuitively, self-1 would like self-2 to make a best guess, but self-2 only wants self-3 to make a best guess.

Turning to the commitment solution, it is clear that the decision-maker cannot be worse off with the experiment g: under commitment, the decision-maker can follow the same

[28] Formally, the experiment f is such that $f_t(x_t|x^{t-1}, k^{t-1}, 0) = f_t(x_t|x^{t-1}, k^{t-1}, 1)$ for all (x^{t-1}, k^{t-1}, x_t). The initial control is $k_{\text{no info}}$. The experiment g is such that $g_1(y_1|0) = g_1(y_1|1) = 1/2$ for all y_1, and, for $t = 2,3$, $g_t(0|y^{t-1}, (k^{t-2}, k_{\text{info}}), 0) = g_t(1|y^{t-1}, (k^{t-2}, k_{\text{info}}), 1) = q > 3/4$, while $g_t(y_t|y^{t-1}, (k^{t-2}, k_{\text{no info}}), 0) = g_t(y_t|y^{t-1}, (k^{t-2}, k_{\text{no info}}), 1) = 1/2$ for all (y^{t-1}, y_t).

action and control plan as under f and ignore any additional information. The experiment g is, in fact, $\delta^{(1)}$-sufficient for f. ■

The example illustrates the well-known observation that information may have a negative value when preferences are time inconsistent. Yet, this does not mean that the ranking of experiments is reversed when moving from the *commitment* solution to the *consistent-planning* solution. In the example, the experiments f and g are *not* comparable by the consistent-planning solution. To see this, observe that it is easy to construct other decision problems, where the decision-maker is strictly better off when the experiment is g. This is because g is strictly more informative than f when the control is k_{info}, and we can always choose problems where the realized control is always k_{info}.

The following proposition shows that the ranking of f and g cannot be reversed when moving from the *commitment* solution to the *consistent-planning* solution, if we consider all controlled problems.

Proposition 6 *Let f and g be two controlled experiments. Fix a fully supported prior p. Suppose that for every controlled problem (A, u, κ) and every profile of discount factors $\boldsymbol{\delta}$,*

$$C_f(A, u, \kappa, p, \boldsymbol{\delta}) \geq C_g(A, u, \kappa, p, \boldsymbol{\delta}) \text{ and } V_g(A, u, \kappa, p, \boldsymbol{\delta}) \geq V_f(A, u, \kappa, p, \boldsymbol{\delta}).$$

Then, for every controlled problem (A, u, κ) and every profile of discount factors $\boldsymbol{\delta}$,

$$C_f(A, u, \kappa, p, \boldsymbol{\delta}) = C_g(A, u, \kappa, p, \boldsymbol{\delta}).$$

The proof is simple. First, observe that the commitment solutions only depend on the discount factor $\delta^{(1)}$, and not on the entire profile $\boldsymbol{\delta}$. Second, for any $\delta^{(1)}$, construct the profile $(\delta^{(1)}, (\delta^{(s)})_{s\geq 2})$ by setting $\delta_t^{(s)} := (\sum_{t'\geq s} \delta_{t'}^{(1)})^{-1}\delta_t^{(1)}$ for all $t \geq s$ whenever $\sum_{t'\geq s} \delta_{t'}^{(1)} > 0$, and define $\delta_t^{(s)}$ arbitrarily otherwise. This makes the continuation preferences of every relevant future self coincide with the continuation preferences of self 1, so the commitment and consistent-planning solutions agree. The result then follows immediately.

We stress that the result necessitates considering all profiles of discount factors. We believe that the result remains true for a fixed profile of discount factors. Clearly, this is so whenever the profile of discount factors guarantees time consistency – a commonly made assumption in economics. More generally, if for each profile of controls k^{T-1}, we

define the experiments $f^{k^{T-1}}$ and $g^{k^{T-1}}$ as

$$f_t^{k^{T-1}}(x_t|x^{t-1},\theta) \coloneqq f_t(x_t|x^{t-1},k^{t-1},\theta), \forall x_t, x^{t-1}, \theta, t,$$
$$g_t^{k^{T-1}}(y_t|y^{t-1},\theta) \coloneqq g_t(y_t|y^{t-1},k^{t-1},\theta), \forall y_t, y^{t-1}, \theta, t,$$

then a ranking reversal would imply that $f^{k^{T-1}}$ and $g^{k^{T-1}}$ are δ-sufficient for each other for all k^{T-1}. To see this, note that if we fix $\overline{k}^{T-1} = (\overline{k}_0, \ldots, \overline{k}_{T-1})$ and let $\kappa_t(\overline{k}_t|k^{t-1}, a_t) = 1$ for all (k^{t-1}, a_t), then the problem is equivalent to the uncontrolled problem, where the experiments are $f^{\overline{k}^{T-1}}$ and $g^{\overline{k}^{T-1}}$. Since the commitment and consistent-planning solutions agree for uncontrolled problems, a ranking reversal necessitates $f^{\overline{k}^{T-1}}$ and $g^{\overline{k}^{T-1}}$ to be δ-sufficient for each other. Since this must be true for all $\overline{k}^{T-1}$, this is a strong requirement.

4 Related Literature and Discussion

As we argue in Section 2, there is a tight relationship between the simplest discounted problems and the static problems. Indeed, the condition of δ-sufficiency coincides with the classical sufficiency condition (Blackwell, 1951, 1953) for the comparison of the mixture experiments $\sum_t \delta_t f^t$ and $\sum_t \delta_t g^t$.[29] The relationship with static problems becomes less tight as we move from the simplest discounted problems to more general ones. For instance, static problems do not allow discussing the sequential comparisons of experiments, evolving states, or past decisions influencing the evolution of the signals.

By comparison, the literature on the comparison of statistical experiments in dynamic problems is much more limited. The seminal paper is Greenshtein (1996), which we reviewed in our history-dependence section (Section 2.5). In contrast with Greenshtein (1996), we focus on discounted problems and extend our analysis to, among others, the sequential comparison of experiments, history-dependent discount factors, time-varying states or controlled flow of information – all absent from Greenshtein (1996).[30]

[29] Few results are available for the comparison of mixture experiments. Torgersen (1970) proves that if f^t is sufficient for g^t for all t, then $\sum_t \delta_t f^t$ is sufficient for $\sum_t \delta_t g^t$. He also proves a partial converse, that is, if the experiment f is sufficient for $\sum_t \delta_t g^t$, then there exist experiments f^t such that (i) $f = \sum_t \delta_t f^t$, and (ii) f^t is sufficient for g^t for all t. A related result is that all binary experiments (a.k.a dichotomies) are the mixtures of fully informative experiments – see Birnbaum (1961).

[30] We have generalized the work of Greenshtein (1996) to sequential problems where the decision-maker controls the flow of information. The statement and proof are available upon request.

We conclude with a discussion of a few other related works. Liang et al. (2022) study the optimal allocation of resources across multiple sources of information in a continuous-time stopping problem. (See also Liang et al. (2018).) More precisely, there is a fixed payoff-relevant state given by a linear combination of unknown parameters $(\theta_1, \ldots, \theta_n)$, with a known Gaussian distribution. At each instant, the decision-maker either allocates his (unit) budget to the n sources of information or takes a decision – taking a decision stops the game. The budget allocation influences the diffusion process the decision-maker observes.[31] Under some assumptions, the authors show that the optimal strategy is history-independent and consists in first allocating the entire budget to the most informative source and, subsequently, to invest into additional sources until all sources are funded – investments are made at fixed times.

Moreover, their analysis shows that the strategy is optimal regardless of the decision problem the decision-maker faces when he stops learning. In other words, their strategy induces an experiment, which is sufficient for all other inducible experiments. In that sense, their work generalizes Theorem 3.1 of Greenshtein (1996).[32] Our work differs from theirs in two important aspects. First, we consider different classes of games and do not restrict attention to Gaussian distributions. Second, there is a natural decoupling between the decision problem the decision-maker faces when he stops learning and the choice of experiments (induced by the budget allocation). Such a decoupling is not always possible in our analysis – see Section 3.1. It is also worth noting that since the decision-maker can allocate his budget at each period, their comparison is sequential.

Lehrer and Wang (2024) also study a stopping problem, where the decision-maker either acquires some information or stops and takes an action. At period t, if the decision-maker decides to acquire some information, he pays a fee and obtains a signal from a static experiment f. (Signals are drawn independently and identically across time.) Their main focus is on how an information provider should design the fee structure. They show that upfront fees are optimal. Intuitively, upfront fees do not distort the incentives to acquire information and ensure that the decision-maker purchases as often as possible the information for sale. They also show that if f is sufficient for g, acquiring information is optimal for the decision-maker under a larger set of beliefs when the information is from f. This does not mean, however, that the decision-maker acquires

[31] In discrete time, if the decision-maker allocates q_i to source i, he observes the noisy signal $\theta_i + \varepsilon_i$, where ε_i is Gaussian with zero mean and variance $1/q_i$. Their modeling is the continuous-time analogue.

[32] It is not a generalization of Theorem 1.1a, which is the one we discussed earlier.

information for longer as beliefs also get updated faster when the information is from the most informative experiment f. This latter observation is reminiscent of Example 2.1 in Greenshtein (1996).

Finally, Mu et al. (2021) study when T independent and identical copies of one static experiment are (eventually) sufficient for T independent and identical copies of another.[33] They provide a complete characterization for dichotomies in terms of the Rényi divergence of the static experiments. For earlier related results, see Torgersen (1970), Torgersen (1991), and Moscarini and Smith (2002).

[33] More precisely, they require the existence of $\overline{T}$ such that for all $T \geq \overline{T}$, T independent and identical copies of one static experiment is sufficient for another.

A Proofs

A.1 Proof of Theorem 1 and 2

We first prove Theorem 2 and then derive Theorem 1 as a corollary.

$(2) \implies (1)$. Immediate.

$(1) \implies (2)$. Consider two behavioral strategies σ and σ', with induced distributions $\mathbb{P}_{p,f,\sigma}$ and $\mathbb{P}_{p,f,\sigma'}$, respectively. From Kuhn's theorem, there exist two mixed strategies $\widehat{\sigma}$ and $\widehat{\sigma}'$, which induce the same distributions over states, signals and actions. From the convexity of the set of mixed strategies and the linearity of the measures,

$$\mathbb{P}_{p,f,\alpha\widehat{\sigma}+(1-\alpha)\widehat{\sigma}'} = \alpha\mathbb{P}_{p,f,\widehat{\sigma}} + (1-\alpha)\mathbb{P}_{p,f,\widehat{\sigma}'}$$

for all $\alpha \in [0,1]$. Therefore, the set

$$\left\{ \sum_t \delta_t \text{marg}_{A_t \times \Theta_t} \mathbb{P}_{p,f,\widehat{\sigma}} : \widehat{\sigma} \text{ a mixed strategy} \right\} \subseteq \Delta(A \times \Theta)$$

is convex. (From Kuhn's theorem, we obtain the same set if we consider all behavioral strategies.)

If (2) does not hold, there exists a behavioral strategy τ such that

$$\sum_t \delta_t \text{marg}_{A_t \times \Theta_t} \mathbb{P}_{p,g,\tau} \notin \left\{ \sum_t \delta_t \text{marg}_{A_t \times \Theta_t} \mathbb{P}_{p,f,\sigma} : \sigma \text{ a behavioral strategy} \right\}.$$

As in Blackwell (1951), we can then apply a separation argument to prove the existence of a payoff function u such that (1) is violated, the required contradiction. Compactness follows from the compactness of the set of behavioral strategies in the product topology and the continuity of the map $\sigma \to \sum_t \delta_t \text{marg}_{A_t \times \Theta_t} \mathbb{P}_{p,f,\sigma}$.

$(3) \implies (2)$. Fix a decision problem (A,u) and a strategy τ. Let $(\gamma_{t'})_{t'}$ be the garblings, with $\gamma_{t'} : X^{t'} \to \Delta(\bigcup_t Y^t)$. Note that summing over all elements in $\bigcup_t Y^t$ is the same as summing over all pairs (t, y^t), where the second coordinate is in Y^t. Throughout, we therefore indicate summations over $\bigcup_t Y^t$ by summing over (t, y^t). For all t', for all $x^{t'}$,

for all a, define $\sigma_{t'}$ as:

$$\sigma_{t'}(a|x^{t'}) := \sum_{(t,y^t)} \mathbb{P}_{p,g,\tau}(\boldsymbol{a}_t = a|y^t)\gamma_{t'}(y^t|x^{t'}).$$

From the above observation, this is well-defined as $\sum_{(t,y^t)} \gamma_{t'}(y^t|x^{t'}) = 1$.

Under the strategy $\sigma = (\sigma_{t'})_{t'}$, the discounted probability of choice a is

$$\begin{aligned}
\sum_{t'} \delta_{t'}\mathbb{P}_{p,f,\sigma}(\boldsymbol{a}_{t'} = a, \boldsymbol{\theta}_{t'} = \theta) = \sum_{t',x^{t'}} \delta_{t'}\mathbb{P}_{p,f,\sigma}(\boldsymbol{a}_{t'} = a|x^{t'}, \boldsymbol{\theta}_{t'} = \theta)\mathbb{P}_{p,f,\sigma}(x^{t'}, \boldsymbol{\theta}_{t'} = \theta) = \\
\sum_{t',x^{t'}} \delta_{t'}\mathbb{P}_{p,f,\sigma}(x^{t'}, \boldsymbol{\theta}_{t'} = \theta)\sigma_{t'}(a|x^{t'}) = \\
\sum_{t',x^{t'}} \delta_{t'}\mathbb{P}_{p,f,\sigma}(x^{t'}, \boldsymbol{\theta}_{t'} = \theta)\left(\sum_{t,y^t} \mathbb{P}_{p,g,\tau}(\boldsymbol{a}_t = a|y^t)\gamma_{t'}(y^t|x^{t'})\right) = \\
\sum_{t,y^t} \mathbb{P}_{p,g,\tau}(\boldsymbol{a}_t = a|y^t)\left(\sum_{t',x^{t'}} \delta_{t'}\mathbb{P}_{p,f,\sigma}(x^{t'}, \boldsymbol{\theta}_{t'} = \theta)\gamma_{t'}(y^t|x^{t'})\right) = \\
\sum_{t,y^t} \mathbb{P}_{p,g,\tau}(\boldsymbol{a}_t = a|y^t)\delta_t\mathbb{P}_{p,g,\tau}(y^t, \boldsymbol{\theta}_t = \theta) = \\
\sum_{t} \delta_t\mathbb{P}_{p,g,\tau}(\boldsymbol{a}_t = a, \boldsymbol{\theta}_t = \theta),
\end{aligned}$$

where

$$\delta_t\mathbb{P}_{p,g,\tau}(y^t, \boldsymbol{\theta}_t = \theta) = \sum_{t',x^{t'}} \delta_{t'}\mathbb{P}_{p,f,\sigma}(x^{t'}, \boldsymbol{\theta}_{t'} = \theta)\gamma_{t'}(y^t|x^{t'}),$$

from δ-sufficiency. This follows from the observation that

$$\begin{aligned}
\mathbb{P}_{p,f,\sigma}(x^{t'}, \boldsymbol{\theta}_{t'} = \theta) &= \sum_{\theta_1,\ldots,\theta_{t'},\ldots,\theta_T:\theta_{t'}=\theta} p(\theta_1, \ldots, \theta_{t'}, \ldots, \theta_T) f^{t'}(x^{t'}|\theta_1, \ldots, \theta_{t'}, \ldots, \theta_T), \\
\mathbb{P}_{p,g,\tau}(y^t, \boldsymbol{\theta}_t = \theta) &= \sum_{\theta_1,\ldots,\theta_t,\ldots,\theta_T:\theta_t=\theta} p(\theta_1, \ldots, \theta_t, \ldots, \theta_T) g^t(y^t|\theta_1, \ldots, \theta_t, \ldots, \theta_T).
\end{aligned}$$

(2) $\implies$ (3). Let $A = \bigcup_{t=1}^T Y^t$ and $\tau_t(y^t|y^t, a^{t-1}) = 1$ for all (t, y^t, a^{t-1}). By construction,

$$\sum_{t'} \delta_{t'}\mathbb{P}_{p,g,\tau}(\boldsymbol{a}_{t'} = y^t, \boldsymbol{\theta}_{t'} = \theta) = \delta_t \sum_{\theta_1,\ldots,\theta_t,\ldots,\theta_T:\theta_t=\theta} p(\theta_1, \ldots, \theta_t, \ldots, \theta_T) g^t(y^t|\theta_1, \ldots, \theta_t, \ldots, \theta_T).$$

From (2), there exists $\sigma = (\sigma_{t'})_{t'}$, where $\sigma_{t'} : A^{t'-1} \times X^{t'} \to \Delta(\bigcup_{t=1}^{T} Y^t)$ such that

$$\sum_{t'} \delta_{t'} \mathbb{P}_{p,g,\tau}(\boldsymbol{a}_{t'} = y^t, \boldsymbol{\theta}_{t'} = \theta) = \sum_{t'} \delta_{t'} \mathbb{P}_{p,f,\sigma}(\boldsymbol{a}_{t'} = y^t, \boldsymbol{\theta}_{t'} = \theta).$$

Since

$$\begin{aligned}\sum_{t'} \delta_{t'} \mathbb{P}_{p,f,\sigma}(\boldsymbol{a}_{t'} = y^t, \boldsymbol{\theta}_{t'} = \theta) &= \sum_{t',x^{t'}} \delta_{t'} \mathbb{P}_{p,f,\sigma}(\boldsymbol{a}_{t'} = y^t | x^{t'}, \boldsymbol{\theta}_{t'} = \theta) \mathbb{P}_{p,f,\sigma}(x^{t'}, \boldsymbol{\theta}_{t'} = \theta) \\ &= \sum_{t',x^{t'}} \delta_{t'} \mathbb{P}_{p,f,\sigma}(x^{t'}, \boldsymbol{\theta}_{t'} = \theta) \gamma_{t'}(y^t | x^{t'}),\end{aligned}$$

with $\gamma_{t'}(y^t|x^{t'}) := \mathbb{P}_{p,f,\sigma}(\boldsymbol{a}_{t'} = y^t | x^{t'}, \boldsymbol{\theta}_{t'} = \theta)$ – the latter probability is independent of θ because the probabilities over actions do not depend on the state conditional on the signals. The result then follows immediately from the observation that

$$\begin{aligned}\mathbb{P}_{p,f,\sigma}(x^{t'}, \boldsymbol{\theta}_{t'} = \theta) &= \sum_{\theta_1,\ldots,\theta_{t'},\ldots,\theta_T : \theta_{t'} = \theta} p(\theta_1, \ldots, \theta_{t'}, \ldots, \theta_T) f^{t'}(x^{t'} | \theta_1, \ldots, \theta_{t'}, \ldots, \theta_T) \\ &= \mathbb{P}_{p,f}(x^{t'}, \boldsymbol{\theta}_{t'} = \theta).\end{aligned}$$

If $T < \infty$, this completes the proof of $1 \iff 2 \iff 3$ in Theorem 2.

If $T = +\infty$, let $A = \bigcup_{t=1}^{T_n} Y^t$, where $T_n < +\infty$ and $T_n \to +\infty$. Fix $\overline{y}^{T_n} \in Y^{T_n}$, and let τ_t choose y^t after observing y^t if $t < T_n$, and $\overline{y}^{T_n}$ otherwise. The previous argument shows that there exist $\gamma^n_{t'} : X^{t'} \to \Delta(\bigcup_{t=1}^{T_n} Y^t)$, $t' = 1, \ldots, +\infty$, such that

$$\delta_t \mathbb{P}_{p,g}(y^t, \boldsymbol{\theta}_t = \theta) = \sum_{t'=1}^{+\infty} \sum_{x^{t'}} \delta_{t'} \mathbb{P}_{p,f}(x^{t'}, \boldsymbol{\theta}_{t'} = \theta) \gamma^n_{t'}(y^t | x^{t'}), \tag{A.1}$$

for all y^t, for all $t < T_n$, for all θ. (Note that we have dropped σ and τ from the notation as the probabilities are independent of the strategies.) We extend the garbling $\gamma^n_{t'}$ into the garblings $\gamma^n_{t'} : X^{t'} \to \Delta(\bigcup_{t=1}^{+\infty} Y^t)$ by letting $\gamma^n_{t'}(y^t|x^{t'}) = 0$ whenever $t > T_n$.

By diagonal extraction, we can choose a subsequence, still indexed by n, such that each $\gamma^n_{t'}(y^t|x^{t'})$ converges. This is possible because there are countably many coordinates.

Denote $\gamma_{t'}$ the limit for each t'. From Fatou's lemma, we have

$$\begin{aligned}\sum_{t\geq 1}\sum_{y^t}\gamma_{t'}(y^t \mid x^{t'}) &= \sum_{t\geq 1}\sum_{y^t}\lim_{n\to\infty}\gamma^n_{t'}(y^t \mid x^{t'}) \\ &\leq \liminf_{n\to\infty}\sum_{t\geq 1}\sum_{y^t}\gamma^n_{t'}(y^t \mid x^{t'}) = 1.\end{aligned}$$

Introduce the dummy symbol $\varnothing$ and let

$$\gamma_{t'}(\varnothing \mid x^{t'}) := 1 - \sum_{t\geq 1}\sum_{y^t}\gamma_{t'}(y^t \mid x^{t'}) \geq 0.$$

We will show that the dummy symbol plays no role.

Fix t, θ, and y^t. Let $S < \infty$. For each $T_n > t$, we have:

$$0 \leq \delta_t \mathbb{P}_{p,g}(y^t, \boldsymbol{\theta}_t = \theta) - \sum_{t'=1}^{S}\sum_{x^{t'}}\delta_{t'}\mathbb{P}_{p,f}(x^{t'}, \boldsymbol{\theta}_{t'} = \theta)\gamma^n_{t'}(y^t|x^{t'}) \leq \sum_{t'>S}\delta_{t'}.$$

Indeed,

$$\begin{aligned}\delta_t \mathbb{P}_{p,g}(y^t, \boldsymbol{\theta}_t = \theta) &= \sum_{t'=1}^{+\infty}\sum_{x^{t'}}\delta_{t'}\mathbb{P}_{p,f}(x^{t'}, \boldsymbol{\theta}_{t'} = \theta)\gamma^n_{t'}(y^t|x^{t'}), \\ &= \sum_{t'=1}^{S}\sum_{x^{t'}}\delta_{t'}\mathbb{P}_{p,f}(x^{t'}, \boldsymbol{\theta}_{t'} = \theta)\gamma^n_{t'}(y^t|x^{t'}) \\ &\quad + \sum_{t'=S+1}^{+\infty}\delta_{t'}\underbrace{\sum_{x^{t'}}\mathbb{P}_{p,f}(x^{t'}, \boldsymbol{\theta}_{t'} = \theta)\gamma^n_{t'}(y^t|x^{t'})}_{\leq 1}.\end{aligned}$$

Taking limits as first $T_n \to +\infty$ and then $S \to +\infty$ gives:

$$\delta_t \mathbb{P}_{p,g}(y^t, \boldsymbol{\theta}_t = \theta) = \sum_{t'=1}^{+\infty}\sum_{x^{t'}}\delta_{t'}\mathbb{P}_{p,f}(x^{t'}, \boldsymbol{\theta}_{t'} = \theta)\gamma_{t'}(y^t|x^{t'})$$

It remains to verify that the kernels are fully supported on $\bigcup_{t=1}^{\infty} Y^t$.

Summing the above equalities over (t, y^t, θ) yields

$$\sum_{\theta}\sum_{t'=1}^{+\infty}\sum_{x^{t'}}\delta_{t'}\mathbb{P}_{p,f}(x^{t'}, \boldsymbol{\theta}_{t'} = \theta)\gamma_{t'}(\varnothing|x^{t'}) = 0.$$

Finally, if $\delta_{t'}\mathbb{P}_{p,f}(x^{t'}, \boldsymbol{\theta}_{t'} = \theta) > 0$ for some $\theta \in \Theta$, then $\gamma_{t'}(\varnothing|x^{t'}) = 0$. If $\delta_{t'}\mathbb{P}_{p,f}(x^{t'}, \boldsymbol{\theta}_{t'} = \theta) = 0$ for all $\theta \in \Theta$, we can choose an arbitrary $\gamma_{t'}(\cdot|x^{t'})$ without affecting the equalities.

(1) $\iff$ (4). Fix a decision problem (A, u). Since decisions affect neither the evolution of states and signals nor the availability of subsequent choices, the decision-maker can optimize separately at each date. Consequently

$$\begin{aligned}\sup_{\sigma} \mathbb{E}_{p,f,\sigma}\left[\sum_t \delta_t u(\boldsymbol{a}_t, \boldsymbol{\theta}_t)\right] &= \sum_t \delta_t \int_{\Delta(\Theta)} \max_{a\in A} \sum_{\theta\in\Theta} q(\theta)u(a,\theta)\,\mu^t_{p,f}(dq) \\ &= \int_{\Delta(\Theta)} \max_{a\in A} \sum_{\theta\in\Theta} q(\theta)u(a,\theta)\,\mu^\delta_{p,f}(dq),\end{aligned}$$

where the second equality follows from the definition of $\mu^\delta_{p,f}$. The same calculation applies to g. These equalities remain valid when $T = +\infty$, since u is bounded and $\sum_t \delta_t = 1$. As (A, u) varies, the integrands range over all finite maxima of affine functions on $\Delta(\Theta)$, which uniformly approximate every continuous convex function on $\Delta(\Theta)$. Thus, (1) is equivalent to (4).

Lastly, Theorem 1 follows as an immediate corollary by restricting attention to distributions p such that $p(\theta_1, \ldots, \theta_t, \ldots, \theta_T) > 0$ if and only if $\theta_t = \theta'_t$ for all (t, t'), i.e., by restricting attention to fully persistent states.

A.2 Proof of Proposition 1

This follows immediately from the discussion in the main text.

A.3 Proof of Theorem 3

(2) $\implies$ (1). Every δ-discounted problem is separable, with $u_t(a, \theta) = \delta_t u(a, \theta)$.

(1) $\implies$ (3). Fix t, and let δ put probability one on date t. The mixture experiments in Theorem 1 then reduce to f^t and g^t. Thus, (1) implies that f^t is sufficient for g^t. Since t was arbitrary, (3) follows.

(3) $\implies$ (2). For each t, let $\gamma_t \colon X^t \to \Delta(Y^t)$ be a garbling from f^t to g^t. Fix a separable problem and a strategy τ. For each y^t, let $\bar{\tau}_t(\cdot \mid y^t)$ be the distribution of the date-t action generated by τ after the signal history y^t, after integrating out past actions. Since signals do not depend on actions, this distribution does not depend on the state.

Define the strategy σ by $\sigma_t(a \mid a^{t-1}, x^t) := \sum_{y^t \in Y^t} \gamma_t(y^t \mid x^t)\bar{\tau}_t(a \mid y^t)$ so that, for every t, a, and θ,

$$\mathbb{P}_{\theta,f,\sigma}(\boldsymbol{a}_t = a) = \sum_{x^t,y^t} f^t(x^t \mid \theta)\gamma_t(y^t \mid x^t)\bar{\tau}_t(a \mid y^t) = \sum_{y^t} g^t(y^t \mid \theta)\bar{\tau}_t(a \mid y^t) = \mathbb{P}_{\theta,g,\tau}(\boldsymbol{a}_t = a).$$

Thus, σ reproduces at every date and in every state the action distribution induced by τ. Since payoffs are separable across dates, the two strategies generate the same expected payoff, delivering (2).

(3) $\iff$ (4). Apply Theorem 0 to f^t and g^t at every date t.

A.4 Proof of Theorem 4

(2) $\implies$ (1) and (3) $\implies$ (2). Immediate.

(3) $\iff$ (4). Greenshtein (1996).

(4) $\implies$ (2). It follows from (3) $\iff$ (4) and (3) $\implies$ (2).

(1) $\implies$ (3). Consider a sequential problem with choice sets $A_1, \ldots, A_n$, set of feasible sequences $\mathcal{A} \subseteq A^n$, and payoff function $u\colon A^n \times \Theta \to \mathbb{R}$. We construct a constrained $\bar{\delta}$-discounted problem that is strategically equivalent to the sequential problem.

We first assume that n is finite, with $n \leq T$. Let $\widehat{A} = \bigcup_{t=1}^n A^t$ a finite set. We define $\widehat{\mathcal{A}}$ as follows: $\widehat{a}^T \in \widehat{\mathcal{A}}$ if and only if there exists $a^n \in \mathcal{A}$ such that $\widehat{a}^t = a^{\min(t,n)}$ for all $t \leq T$. In words, each decision completes the previous one until date n, after which the decision-maker must repeat the decision made at date n. Finally, $\widehat{u}(\widehat{a}, \theta) = 0$ if $\widehat{a} \in A^t$ with $t < n$, and $\widehat{u}(\widehat{a}, \theta) = u(a^n, \theta) / \sum_{s=n}^T \bar{\delta}_s$ if $\widehat{a} \in A^n$. The denominator is strictly positive because $\bar{\delta}$ has full support. Thus, the discounted payoff from the profile $\widehat{a}^T \in \widehat{\mathcal{A}}$ is

$$\sum_{t=n}^T \bar{\delta}_t \frac{u(a^n, \theta)}{\sum_{s=n}^T \bar{\delta}_s} = u(a^n, \theta),$$

with a^n the element in $\mathcal{A}$ corresponding to $\widehat{a}^T$. At each period $t \leq n$, choosing a^t in the constrained discounted problem is equivalent to choosing a_t in the original sequential problem, with the same information. Hence, strategies in the two problems correspond and have the same expected payoff in every state.

By (1), f is therefore more valuable than g in every n-period sequential problem. If $T < +\infty$, taking $n = T$ proves (3).

Suppose now that $T = +\infty$. For each finite n, consider the set of infinite sequences of kernels $\gamma_t \colon X^t \times Y^{t-1} \to \Delta(Y_t)$ satisfying

$$g^t(y^t \mid \theta) = \sum_{x^t} f^t(x^t \mid \theta) \prod_{s=1}^{t} \gamma_s(y_s \mid x^s, y^{s-1}) \tag{A.2}$$

for every $t \leq n$, y^t, and θ. The finite-horizon characterization of Greenshtein (1996) implies that each such set is non-empty: the kernels through date n can be extended arbitrarily to later dates. These sets are nested and closed in the compact space

$$\prod_{t \geq 1} \prod_{x^t \in X^t, y^{t-1} \in Y^{t-1}} \Delta(Y_t),$$

since each side of (A.2) is a continuous function of finitely many probabilities $\gamma_s(y_s \mid x^s, y^{s-1})$. Consequently, their intersection is non-empty. Any sequence of kernels in this intersection satisfies (A.2) at every date, establishing (4), hence proving (3).

Remark: In the simpler case, where $\mathcal{A} = \prod_{t=1}^{T} A_t$, with $A_t \subseteq A$ for all t, Δ-sufficiency allows the decision-maker to reproduce the distribution over decisions at each period, respecting the choice sets A_t and preserving the expected discounted payoff. Conversely, fix a period t. Allow an arbitrary finite set of decisions and arbitrary payoffs at that period, and require the decision-maker to choose the same additional decision, with payoff zero, at every other period. Since $\bar{\delta}_t > 0$, comparison in these problems implies that f^t is more valuable than g^t in all static problems. By Theorem 0, f^t is sufficient for g^t. Since t was arbitrary, f is Δ-sufficient for g, which by Theorem 3 is equivalent to f being more valuable than g in all separable problems.

A.5 Proof of Theorem 5

Proposition 7 *Suppose $T < +\infty$. For every decision problem $\widehat{u} \colon \widehat{A}_1 \times \cdots \times \widehat{A}_T \times \Theta \to \mathbb{R}$ with $\widehat{u} \geq 0$, there exist a real number $M(\widehat{u})$ and a history-dependent discounted problem $\phi(\widehat{u}) = (A, u, \delta)$ of length $T+1$ such that, for every experiment f, the following statements hold:*

1. *For every strategy $\widehat{\sigma}$ adapted to f, there exists a strategy σ adapted to f^+ such that*

$$\mathbb{E}_{\theta, f^+, \sigma}\left[\sum_{t=1}^{T+1} \delta_t(\boldsymbol{a}^{t-1}) u(\boldsymbol{a}_t, \theta)\right] = \mathbb{E}_{\theta, f, \widehat{\sigma}}\left[\widehat{u}(\boldsymbol{a}_1, \ldots, \boldsymbol{a}_T, \theta)\right] + T \times M(\widehat{u}), \qquad \forall\, \theta.$$

2. *For every strategy σ adapted to f^+, there exists a strategy $\widehat{\sigma}$ adapted to f such that*

$$\mathbb{E}_{\theta,f^+,\sigma}\left[\sum_{t=1}^{T+1}\delta_t(\boldsymbol{a}^{t-1})u(\boldsymbol{a}_t,\theta)\right]\leq\mathbb{E}_{\theta,f,\widehat{\sigma}}\left[\widehat{u}(\boldsymbol{a}_1,\ldots,\boldsymbol{a}_T,\theta)\right]+T\times M(\widehat{u}),\qquad\forall\theta.$$

Proof of Proposition 7 The proof is constructive. Fix the decision problem $\widehat{u}\colon \widehat{A}_1\times\cdots\times\widehat{A}_T\times\Theta\to\mathbb{R}$, with $\widehat{u}\geq 0$. To ease notation, we do not make explicit the dependence on $\widehat{u}$ in what follows. Let $M>2\|\widehat{u}\|_\infty$. We define the history-dependent discounted problem ϕ as follows:

- The set of decisions is $A=\bigcup_{t=1}^{T}(\widehat{A}_1\times\cdots\times\widehat{A}_t)\cup(\widehat{A}_1\times\cdots\times\widehat{A}_T\times\{a_\infty\})$, where a_∞ is a fixed (arbitrary) decision.
- The payoff function $u:A\times\Theta\to\mathbb{R}$ is

$$u(a,\theta)=\begin{cases}M, & a\notin\widehat{A}_1\times\cdots\times\widehat{A}_T,\\ \widehat{u}(a,\theta), & a\in\widehat{A}_1\times\cdots\times\widehat{A}_T.\end{cases}$$

- For every $t\in\{1,\ldots,T+1\}$, the discount factor $\delta_t\colon A^{t-1}\to[0,1]$ is

$$\delta_t(a^{t-1})=\begin{cases}0, & \text{if } a^{t-1}\text{ is incoherent},\\ 1, & \text{if } a^{t-1}\text{ is coherent},\end{cases}$$

  where $a^{t-1}=(a_1,\ldots,a_{t-1})$ is *coherent* if $a_s\in\widehat{A}_1\times\cdots\times\widehat{A}_s$ for every $1\leq s<t$, and the restriction of a_s to its first $s-1$ coordinates is a_{s-1} for every $2\leq s<t$. In words, each decision completes the preceding one. For every $t>T+1$, let $\delta_t:=0$.

Fix an experiment f and a strategy $\widehat{\sigma}$ adapted to f. We first construct a strategy σ adapted to f^+. At a coherent history $a^{t-1}=(\widehat{a}^1,\ldots,\widehat{a}^{t-1})$ ($t\leq T$), let

$$\sigma_t\left(a_t\mid x^t,(\widehat{a}^1,\ldots,\widehat{a}^{t-1})\right)=\begin{cases}\widehat{\sigma}_t(\widehat{a}_t\mid x^t,\widehat{a}^{t-1}), & \text{if } a_t=(\widehat{a}^{t-1},\widehat{a}_t)\text{ for some }\widehat{a}_t\in\widehat{A}_t,\\ 0, & \text{otherwise}.\end{cases}$$

At date $T+1$, after a coherent history corresponding to $\widehat{a}^T$, let σ choose $(\widehat{a}^T,a_\infty)$. At all incoherent histories, define the strategy arbitrarily. By construction,

$$\mathbb{E}_{\theta,f^+,\sigma}\left[\sum_{t=1}^{T+1}\delta_t(\boldsymbol{a}^{t-1})u(\boldsymbol{a}_t,\theta)\right]=(T-1)\times M+\mathbb{E}_{\theta,f,\widehat{\sigma}}\left[\widehat{u}(\boldsymbol{a}_1,\ldots,\boldsymbol{a}_T,\theta)\right]+M$$

for every θ. This proves the first statement.

We now prove the second statement. Fix a strategy σ adapted to f^+. We construct another strategy σ^* adapted to f^+ by backward induction. At date $T+1$, after every coherent history $a^T = (a_1, \ldots, a_T)$ with $a_T = \widehat{a}^T$, let $\sigma^*_{T+1}((\widehat{a}^T, a_\infty)|a^T, x^{T+1}) = 1$ for all x^{T+1}. At incoherent histories, let σ^*_{T+1} agree with σ_{T+1}.

At the coherent history (a^{t-1}, x^t) for $1 \leq t \leq T$ (where $a^{t-1} = (\widehat{a}^1, \ldots, \widehat{a}^{t-1})$), fix an arbitrary $\widehat{a}^*_t \in \widehat{A}_t$, let $a^*_t := (\widehat{a}^{t-1}, \widehat{a}^*_t)$ and define

$$\sigma^*_t(a^*_t|a^{t-1}, x^t) = \sigma_t(a^*_t|a^{t-1}, x^t) + \sum_{a_t \in A:(a^{t-1},a_t) \text{ incoherent}} \sigma_t(a_t|a^{t-1}, x^t),$$

$$\sigma^*_t(a_t|a^{t-1}, x^t) = \sigma_t(a_t|a^{t-1}, x^t), \forall a_t \neq a^*_t \text{ such that } (a^{t-1}, a_t) \text{ is coherent,}$$

so that $\sigma^*_t(a_t|a^{t-1}, x^t) = 0$ whenever (a^{t-1}, a_t) is incoherent. At incoherent histories, let $\sigma^*_t = \sigma_t$.

Under σ, any path that first chooses an incoherent action a_t at period $t \leq T$ receives $u(a_t, \theta) \leq \max(M, \|\widehat{u}\|_\infty) = M$ at date t and 0 at all later dates (since $\delta_s(a^{s-1}) = 0$ for all $s > t$), whereas σ^* induces a coherent history, hence a continuation payoff of $(T-t)M + \widehat{u}(\widehat{a}^T, \theta) + M \geq M \geq u(a_t, \theta)$. Since σ^* ensures coherent histories with probability one, the strategy $\widehat{\sigma}$ defined by

$$\widehat{\sigma}_t(\widehat{a}_t \mid x^t, \widehat{a}^{t-1}) \coloneqq \sigma^*_t\left((\widehat{a}^{t-1}, \widehat{a}_t) \mid x^t, (\widehat{a}^1, \ldots, \widehat{a}^{t-1})\right)$$

is well-defined (and adapted to f). For every θ, we have

$$\begin{aligned}\mathbb{E}_{\theta, f^+, \sigma}\left[\sum_{t=1}^{T+1} \delta_t(\boldsymbol{a}^{t-1}) u(\boldsymbol{a}_t, \theta)\right] &\leq \mathbb{E}_{\theta, f^+, \sigma^*}\left[\sum_{t=1}^{T+1} \delta_t(\boldsymbol{a}^{t-1}) u(\boldsymbol{a}_t, \theta)\right] \\ &= (T-1) \times M + \mathbb{E}_{\theta, f, \widehat{\sigma}}\left[\widehat{u}(\boldsymbol{a}_1, \ldots, \boldsymbol{a}_T, \theta)\right] + M,\end{aligned}$$

proving the second statement. ■

COMMENT: At period T the decision-maker can use all available information to choose the best possible joint decision. In general, this gives a strictly higher payoff than choosing sequentially. The $T+1$-th period is then crucial in insuring that if the decision-maker does so, he is punished in the continuation by not getting the bonus M.

Proof of Theorem 5 (1) $\implies$ (2) and (3) $\implies$ (1). Immediate.

(2) $\implies$ (3). Suppose first that $T < +\infty$. Fix a sequential decision problem $\widehat{u}\colon \widehat{A}_1 \times \cdots \times \widehat{A}_T \times \Theta \to \mathbb{R}$ and a strategy $\widehat{\tau}$ adapted to g. Without loss of generality, assume that $\widehat{u} \geq 0$. Let $M = M(\widehat{u})$ and $\phi(\widehat{u}) = (A, u, \delta)$ be given by Proposition 7.

The first statement of Proposition 7, applied to g, delivers a strategy τ adapted to g^+. By (2), there exists a strategy σ adapted to f^+ whose expected payoff is at least that of τ in every state. The second statement of Proposition 7, applied to f, then delivers a strategy $\widehat{\sigma}$ adapted to f such that, for every θ,

$$\begin{aligned}\mathbb{E}_{\theta,f,\widehat{\sigma}}\left[\widehat{u}(\boldsymbol{a}_1,\ldots,\boldsymbol{a}_T,\theta)\right] + T \times M &\geq \mathbb{E}_{\theta,f^+,\sigma}\left[\sum_{t=1}^{T+1}\delta_t(\boldsymbol{a}^{t-1})u(\boldsymbol{a}_t,\theta)\right]\\ &\geq \mathbb{E}_{\theta,g^+,\tau}\left[\sum_{t=1}^{T+1}\delta_t(\boldsymbol{a}^{t-1})u(\boldsymbol{a}_t,\theta)\right]\\ &= \mathbb{E}_{\theta,g,\widehat{\tau}}\left[\widehat{u}(\boldsymbol{a}_1,\ldots,\boldsymbol{a}_T,\theta)\right] + T \times M.\end{aligned}$$

Canceling $T \times M$, f is more valuable than g in every sequential problem. The characterization of Greenshtein (1996) then implies that f is sufficient for g.

Suppose now that $T = +\infty$. The above arguments imply that f is more valuable than g in all sequential problems of length $n < \infty$. By the finite-horizon characterization of Greenshtein (1996), there exists an adapted garbling linking g^n and f^n. Since n was arbitrary, the compactness argument in the proof of Theorem 4 implies the sufficiency of f for g. ■

A.6 Proof of Theorem 6

(2) $\implies$ (1). Immediate.

(1) $\implies$ (2). As in the proof of Theorem 1 and 2.

(3) $\implies$ (2). Consider the decision problem $(A, u, \kappa_t)_t$. Fix any strategy τ. Observe that

$$\begin{aligned}&\mathbb{P}_{\theta,g,\kappa,\tau}(a_t|y^t,k^{t-1}) =\\ &\frac{\sum_{a^{t-1}} g_1(y_1|\theta)\tau_1(a_1|y_1)\kappa_1(k_1|a_1,k^0)\times\cdots\times g_t(y_t|y^{t-1},k^{t-1},\theta)\tau_t(a_t|a^{t-1},y^t,k^{t-1})}{\sum_{a^t} g_1(y_1|\theta)\tau_1(a_1|y_1)\kappa_1(k_1|a_1,k^0)\times\cdots\times g_t(y_t|y^{t-1},k^{t-1},\theta)\tau_t(a_t|a^{t-1},y^t,k^{t-1})} =\\ &\frac{\sum_{a^{t-1}} \tau_1(a_1|y_1)\kappa_1(k_1|a_1,k^0)\times\cdots\times \tau_t(a_t|a^{t-1},y^t,k^{t-1})}{\sum_{a^t} \tau_1(a_1|y_1)\kappa_1(k_1|a_1,k^0)\times\cdots\times \tau_t(a_t|a^{t-1},y^t,k^{t-1})},\end{aligned}$$

whenever the denominator is strictly positive, i.e., $\mathbb{P}_{\theta,g,\kappa,\tau}(y^t, k^{t-1}) > 0$. The last ratio does not depend on θ. Thus, if (y^t, k^{t-1}) has positive probability in at least one state, the conditional distribution of $\boldsymbol{a}_t$ given (y^t, k^{t-1}) is the same in every state in which this history has positive probability. We denote this common distribution by

$$\overline{\mathbb{P}}_{\theta,g,\kappa,\tau}(\boldsymbol{a}_t \mid y^t, k^{t-1}).$$

and use it also in states in which the history has zero probability. If the history (y^t, k^{t-1}) has zero probability in every state θ, we let $\overline{\mathbb{P}}_{\theta,g,\kappa,\tau}(\boldsymbol{a}_t \mid y^t, k^{t-1})$ be an arbitrary distribution over A, the same for every θ.

Define $\xi_{t'} \colon \bigcup_t (Y^t \times K^{t-1}) \times K^{t'-1} \to \Delta(K_{t'})$ by

$$\xi_{t'}(k_{t'} \mid (y^t, l^{t-1}), k^{t'-1}) \coloneqq \sum_a \overline{\mathbb{P}}_{\theta,g,\kappa,\tau}(\boldsymbol{a}_t = a \mid y^t, l^{t-1}) \kappa_{t'}(k_{t'} \mid a, k^{t'-1}).$$

By construction, $\xi_{t'}$ is independent of θ.

Let $(\gamma_t)_t$ be the kernels satisfying δ-sufficiency, which correspond to the kernels $(\xi_t)_t$ we have constructed. Define $\sigma_{t'} : X^{t'} \times K^{t'-1} \to \Delta(A)$ as follows:

$$\sigma_{t'}(a|x^{t'}, k^{t'-1}) = \sum_t \sum_{y^t, l^{t-1}} \overline{\mathbb{P}}_{\theta,g,\kappa,\tau}(\boldsymbol{a_t} = a|y^t, l^{t-1}) \gamma_{t'}(y^t, l^{t-1}|x^{t'}, k^{t'-1}).$$

For future reference, notice that:

$$\begin{aligned}\sum_a \sigma_{t'}(a|x^{t'}, k^{t'-1}) \kappa_{t'}(k_{t'}|a, k^{t'-1}) &= \sum_a \sum_{t, y^t, l^{t-1}} \overline{\mathbb{P}}_{\theta,g,\kappa,\tau}(\boldsymbol{a_t} = a|y^t, l^{t-1}) \gamma_{t'}(y^t, l^{t-1}|x^{t'}, k^{t'-1}) \kappa_{t'}(k_{t'}|a, k^{t'-1}), \\ &= \sum_{t, y^t, l^{t-1}} \xi_{t'}(k_{t'}|(y^t, l^{t-1}), k^{t'-1}) \gamma_{t'}(y^t, l^{t-1}|x^{t'}, k^{t'-1}).\end{aligned}$$

We now show that $\sum_t \delta_t \text{marg}_{A_t} \mathbb{P}_{\theta,g,\kappa,\tau} = \sum_t \delta_t \text{marg}_{A_t} \mathbb{P}_{\theta,f,\kappa,\sigma}$. We have:

$$\sum_{t',x^{t'},k^{t'-1}} \delta_{t'} \mathbb{P}_{\theta,f,\kappa,\sigma}(x^{t'}, k^{t'-1}) \mathbb{P}_{\theta,f,\kappa,\sigma}(\boldsymbol{a_{t'}} = a | x^{t'}, k^{t'-1}) =$$

$$\sum_{t',x^{t'},k^{t'-1}} \delta_{t'} \mathbb{P}_{\theta,f,\kappa,\sigma}(x^{t'}, k^{t'-1}) \sigma_{t'}(a | x^{t'}, k^{t'-1}) =$$

$$\sum_{t',x^{t'},k^{t'-1}} \delta_{t'} \mathbb{P}_{\theta,f,\kappa,\sigma}(x^{t'}, k^{t'-1}) \left(\sum_{t,y^t,l^{t-1}} \overline{\mathbb{P}}_{\theta,g,\kappa,\tau}(\boldsymbol{a_t} = a | y^t, l^{t-1}) \gamma_{t'}(y^t, l^{t-1} | x^{t'}, k^{t'-1}) \right) =$$

$$\sum_{t,y^t,l^{t-1}} \overline{\mathbb{P}}_{\theta,g,\kappa,\tau}(\boldsymbol{a_t} = a | y^t, l^{t-1}) \left(\sum_{t',x^{t'},k^{t'-1}} \delta_{t'} \mathbb{P}_{\theta,f,\kappa,\sigma}(x^{t'}, k^{t'-1}) \gamma_{t'}(y^t, l^{t-1} | x^{t'}, k^{t'-1}) \right).$$

We argue that δ-sufficiency implies that the term in parentheses equals $\delta_t \mathbb{P}_{\theta,g,\kappa,\tau}(y^t, l^{t-1})$. By construction, we have:

$$\mathbb{P}_{\theta,f,\kappa,\sigma}(x^{t'}, k^{t'-1}) = f_1(x_1|\theta) \times \cdots \times f_{t'}(x_{t'} | x^{t'-1}, k^{t'-1}, \theta) \times$$

$$\left(\sum_{a_1} \sigma_1(a_1|x^1) \kappa_1(k_1|a_1) \right) \times \cdots \times \left(\sum_{a_{t'-1}} \sigma_{t'-1}(a_{t'-1} | x^{t'-1}, k^{t'-2}) \kappa_{t'-1}(k_{t'-1} | a_{t'-1}, k^{t'-2}) \right) =$$

$$f_1(x_1|\theta) \times \cdots \times f_{t'}(x_{t'} | x^{t'-1}, k^{t'-1}, \theta) \times \left(\sum_{t,y^t,l^{t-1}} \xi_1(k_1 | (y^t, l^{t-1}), k^0) \gamma_1(y^t, l^{t-1} | x^1, k^0) \right) \times \cdots \times$$

$$\left(\sum_{t,y^t,l^{t-1}} \xi_{t'-1}(k_{t'-1} | (y^t, l^{t-1}), k^{t'-2}) \gamma_{t'-1}(y^t, l^{t-1} | x^{t'-1}, k^{t'-2}) \right) = \mathbb{P}_{\theta,f,\,\xi \circ \gamma}(x^{t'}, k^{t'-1}).$$

The condition of δ-sufficiency then implies that:

$$\sum_{t',x^{t'},k^{t'-1}} \delta_{t'} \mathbb{P}_{\theta,f,\kappa,\sigma}(x^{t'}, k^{t'-1}) \gamma_{t'}(y^t, l^{t-1} | x^{t'}, k^{t'-1}) = \delta_t \mathbb{P}_{\theta,g,\xi}(y^t, l^{t-1})$$

which is equal to $\delta_t \mathbb{P}_{\theta,g,\kappa,\tau}(y^t, l^{t-1})$ whenever $\mathbb{P}_{\theta,g,\kappa,\tau}(y^t, l^{t-1}) > 0$. To see this, note that

$$\mathbb{P}_{\theta,g,\kappa,\tau}(y^t, l^{t-1}) = \sum_{a^{t-1}} \mathbb{P}_{\theta,g,\kappa,\tau}(a^{t-1}, y^t, l^{t-1})$$

$$= \sum_{a^{t-1}} \mathbb{P}_{\theta,g,\kappa,\tau}(y_t|a^{t-1}, y^{t-1}, l^{t-1}) \mathbb{P}_{\theta,g,\kappa,\tau}(l_{t-1}|a^{t-1}, y^{t-1}, l^{t-2}) \mathbb{P}_{\theta,g,\kappa,\tau}(a^{t-1}|y^{t-1}, l^{t-2}) \mathbb{P}_{\theta,g,\kappa,\tau}(y^{t-1}, l^{t-2})$$

$$= \sum_{a^{t-1}} g_t(y_t|y^{t-1}, l^{t-1}, \theta) \kappa_{t-1}(l_{t-1}|a_{t-1}, l^{t-2}) \mathbb{P}_{\theta,g,\kappa,\tau}(a^{t-1}|y^{t-1}, l^{t-2}) \mathbb{P}_{\theta,g,\kappa,\tau}(y^{t-1}, l^{t-2})$$

$$= g_t(y_t|y^{t-1}, l^{t-1}, \theta) \left(\sum_{a_{t-1}} \kappa_{t-1}(l_{t-1}|a_{t-1}, l^{t-2}) \mathbb{P}_{\theta,g,\kappa,\tau}(a_{t-1}|y^{t-1}, l^{t-2}) \right) \mathbb{P}_{\theta,g,\kappa,\tau}(y^{t-1}, l^{t-2})$$

$$= g_t(y_t|y^{t-1}, l^{t-1}, \theta) \xi_{t-1}(l_{t-1}|(y^{t-1}, l^{t-2}), l^{t-2}) \mathbb{P}_{\theta,g,\kappa,\tau}(y^{t-1}, l^{t-2}),$$

so that the result follows by induction.

$(2) \implies (3)$. Fix $\xi \in \Xi$ and a finite $n \leq T$. Let

$$A := \begin{cases} \bigcup_{t=1}^{n} (Y^t \times K^{t-1}) \cup \{(y^{n+1}, k^n)\} & \text{if } n < T, \\ \bigcup_{t=1}^{n} (Y^t \times K^{t-1}) & \text{if } n = T, \end{cases}$$

for some fixed (y^{n+1}, k^n). For every date $s \leq T$, define $\kappa_s(k_s \mid (y^t, l^{t-1}), k^{s-1}) := \xi_s(k_s \mid (y^t, l^{t-1}), k^{s-1})$ for every $(y^t, l^{t-1}) \in A$. At all $t \leq n$, let $\tau_t((y^t, k^{t-1}) \mid a^{t-1}, y^t, k^{t-1}) = 1$ for all (a^{t-1}, y^t, k^{t-1}). At $t > n$, let $\tau_t((y^{n+1}, k^n)|a^{t-1}, y^t, k^{t-1}) = 1$ for all (a^{t-1}, y^t, k^{t-1}).

We show that $\mathbb{P}_{\theta,g,\kappa,\tau}$ and $\mathbb{P}_{\theta,g,\xi}$ assign the same probability to every history (y^t, l^{t-1}) with $t \leq n$. Both processes have first-period signal distribution $g_1(\cdot \mid \theta)$. Suppose they assign the same probability to every history (y^s, l^{s-1}). After such a history, the strategy τ chooses the decision (y^s, l^{s-1}), so the definition of κ_s implies that both processes draw l_s according to $\xi_s(\cdot \mid (y^s, l^{s-1}), l^{s-1})$.

Both processes then draw y_{s+1} according to $g_{s+1}(\cdot \mid y^s, l^s, \theta)$. Accordingly, they assign the same probability to every history (y^{s+1}, l^s). Repeating this argument from the first period onward, we obtain $\mathbb{P}_{\theta,g,\kappa,\tau}(y^t, l^{t-1}) = \mathbb{P}_{\theta,g,\xi}(y^t, l^{t-1})$ for every $t \leq n$, θ, and (y^t, l^{t-1}). Consequently, the discounted probability of the decision (y^t, l^{t-1}) under g is $\delta_t \mathbb{P}_{\theta,g,\xi}(y^t, l^{t-1})$.

By (2), there exists a strategy σ under f such that

$$\delta_t \mathbb{P}_{\theta,g,\xi}(y^t, l^{t-1}) = \sum_{s=1}^{T} \delta_s \mathbb{P}_{\theta,f,\kappa,\sigma} \left(\boldsymbol{a}_s = (y^t, l^{t-1}) \right)$$

for every $t \leq n$, θ, and (y^t, l^{t-1}).

At a history (x^s, k^{s-1}) that has positive probability in at least one state, use the common conditional distribution of $\boldsymbol{a}_s$ given this history. At a history that has zero probability in every state, choose any distribution over A. Define

$$\gamma_s(y^t, l^{t-1} \mid x^s, k^{s-1}) := \overline{\mathbb{P}}_{\theta,f,\kappa,\sigma}\left(\boldsymbol{a}_s = (y^t, l^{t-1}) \mid x^s, k^{s-1}\right).$$

By construction, γ_s does not depend on θ. (See the proof of $(3) \implies (2)$ for details.)

We next show that $\mathbb{P}_{\theta,f,\kappa,\sigma}$ and $\mathbb{P}_{\theta,f,\xi\circ\gamma}$ assign the same probability to every history (x^s, k^{s-1}), $s \leq T$. Conditional on such a history, the distribution of k_s under (f, κ, σ) is

$$\sum_{t,y^t,l^{t-1}} \gamma_s(y^t, l^{t-1} \mid x^s, k^{s-1})\xi_s(k_s \mid (y^t, l^{t-1}), k^{s-1}) = \sum_{a\in A} \overline{\mathbb{P}}_{\theta,f,\kappa,\sigma}(\boldsymbol{a}_s = a \mid x^s, k^{s-1})\kappa_s(k_s \mid a, k^{s-1}),$$

and is, therefore, $(\xi_s \circ \gamma_s)(k_s \mid x^s, k^{s-1})$. Both processes have first-period signal distribution $f_1(\cdot \mid \theta)$. Suppose they assign the same probability to every history (x^s, k^{s-1}). After such a history, they draw k_s according to the same distribution and then draw x_{s+1} according to $f_{s+1}(\cdot \mid x^s, k^s, \theta)$. Consequently, they assign the same probability to every history (x^{s+1}, k^s). Repeating this argument from the first period onward, we obtain $\mathbb{P}_{\theta,f,\kappa,\sigma}(x^s, k^{s-1}) = \mathbb{P}_{\theta,f,\xi\circ\gamma}(x^s, k^{s-1})$ for every $s \leq T$, θ, and (x^s, k^{s-1}) and, hence,

$$\begin{aligned}\delta_t\mathbb{P}_{\theta,g,\xi}(y^t, l^{t-1}) &= \sum_{s=1}^{T} \delta_s \sum_{x^s,k^{s-1}} \mathbb{P}_{\theta,f,\kappa,\sigma}(x^s, k^{s-1})\gamma_s(y^t, l^{t-1} \mid x^s, k^{s-1}) \\ &= \sum_{s=1}^{T} \delta_s \sum_{x^s,k^{s-1}} \mathbb{P}_{\theta,f,\xi\circ\gamma}(x^s, k^{s-1})\gamma_s(y^t, l^{t-1} \mid x^s, k^{s-1})\end{aligned}$$

for all $t \leq n$, for all (y^t, l^{t-1}), for all θ. Therefore, if $T = n$, we have proved (3).

Assume now that $T = +\infty$. If δ has a bounded support, the result follows from the above argument. So, suppose that δ has an unbounded support.

From the above arguments, for each $n < \infty$, there exists γ^n such that

$$\delta_t\mathbb{P}_{\theta,g,\xi}(y^t, l^{t-1}) = \sum_{s=1}^{+\infty} \delta_s \sum_{x^s,k^{s-1}} \mathbb{P}_{\theta,f,\xi\circ\gamma^n}(x^s, k^{s-1})\gamma_s^n(y^t, l^{t-1} \mid x^s, k^{s-1}),$$

for all (θ, y^t, l^{t-1}), for all $t \leq n$. (We view each γ_s^n as a mapping from $X^s \times K^{s-1}$ to $\Delta\left(\bigcup_{t=1}^{\infty}(Y^t \times K^{t-1})\right)$.)

By diagonal extraction, we can choose a subsequence, still indexed by n, such that each $\gamma_s^n(y^t, l^{t-1}|x^s, k^{s-1})$ and each $(\xi_s \circ \gamma_s^n)(k_s \mid x^s, k^{s-1})$ converges. This is possible because there are countably many coordinates. Denote by $\gamma_s(y^t, l^{t-1}|x^s, k^{s-1})$ and $\overline{\xi}_s^x(k_s \mid x^s, k^{s-1})$ the corresponding limits, and write $\overline{\xi}^x = (\overline{\xi}_s^x)_{s\geq 1}$. Since K_s is finite, $\overline{\xi}_s^x(\cdot \mid x^s, k^{s-1})$ is a probability distribution. From Fatou's lemma, we have

$$\begin{aligned}
&\sum_{t\geq 1}\sum_{y^t,l^{t-1}} \gamma_s(y^t, l^{t-1} \mid x^s, k^{s-1}) \\
&\quad = \sum_{t\geq 1}\sum_{y^t,l^{t-1}} \lim_{n\to\infty} \gamma_s^n(y^t, l^{t-1} \mid x^s, k^{s-1}) \\
&\quad \leq \liminf_{n\to\infty} \sum_{t\geq 1}\sum_{y^t,l^{t-1}} \gamma_s^n(y^t, l^{t-1} \mid x^s, k^{s-1}) = 1.
\end{aligned}$$

Thus, $\gamma_s(\cdot|x^s, k^{s-1})$ is a sub-probability. We argue later on that it is a probability at all relevant histories (x^s, k^{s-1}).

Similarly, for each s, (x^s, k^{s-1}), and k_s, Fatou's lemma also gives

$$\begin{aligned}
(\xi_s \circ \gamma_s)(k_s \mid x^s, k^{s-1}) &= \sum_{t\geq 1}\sum_{y^t,l^{t-1}} \gamma_s(y^t, l^{t-1} \mid x^s, k^{s-1})\xi_s(k_s \mid (y^t, l^{t-1}), k^{s-1}) \\
&\leq \lim_{n\to\infty} (\xi_s \circ \gamma_s^n)(k_s \mid x^s, k^{s-1}) = \overline{\xi}_s^x(k_s \mid x^s, k^{s-1}).
\end{aligned}$$

Starting from the first-period distribution $f_1(\cdot \mid \theta)$, induction over dates therefore yields

$$\mathbb{P}_{\theta,f,\xi\circ\gamma^n}(x^s, k^{s-1}) \longrightarrow \mathbb{P}_{\theta,f,\overline{\xi}^x}(x^s, k^{s-1})$$

for every s, θ, and (x^s, k^{s-1}). The same subsequence works for every θ, since the converging kernels are state-independent.

Fix t, θ, and (y^t, l^{t-1}). Let $S < \infty$. For each $n \geq t$, we have:

$$0 \leq \delta_t \mathbb{P}_{\theta,g,\xi}(y^t, l^{t-1}) - \sum_{s=1}^{S} \delta_s \sum_{x^s,k^{s-1}} \mathbb{P}_{\theta,f,\xi\circ\gamma^n}(x^s, k^{s-1})\gamma_s^n(y^t, l^{t-1} \mid x^s, k^{s-1}) \leq \sum_{s>S} \delta_s. \qquad (*)$$

Indeed,

$$\begin{aligned}\delta_t \mathbb{P}_{\theta,g,\xi}(y^t, l^{t-1}) &= \sum_{s=1}^{\infty} \delta_s \sum_{x^s,k^{s-1}} \mathbb{P}_{\theta,f,\xi\circ\gamma^n}(x^s,k^{s-1})\gamma_s^n(y^t,l^{t-1} \mid x^s,k^{s-1}) \\ &= \sum_{s=1}^{S} \delta_s \sum_{x^s,k^{s-1}} \mathbb{P}_{\theta,f,\xi\circ\gamma^n}(x^s,k^{s-1})\gamma_s^n(y^t,l^{t-1} \mid x^s,k^{s-1}) \\ &\quad + \sum_{s>S} \delta_s \underbrace{\sum_{x^s,k^{s-1}} \mathbb{P}_{\theta,f,\xi\circ\gamma^n}(x^s,k^{s-1})\gamma_s^n(y^t,l^{t-1} \mid x^s,k^{s-1})}_{\leq 1}.\end{aligned}$$

Taking limits first as $n \to +\infty$ and then $S \to +\infty$ gives:

$$\delta_t \mathbb{P}_{\theta,g,\xi}(y^t, l^{t-1}) = \sum_{s=1}^{\infty} \delta_s \sum_{x^s,k^{s-1}} \mathbb{P}_{\theta,f,\overline{\xi}^x}(x^s,k^{s-1})\gamma_s(y^t,l^{t-1} \mid x^s,k^{s-1}).$$

Thus all the required identities hold. Finally, we verify that the kernels are indeed probabilities supported on $\bigcup_{t=1}^{\infty}(Y^t \times K^{t-1})$, and therefore that $\overline{\xi}^x$ coincides with $\xi \circ \gamma$ at all relevant histories.

Summing the above equalities over (t, y^t, l^{t-1}) implies that for every θ,

$$\sum_{s=1}^{\infty} \delta_s \sum_{x^s,k^{s-1}} \mathbb{P}_{\theta,f,\overline{\xi}^x}(x^s,k^{s-1})\gamma_s\left(\bigcup_{t=1}^{\infty}(Y^t \times K^{t-1}) \mid x^s,k^{s-1}\right) = 1.$$

Finally, since $\delta_s > 0$ for all s, $\gamma_s\left(\bigcup_{t=1}^{\infty}(Y^t \times K^{t-1}) \mid x^s,k^{s-1}\right) = 1$ for all (x^s,k^{s-1}) such that $\mathbb{P}_{\theta,f,\overline{\xi}^x}(x^s,k^{s-1}) > 0$ for some θ. At every such history, both $(\xi_s \circ \gamma_s)(\cdot \mid x^s,k^{s-1})$ and $\overline{\xi}_s^x(\cdot \mid x^s,k^{s-1})$ are probability distributions on K_s, and the former is pointwise weakly below the latter. Hence they coincide.

If $\mathbb{P}_{\theta,f,\overline{\xi}^x}(x^s,k^{s-1}) = 0$ for all θ, then we can choose an arbitrary probability $\gamma_s(\cdot|x^s,k^{s-1})$ without affecting the equalities. Therefore $\overline{\xi}_s^x = \xi_s \circ \gamma_s$ at all relevant histories, and $\mathbb{P}_{\theta,f,\overline{\xi}^x} = \mathbb{P}_{\theta,f,\xi\circ\gamma}$.

A.7 Proof of Proposition 3

Put

$$q := \min_{k\in K} \frac{1}{\max\{1, R_k^*\}}, \qquad C := \sum_{t=1}^{T} \delta_t q^{t-1},$$

with $1/(+\infty) = 0$ and $0^0 = 1$. Then $0 \leq q, C \leq 1$, and (4) is equivalent to $R_{k_0}^* \leq C$.

The static kernels. The definition of R_k^* gives, for every static decision problem,

$$V^*_{p,h_f(k_0,\cdot)}(u) \geq CV^*_{p,h_g(k_0,\cdot)}(u) + (1-C)V^*_{p,h^*}(u),$$
$$qV^*_{p,h_f(k,\cdot)}(u) + (1-q)V^*_{p,h^*}(u) \geq V^*_{p,h_g(k,\cdot)}(u) \qquad (k \in K).$$

For the second inequality, use $qR_k^* \leq 1$ when $q > 0$; when $q = 0$, full information is sufficient. Mixtures here are fictitious kernels on disjoint signal spaces. Their static values are the corresponding weighted averages. Blackwell's theorem therefore gives a stochastic kernel $\beta\colon X \to \Delta(Y \sqcup \Theta)$ satisfying

$$\sum_x h_f(x \mid k_0, \theta)\beta(y \mid x) = Ch_g(y \mid k_0, \theta), \tag{A.3}$$

$$\sum_x h_f(x \mid k_0, \theta)\beta(\widehat{\theta} \mid x) = (1-C)\mathbf{1}_{\{\widehat{\theta}=\theta\}}. \tag{A.4}$$

The entries of β are nonnegative, with $\sum_y \beta(y \mid x) + \sum_{\widehat{\theta}} \beta(\widehat{\theta} \mid x) = 1$ for every x. These equalities and (A.3)–(A.4) form a finite linear feasibility problem. The reports y and $\widehat{\theta}$ are generated from the actual signal x. In particular, (A.4) implies

$$\beta(\widehat{\theta} \mid x) > 0, \quad h_f(x \mid k_0, \theta) > 0 \quad \Longrightarrow \quad \widehat{\theta} = \theta.$$

Hence a state report uses information already contained in x. When $C < 1$, the hypothesis requires initial signals identifying the state with positive probability in every state.

Similarly, restricting a garbling from $qh_f(k,\cdot) + (1-q)h^*$ to $h_g(k,\cdot)$ to its X-rows gives stochastic kernels $\lambda_k\colon X \to \Delta(Y)$ such that

$$q\sum_x h_f(x \mid k, \theta)\lambda_k(y \mid x) \leq h_g(y \mid k, \theta). \tag{A.5}$$

Together with nonnegativity and $\sum_y \lambda_k(y \mid x) = 1$, (A.5) is another finite linear feasibility problem. If $q = 0$, any stochastic λ_k will do.

The discount weights. For $C > 0$, define

$$w_t := \frac{\sum_{s=t}^T \delta_s q^{s-1}}{C\sum_{s=t}^T \delta_s} \tag{A.6}$$

when the denominator is positive, and set $w_t = 0$ otherwise. If $C = 0$, set $w_1 = 1$ and

$w_t = 0$ for $t > 1$. For finite T, also set $w_{T+1} = 0$.

The construction uses $1 = w_1 \geq w_t \geq w_{t+1} \geq 0$. To verify this, suppose first that $C > 0$ and the following denominators are positive. Since $q^{s-1} \leq q^{t-1}$ for every $s > t$,

$$Cw_{t+1} = \frac{\sum_{s>t} \delta_s q^{s-1}}{\sum_{s>t} \delta_s} \leq q^{t-1},$$
$$C(w_t - w_{t+1}) = \frac{\delta_t \left(q^{t-1} - Cw_{t+1}\right)}{\sum_{s=t}^{T} \delta_s} \geq 0.$$

If there is no positive discount weight after t, then $w_{t+1} = 0$ and the conclusion follows directly. Also, $w_1 = 1$ by the definition of C; the case $C = 0$ follows from the specified weights. The equality

$$Cw_t \sum_{s=t}^{T} \delta_s = \sum_{s=t}^{T} \delta_s q^{s-1}$$

holds at both t and $t + 1$. Subtracting these two equalities yields

$$C\left[\delta_t w_t + \left(\sum_{s>t} \delta_s\right)(w_t - w_{t+1})\right] = \delta_t q^{t-1}. \tag{A.7}$$

For $C = 0$, every $\delta_t q^{t-1}$ is zero, so this identity still holds.

The behavioral strategy. Fix (A, u, κ) and any strategy τ under g. Subtracting the full-information payoff in each state leaves the comparison unchanged, since $\sum_t \delta_t = 1$. We may therefore assume $\max_a u(a, \theta) = 0$. For each report $\widehat{\theta}$, choose a maximizing action $a^*(\widehat{\theta})$.

At each actual history (a^{t-1}, k^{t-1}, x^t), we define nonnegative weights $M_t(\cdot \mid a^{t-1}, k^{t-1}, x^t)$ on the disjoint union $Y^t \sqcup A$. The y^t-component applies τ_t to that report history; the a-component chooses a at every remaining date. All history arguments are displayed below. At the first date, a^0 and k^0 contain no past actions or endogenous controls; the initial control k_0 is fixed. Let

$$\begin{aligned} M_1(y_1 \mid a^0, k^0, x_1) &= \beta(y_1 \mid x_1), \\ M_1(a \mid a^0, k^0, x_1) &= \sum_{\widehat{\theta}} \beta(\widehat{\theta} \mid x_1) \mathbf{1}_{\{a = a^*(\widehat{\theta})\}}. \end{aligned} \tag{A.8}$$

Set

$$D_t(a^{t-1}, k^{t-1}, x^t) := \sum_{y^t \in Y^t} M_t(y^t \mid a^{t-1}, k^{t-1}, x^t) + \sum_{a \in A} M_t(a \mid a^{t-1}, k^{t-1}, x^t).$$

If $D_t(a^{t-1}, k^{t-1}, x^t) > 0$, define the behavioral strategy by

$$\begin{aligned} &\sigma_t(a_t \mid a^{t-1}, k^{t-1}, x^t) \\ &\quad = \frac{M_t(a_t \mid a^{t-1}, k^{t-1}, x^t)}{D_t(a^{t-1}, k^{t-1}, x^t)} \\ &\qquad + \sum_{y^t} \frac{M_t(y^t \mid a^{t-1}, k^{t-1}, x^t)}{D_t(a^{t-1}, k^{t-1}, x^t)} \tau_t(a_t \mid a^{t-1}, y^t, k^{t-1}). \end{aligned} \tag{A.9}$$

If $D_t(a^{t-1}, k^{t-1}, x^t) = 0$, choose any fixed distribution on A. Nonnegativity and the definition of $D_t(a^{t-1}, k^{t-1}, x^t)$ imply that (A.9) sums to one over a_t. Here a_t labels a candidate action; it is not part of the conditioning history of either weight in (A.9).

For $t < T$ and each continuation (a_t, k_t, x_{t+1}), define the weights at (a^t, k^t, x^{t+1}) recursively by

$$\begin{aligned} &M_{t+1}(y^{t+1} \mid a^t, k^t, x^{t+1}) \\ &\quad = \frac{w_{t+1}}{w_t} M_t(y^t \mid a^{t-1}, k^{t-1}, x^t) \tau_t(a_t \mid a^{t-1}, y^t, k^{t-1}) \\ &\qquad \times \lambda_{k_t}(y_{t+1} \mid x_{t+1}), \end{aligned} \tag{A.10}$$

$$\begin{aligned} &M_{t+1}(a \mid a^t, k^t, x^{t+1}) \\ &\quad = \mathbf{1}_{\{a=a_t\}} M_t(a_t \mid a^{t-1}, k^{t-1}, x^t) \\ &\qquad + \mathbf{1}_{\{a=a_t\}} \left(1 - \frac{w_{t+1}}{w_t}\right) \sum_{y^t} M_t(y^t \mid a^{t-1}, k^{t-1}, x^t) \tau_t(a_t \mid a^{t-1}, y^t, k^{t-1}). \end{aligned} \tag{A.11}$$

Set $w_{t+1}/w_t = 0$ when $w_t = 0$. Equation (A.10) assigns the fraction w_{t+1}/w_t of each report contribution to generating a new report. Equation (A.11) assigns the remainder to repeating a_t; its indicator preserves each action already assigned to later dates. The new k_t, x_{t+1} enter (A.10) only through $\lambda_{k_t}(y_{t+1} \mid x_{t+1})$; they do not enter the right-hand side of (A.11). Both recursions use the old history and the indicated continuation. No formula uses the true state θ as an input. Hence (A.9) is adapted to f. The maps κ_t continue to generate actual controls at every date.

The induced distribution. Write $\mathbb{P}_{\theta,f,\kappa,\sigma}$ for the induced history law. Summing (A.10)–

(A.11) and using $\sum_{y_{t+1}} \lambda_{k_t}(y_{t+1} \mid x_{t+1}) = 1$ yields

$$\begin{aligned} &D_{t+1}(a^t, k^t, x^{t+1}) \\ &\quad = M_t(a_t \mid a^{t-1}, k^{t-1}, x^t) \\ &\qquad + \sum_{y^t} M_t(y^t \mid a^{t-1}, k^{t-1}, x^t) \tau_t(a_t \mid a^{t-1}, y^t, k^{t-1}) \\ &\quad = D_t(a^{t-1}, k^{t-1}, x^t)\, \sigma_t(a_t \mid a^{t-1}, k^{t-1}, x^t). \end{aligned} \tag{A.12}$$

Thus, holding a^t, k^{t-1}, x^t fixed, the total $D_{t+1}(a^t, k^t, x^{t+1})$ is independent of the new k_t, x_{t+1}. This equality concerns the sum of the weights; the individual report weights in (A.10) generally depend on both new coordinates. We claim that, for all histories and states,

$$\begin{aligned} &\mathbb{P}_{\theta,f,\kappa,\sigma}(a^{t-1}, k^{t-1}, x^t) \\ &\quad = D_t(a^{t-1}, k^{t-1}, x^t) \left[\prod_{s=1}^{t} h_f(x_s \mid k_{s-1}, \theta)\right] \left[\prod_{s=1}^{t-1} \kappa_s(k_s \mid a_s, k^{s-1})\right]. \end{aligned} \tag{A.13}$$

Empty products are one. At $t = 1$, the claim follows from $D_1(a^0, k^0, x_1) = 1$, the row sum of β. Suppose it holds at t. The probability of a continuation is the left-hand side of (A.13) multiplied by

$$\sigma_t(a_t \mid a^{t-1}, k^{t-1}, x^t)\, \kappa_t(k_t \mid a_t, k^{t-1})\, h_f(x_{t+1} \mid k_t, \theta).$$

Substituting (A.12) gives (A.13) at $t + 1$. If $D_t(a^{t-1}, k^{t-1}, x^t) = 0$, all the nonnegative weights $M_t(\cdot \mid a^{t-1}, k^{t-1}, x^t)$ vanish. The recursions then give $M_{t+1}(\cdot \mid a^t, k^t, x^{t+1}) = 0$ at every continuation. The history has zero probability by the induction hypothesis. Hence the same conclusion holds regardless of the arbitrary rule specified there. The control factors multiply every component $M_t(\cdot \mid a^{t-1}, k^{t-1}, x^t)$ by the same factor in (A.13). This is why they need not enter the weight recursions, although those weights depend on the realized controls.

The payoff comparison. Define the probability distribution

$$\begin{aligned} B_{\theta,t}(y^t, a^t, k^{t-1}) \coloneqq h_g(y_1 \mid k_0, \theta) &\prod_{s=1}^{t} \tau_s(a_s \mid a^{s-1}, y^s, k^{s-1}) \\ &\times \prod_{s=1}^{t-1} \left[\kappa_s(k_s \mid a_s, k^{s-1}) \sum_x h_f(x \mid k_s, \theta) \lambda_{k_s}(y_{s+1} \mid x)\right]. \end{aligned}$$

It uses $h_g(k_0, \cdot)$ initially, h_f garbled by λ_k subsequently, and the action and control maps τ, κ throughout. Multiplying (A.5) at each date after the first gives

$$q^{t-1} B_{\theta,t}(y^t, a^t, k^{t-1}) \leq \mathbb{P}_{\theta,g,\kappa,\tau}(y^t, a^t, k^{t-1}). \tag{A.14}$$

Iteration of (A.10), with (A.8), gives the following explicit expression for the report-history weights:

$$\begin{aligned} &M_t(y^t \mid a^{t-1}, k^{t-1}, x^t) \\ &\quad = w_t \beta(y_1 \mid x_1) \prod_{s=1}^{t-1} \left[\tau_s(a_s \mid a^{s-1}, y^s, k^{s-1}) \lambda_{k_s}(y_{s+1} \mid x_{s+1})\right]. \end{aligned}$$

Fix y^t, a^t, k^{t-1} and sum over x^t. At each history with positive $D_t(a^{t-1}, k^{t-1}, x^t)$, its factor in (A.13) cancels the denominator below. If that factor is zero, the corresponding $M_t(y^t \mid a^{t-1}, k^{t-1}, x^t)$ is also zero. Substituting the displayed product therefore gives

$$\begin{aligned} &\sum_{x^t} \frac{\mathbb{P}_{\theta,f,\kappa,\sigma}(a^{t-1}, k^{t-1}, x^t)}{D_t(a^{t-1}, k^{t-1}, x^t)} M_t(y^t \mid a^{t-1}, k^{t-1}, x^t) \tau_t(a_t \mid a^{t-1}, y^t, k^{t-1}) \\ &\quad = C w_t B_{\theta,t}(y^t, a^t, k^{t-1}), \end{aligned} \tag{A.15}$$

where a summand with $D_t(a^{t-1}, k^{t-1}, x^t) = 0$ is zero. Indeed, the τ and κ factors are fixed in this sum. The x_1-sum is (A.3), and each x_{s+1}-sum is $\sum_x h_f(x \mid k_s, \theta) \lambda_{k_s}(y_{s+1} \mid x)$, exactly the factors defining $B_{\theta,t}$. Thus the report-history term in σ_t contributes $C w_t \mathbb{E}_{B_{\theta,t}}[u(\mathbf{a}_t, \theta)]$ to date-t expected payoff.

By (A.11), the contribution newly assigned after a date $s < T$ to repeating a_s has distribution $C(w_s - w_{s+1}) B_{\theta,s}$ over (y^s, a^s, k^{s-1}). The recursion is linear in the weights, so each contribution can be followed separately. At every later date its action is a_s, by the indicator in (A.11).

At a transition from t to $t+1$, with $s \leq t < T$, its probability is multiplied only by

$$\kappa_t(k_t \mid a_s, k^{t-1}) h_f(x_{t+1} \mid k_t, \theta).$$

Summing these factors over k_t, x_{t+1} gives one. Repeating this argument shows that its total weight is preserved at every later date.

The initial action weights in (A.8) contribute zero payoff: by (A.4) they select $a^*(\theta)$ in

state θ. Consequently,

$$\begin{aligned}\mathbb{E}_{\theta,f,\kappa,\sigma}[u(\mathbf{a}_t,\theta)] = Cw_t\mathbb{E}_{B_{\theta,t}}[u(\mathbf{a}_t,\theta)] \\ + C\sum_{s<t}(w_s - w_{s+1})\mathbb{E}_{B_{\theta,s}}[u(\mathbf{a}_s,\theta)].\end{aligned} \tag{A.16}$$

An action selected from a report history therefore contributes its current payoff and, when assigned to all later dates, the same payoff with the remaining discount weights. Equation (A.7) makes their combined coefficient $\delta_s q^{s-1}$. Indeed,

$$\begin{aligned}&\mathbb{E}_{\theta,f,\kappa,\sigma}\left[\sum_t \delta_t u(\mathbf{a}_t,\theta)\right] \\ &= C\sum_s \left[\delta_s w_s + \left(\sum_{t>s}\delta_t\right)(w_s - w_{s+1})\right]\mathbb{E}_{B_{\theta,s}}[u(\mathbf{a}_s,\theta)] \\ &= \sum_s \delta_s q^{s-1}\mathbb{E}_{B_{\theta,s}}[u(\mathbf{a}_s,\theta)] \\ &\geq \mathbb{E}_{\theta,g,\kappa,\tau}\left[\sum_s \delta_s u(\mathbf{a}_s,\theta)\right].\end{aligned}$$

The last inequality multiplies (A.14) by $u(a_s,\theta) \leq 0$, reversing the inequality, and sums over histories. For $T = +\infty$, non-negativity of $-u$ justifies rearranging the sums. Restoring the subtracted state-wise constants proves the payoff comparison for every (A,u,κ) and τ, and hence δ-sufficiency.

A.8 Proof of Theorem 7

(2) $\implies$ (1). Immediate. It is worth noting that

$$\begin{aligned}\mathbb{E}_{\theta,f,\kappa,\sigma}\left[\sum_{t=1}^{\infty} u_t(\boldsymbol{a}_t,\theta)\right] &= \sum_{t=1}^{\infty}\sum_{a\in A_t} u_t(a,\theta)\mathrm{marg}_{A_t}\mathbb{P}_{\theta,f,\kappa,\sigma}(a), \\ &= \sum_{t=1}^{\infty}\sum_{a\in A_t} u_t(a,\theta)\mathrm{marg}_{A_t}\mathbb{P}_{\theta,g,\kappa,\tau}(a), \\ &= \mathbb{E}_{\theta,g,\kappa,\tau}\left[\sum_{t=1}^{\infty} u_t(\boldsymbol{a}_t,\theta)\right],\end{aligned}$$

because $\sum_{t=1}^{\infty}\max_{a\in A_t,\theta\in\Theta}|u_t(a_t,\theta)| < +\infty$.

(1) $\implies$ (2). The proof is by contradiction. Fix $(A_t,\kappa_t)_t$ and a finite $n \leq T$. Suppose that

there exists τ adapted to g such that

$$(\mathbb{P}_{\theta,g,\kappa,\tau}(\boldsymbol{a}_t = a))_{t\leq n,\ \theta\in\Theta,\ a\in A_t} \notin \left\{(\mathbb{P}_{\theta,f,\kappa,\sigma}(\boldsymbol{a}_t = a))_{t\leq n,\ \theta\in\Theta,\ a\in A_t} : \sigma \text{ is a strategy adapted to } f\right\},$$

where the latter set is non-empty, compact and convex.[34]

By the strict separating hyperplane theorem, there exist numbers $u_t(a,\theta)$ such that

$$\sum_{\theta\in\Theta}\mathbb{E}_{\theta,f,\kappa,\sigma}\left[\sum_{t=1}^{n}u_t(\boldsymbol{a}_t,\theta)\right] < \sum_{\theta\in\Theta}\mathbb{E}_{\theta,g,\kappa,\tau}\left[\sum_{t=1}^{n}u_t(\boldsymbol{a}_t,\theta)\right]$$

for every σ. Setting $u_t = 0$ for $t > n$ defines a separable problem, satisfying

$$\sum_{t=1}^{\infty}\max_{a\in A_t,\theta\in\Theta}|u_t(a_t,\theta)| < +\infty.$$

We thus have a contradiction with (1) when $T < \infty$.

When $T = +\infty$, the above arguments imply the existence of σ^n such that $\operatorname{marg}_{A_t}\mathbb{P}_{\theta,g,\kappa,\tau} = \operatorname{marg}_{A_t}\mathbb{P}_{\theta,f,\kappa,\sigma^n}$ for all $t \leq n$, for all θ. The statement then follows from the compactness of the space of behavioral strategies in the product topology and continuity of each marginal $\sigma \mapsto \operatorname{marg}_{A_t}\mathbb{P}_{\theta,f,\kappa,\sigma}$ in the first t coordinates of σ.

$(2) \Rightarrow (3)$. Fix ξ. Let $A_t := Y^t \times K^{t-1}$ and $\kappa_t(k_t|(y^t,l^{t-1}),k^{t-1}) = \xi_t(k_t|(y^t,l^{t-1}),k^{t-1})$ for all $t \geq 1$. Let τ choose the action (y^t,l^{t-1}) with probability one after observing (y^t,l^{t-1}), regardless of past actions.

Recall from the proof of Theorem 6 that the distribution of $\boldsymbol{a}_t$ conditional on (y^t,k^{t-1}) is the same in every state in which this history has positive probability. We denote this common distribution by $\overline{\mathbb{P}}_{\theta,g,\kappa,\tau}(\boldsymbol{a}_t \in \cdot \mid y^t,k^{t-1})$ and use it for every θ, including states in which the history has zero probability. At a history impossible in every state, we choose any distribution over A_t, again the same for every θ (and use the same convention under f and σ).

By (2), there exists a strategy σ under f that induces the same distribution over actions as τ at every date and in every state. Fix such a σ. Define

$$\gamma_t(y^t,l^{t-1} \mid x^t,k^{t-1}) := \overline{\mathbb{P}}_{\theta,f,\kappa,\sigma}\left(\boldsymbol{a}_t = (y^t,l^{t-1}) \mid x^t,k^{t-1}\right),$$

[34] Compactness follows because the behavioral strategy space is a compact product of probability simplexes and each displayed probability depends continuously on finitely many strategy coordinates. Convexity follows by randomizing between strategies at the outset and using a realization-equivalent behavioral strategy.

where $\overline{\mathbb{P}}_{\theta,f,\kappa,\sigma}$ is defined as in the proof of Theorem 6. We stress that they are independent of θ, hence the γ_t are well-defined.

We argue that $\mathbb{P}_{\theta,f,\kappa,\sigma}(x^t,k^{t-1}) = \mathbb{P}_{\theta,f,\xi\circ\gamma}(x^t,k^{t-1})$ for every t, θ, and (x^t,k^{t-1}). The proof is by induction. This is clearly true at period 1 since the first-period signal has distribution $f_1(\cdot \mid \theta)$ under both measures. By induction, suppose they assign the same probability to every history (x^s,k^{s-1}). After such a history, they both draw k_s according to $(\xi_s \circ \gamma_s)(\cdot \mid x^s,k^{s-1})$ and then draw x_{s+1} according to $f_{s+1}(\cdot \mid x^s,k^s,\theta)$. They, therefore, assign the same probability to every history (x^{s+1},k^s) as well.

Under τ, the date-t action distribution is $\mathbb{P}_{\theta,g,\xi}(y^t,l^{t-1})$. Consequently, (2) implies

$$\mathbb{P}_{\theta,g,\xi}(y^t,l^{t-1}) = \mathbb{P}_{\theta,f,\kappa,\sigma}\left(\boldsymbol{a}_t = (y^t,l^{t-1})\right) = \sum_{x^t,k^{t-1}} \gamma_t(y^t,l^{t-1} \mid x^t,k^{t-1})\mathbb{P}_{\theta,f,\xi\circ\gamma}(x^t,k^{t-1}),$$

which is (3).

$(3) \Rightarrow (2)$. Fix a separable problem $(A_t,u_t,\kappa_t)_t$ and a strategy τ under g. Define

$$\xi_t(k_t \mid (y^t,l^{t-1}),k^{t-1}) := \sum_{a\in A_t} \overline{\mathbb{P}}_{\theta,g,\kappa,\tau}(\boldsymbol{a}_t = a \mid y^t,l^{t-1})\kappa_t(k_t \mid a,k^{t-1}),$$

and extend ξ_t arbitrarily to simulated histories of other lengths. Let γ satisfy (3) for this ξ. Define σ by

$$\sigma_t(a \mid a^{t-1},x^t,k^{t-1}) := \sum_{y^t,l^{t-1}} \gamma_t(y^t,l^{t-1} \mid x^t,k^{t-1})\overline{\mathbb{P}}_{\theta,g,\kappa,\tau}(\boldsymbol{a}_t = a \mid y^t,l^{t-1}).$$

for all (a^{t-1},x^t,k^{t-1}). Note that σ is independent of the profiles of past actions.

We first show that $\mathbb{P}_{\theta,g,\kappa,\tau}$ and $\mathbb{P}_{\theta,g,\,\xi}$ assign the same probability to every history (y^t,l^{t-1}). Both have first-period signal distribution $g_1(\cdot \mid \theta)$. Suppose they agree on the probability of every (y^s,l^{s-1}). At each such history of positive probability, the definition of ξ_s implies that both draw l_s according to $\xi_s(\cdot \mid (y^s,l^{s-1}),l^{s-1})$. Both then draw y_{s+1} according to $g_{s+1}(\cdot \mid y^s,l^s,\theta)$ and, therefore, agree on the probability of every (y^{s+1},l^s). Accordingly,

$$\mathbb{P}_{\theta,g,\kappa,\tau}(y^t,l^{t-1}) = \mathbb{P}_{\theta,g,\xi}(y^t,l^{t-1}) \qquad \forall\, t,\theta,\ \&\ (y^t,l^{t-1}). \tag{A.17}$$

We next show that $\mathbb{P}_{\theta,f,\kappa,\sigma}$ and $\mathbb{P}_{\theta,f,\xi\circ\gamma}$ assign the same probability to every history (x^t,k^{t-1}).

The definitions of σ_t and ξ_t imply

$$\sum_{a \in A_t} \sigma_t(a \mid x^t, k^{t-1}) \kappa_t(k_t \mid a, k^{t-1}) = (\xi_t \circ \gamma_t)(k_t \mid x^t, k^{t-1}). \tag{A.18}$$

Both measures have first-period signal distribution $f_1(\cdot \mid \theta)$. Suppose they agree on the probability of every (x^s, k^{s-1}). By (A.18), both draw k_s according to $(\xi_s \circ \gamma_s)(\cdot \mid x^s, k^{s-1})$ at each such history of positive probability. Both then draw x_{s+1} according to $f_{s+1}(\cdot \mid x^s, k^s, \theta)$. Accordingly, they agree on the probability of every (x^{s+1}, k^s) and so

$$\mathbb{P}_{\theta, f, \kappa, \sigma}(x^t, k^{t-1}) = \mathbb{P}_{\theta, f, \xi \circ \gamma}(x^t, k^{t-1}) \qquad \forall\, t, \theta, \ \& \ (x^t, k^{t-1}). \tag{A.19}$$

We can now verify that σ reproduces the action distributions generated by τ:

$$\begin{aligned}
\mathbb{P}_{\theta, f, \kappa, \sigma}(\boldsymbol{a}_t = a) &= \sum_{y^t, l^{t-1}} \overline{\mathbb{P}}_{\theta, g, \kappa, \tau}(\boldsymbol{a}_t = a \mid y^t, l^{t-1}) \sum_{x^t, k^{t-1}} \gamma_t(y^t, l^{t-1} \mid x^t, k^{t-1}) \mathbb{P}_{\theta, f, \xi \circ \gamma}(x^t, k^{t-1}) \\
&= \sum_{y^t, l^{t-1}} \overline{\mathbb{P}}_{\theta, g, \kappa, \tau}(\boldsymbol{a}_t = a \mid y^t, l^{t-1}) \mathbb{P}_{\theta, g, \xi}(y^t, l^{t-1}) = \mathbb{P}_{\theta, g, \kappa, \tau}(\boldsymbol{a}_t = a).
\end{aligned}$$

The first equality uses the definition of σ_t and (A.19); the second uses (3); and the third uses (A.17) and the law of total probability. As this holds for every t, $a \in A_t$, and θ, we conclude (2). ■

A.9 Proof of Proposition 4

Proof Assume that, for every $k \in K$, $h_f(\cdot \mid k, \cdot)$ Blackwell dominates $h_g(\cdot \mid k, \cdot)$. For each $k \in K$, choose a Markov kernel $\lambda_k \colon X \to \Delta(Y)$ such that

$$h_g(y \mid k, \theta) = \sum_{x \in X} h_f(x \mid k, \theta) \lambda_k(y \mid x) \quad \forall\, y, k, \theta. \tag{A.20}$$

Fix an arbitrary test kernel ξ. We construct kernels $\gamma_t \colon X^t \times K^{t-1} \to \Delta(Y^t \times K^{t-1})$ satisfying the Δ-sufficiency equation (Theorem 7 Item (3)). When the simulated and realized control histories coincide, write $\xi_s(k_s \mid y^s, k^{s-1}) := \xi_s(k_s \mid (y^s, k^{s-1}), k^{s-1})$. For $x^t \in X^t$, $y^t \in Y^t$, and $k^{t-1} \in K^{t-1}$, define

$$M_t(y^t, k^{t-1} \mid x^t) := \prod_{s=1}^{t} \lambda_{k_{s-1}}(y_s \mid x_s) \prod_{s=1}^{t-1} \xi_s(k_s \mid y^s, k^{s-1}),$$

where the empty product is one, and let $D_t(x^t, k^{t-1}) := \sum_{\hat{y}^t \in Y^t} M_t(\hat{y}^t, k^{t-1} \mid x^t)$.

Define γ_t as follows. If $D_t(x^t, k^{t-1}) > 0$, set

$$\gamma_t(y^t, l^{t-1} \mid x^t, k^{t-1}) := \begin{cases} \dfrac{M_t(y^t, k^{t-1} \mid x^t)}{D_t(x^t, k^{t-1})}, & \text{if} \quad l^{t-1} = k^{t-1}, \\ 0, & \text{otherwise.} \end{cases}$$

If $D_t(x^t, k^{t-1}) = 0$, choose an arbitrary $\nu_t \in \Delta(Y^t)$ and set

$$\gamma_t(y^t, l^{t-1} \mid x^t, k^{t-1}) := \begin{cases} \nu_t(y^t), & \text{if } l^{t-1} = k^{t-1}, \\ 0, & \text{otherwise.} \end{cases}$$

We first record the distribution induced by f and $\xi \circ \gamma$.

Claim 1 *For every t, $\mathbb{P}_{\theta, f, \xi \circ \gamma}(x^t, k^{t-1}) = \left[\prod_{s=1}^{t} h_f(x_s \mid k_{s-1}, \theta)\right] D_t(x^t, k^{t-1})$.*

Proof of Claim 1 For $t = 1$, this is immediate since $D_1(x_1, k_0) = \sum_{y_1} \lambda_{k_0}(y_1 \mid x_1) = 1$. Suppose the claim holds at t. If $D_t(x^t, k^{t-1}) > 0$, then, by the definition of composition,

$$\begin{aligned} (\xi_t \circ \gamma_t)(k_t \mid x^t, k^{t-1}) &= \sum_{y^t, l^{t-1}} \gamma_t(y^t, l^{t-1} \mid x^t, k^{t-1}) \xi_t(k_t \mid (y^t, l^{t-1}), k^{t-1}) \\ &= \frac{\sum_{y^t} M_t(y^t, k^{t-1} \mid x^t) \xi_t(k_t \mid y^t, k^{t-1})}{D_t(x^t, k^{t-1})} \\ &= \frac{\sum_{y^t} \prod_{s=1}^{t} \lambda_{k_{s-1}}(y_s \mid x_s) \prod_{s=1}^{t} \xi_s(k_s \mid y^s, k^{s-1})}{D_t(x^t, k^{t-1})} \\ &= \frac{\sum_{y^{t+1}} \prod_{s=1}^{t+1} \lambda_{k_{s-1}}(y_s \mid x_s) \prod_{s=1}^{t} \xi_s(k_s \mid y^s, k^{s-1})}{D_t(x^t, k^{t-1})} \\ &= \frac{D_{t+1}(x^{t+1}, k^t)}{D_t(x^t, k^{t-1})}, \end{aligned}$$

where the third and fourth equalities follow from the definition of $M_t(y^t, k^{t-1} \mid x^t)$ and $\sum_{y_{t+1}} \lambda_{k_t}(y_{t+1} \mid x_{t+1}) = 1$. Hence, the claim holds at $t + 1$. If $D_t(x^t, k^{t-1}) = 0$, then $M_t(y^t, k^{t-1} \mid x^t) = 0$ for every y^t, so $D_{t+1}(x^{t+1}, k^t) = 0$ for every continuation (x_{t+1}, k_t). Since (x^t, k^{t-1}) has zero probability by the induction hypothesis, the claim again holds at $t + 1$. ■

We now verify Δ-sufficiency. Fix $y^t \in Y^t$ and $l^{t-1} \in K^{t-1}$. By Claim 1 and the definition

of γ_t,

$$\begin{aligned}
&\sum_{x^t,k^{t-1}} \mathbb{P}_{\theta,f,\xi\circ\gamma}(x^t,k^{t-1})\gamma_t(y^t,l^{t-1}\mid x^t,k^{t-1}) \\
&= \sum_{x^t}\left[\prod_{s=1}^{t} h_f(x_s\mid l_{s-1},\theta)\right] M_t(y^t,l^{t-1}\mid x^t) \\
&= \sum_{x^t}\left[\prod_{s=1}^{t} h_f(x_s\mid l_{s-1},\theta)\lambda_{l_{s-1}}(y_s\mid x_s)\right]\prod_{s=1}^{t-1}\xi_s(l_s\mid y^s,l^{s-1}) \\
&= \prod_{s=1}^{t} h_g(y_s\mid l_{s-1},\theta)\prod_{s=1}^{t-1}\xi_s(l_s\mid y^s,l^{s-1}) = \mathbb{P}_{\theta,g,\xi}(y^t,l^{t-1}).
\end{aligned}$$

The third equality uses (A.20) for each realized control l_{s-1}. We conclude that f is Δ-sufficient for g. ■

A.10 Proof of Proposition 5

Sufficiency follows from the text. We prove necessity by contraposition. Suppose that there exists an arm u^* such that $h_f(\cdot|u^*,\cdot)$ does not Blackwell dominate $h_g(\cdot|u^*,\cdot)$. By Blackwell's theorem, there exist a full-support prior p and a finite static decision problem with payoff vectors $\bar{u}^1,\dots,\bar{u}^n \in \mathbb{R}^{|\Theta|}$ such that $\mathbb{E}_g[\varphi(\mu)] > \mathbb{E}_f[\varphi(\mu)]$, with $\varphi(\mu) = \max_{j=1,\dots,n}\bar{u}^j\cdot\mu$ and $\mathbb{E}_i$ the expectation with respect to the distribution over posteriors induced by $h_i(\cdot|u^*,\cdot)$, $i=f,g$.

Choose a supporting vector $s\in\partial\varphi(p)$, and define $\psi(\mu) \coloneqq \varphi(\mu)-\varphi(p)-s\cdot(\mu-p)$. By construction, $\mathbb{E}_g[\psi(\mu)]-\mathbb{E}_f[\psi(\mu)]>0$. For $\eta>0$, define

$$\psi_\eta(\mu) \coloneqq \max\left\{0, \max_{j=1,\dots,n}\left[(\bar{u}^j-s)\cdot\mu-\varphi(p)+s\cdot p-\eta\right]\right\}.$$

As $\eta\downarrow 0$, $\psi_\eta(\mu)$ converges point-wise to $\psi(\mu)$. Since the signal sets are finite, expectations are continuous in η. Hence, for sufficiently small $\eta>0$, $\mathbb{E}_g[\psi_\eta(\mu)]>\mathbb{E}_f[\psi_\eta(\mu)]$.

Now construct a finite bandit menu. The menu contains the arm u^* and, for each $j=1,\dots,n$, a continuation arm

$$u^j \coloneqq u^* + \bar{u}^j - s + [-\varphi(p)+s\cdot p-\eta]\,\mathbf{1},$$

where $\mathbf{1}$ is the vector that equals one in every state. At the prior p,

$$u^j\cdot p = u^*\cdot p + \bar{u}^j\cdot p - \varphi(p) - \eta \le u^*\cdot p-\eta.$$

Thus, u^* is strictly better than u^j, $j \neq 0$ at the prior.

Consider the two-period discount factor $\delta = (1-\varepsilon, \varepsilon, 0, \ldots)$. For $\varepsilon > 0$ sufficiently small, choosing u^* at the first period is optimal under both experiments.[35]

After u^* is chosen first and the reward is observed, the second-period value at posterior μ is

$$\max\{u^* \cdot \mu, u^1 \cdot \mu, \ldots, u^n \cdot \mu\} = u^* \cdot \mu + \psi_\eta(\mu).$$

Therefore, under experiment $i \in \{f, g\}$, the Bayesian value in the constructed two-period bandit problem is

$$(1-\varepsilon) u^* \cdot p + \varepsilon \, \mathbb{E}_i \left[u^* \cdot \mu + \psi_\eta(\mu) \right].$$

Since posterior beliefs average to p, $\mathbb{E}_i[u^* \cdot \mu] = u^* \cdot p$ under both experiments, and so the value difference between g and f is $\varepsilon \left(\mathbb{E}_g[\psi_\eta(\mu)] - \mathbb{E}_f[\psi_\eta(\mu)] \right) > 0$, so f does not dominate g in every bandit problem.

References

BERGEMANN, D. AND U. HEGE (2005): "The Financing of Innovation: Learning and Stopping," *RAND Journal of Economics*, 36, 719–752.

BERGEMANN, D. AND J. VÄLIMÄKI (1996): "Learning and Strategic Pricing," *Econometrica*, 64, 1125–1149.

BIRNBAUM, A. (1961): "On the foundations of statistical inference: Binary experiments," *The Annals of Mathematical Statistics*, 32, 414–435.

BLACKWELL, D. (1951): "Comparison of experiments," in *Proceedings of the Second Berkeley Symposium on Mathematical Statistics*, 93–102.

——— (1953): "Equivalent comparisons of experiments," *The Annals of Mathematical Statistics*, 265–272.

BLACKWELL, D. A. AND M. A. GIRSHICK (1954): *Theory of games and statistical decisions*, Wiley.

[35] To see this explicitly, let $C \coloneqq \max_{u \in \{u^*, u^1, \ldots, u^n\}} \max_{\theta \in \Theta} |u(\theta)|$. Since the largest continuation payoff is at most $2C$, choosing ε so that $(1-\varepsilon)\eta > 2\varepsilon C$ guarantees that u^0 is optimal at the first period, regardless of the experiment.

BROOKS, B., A. FRANKEL, AND E. KAMENICA (2024): "Comparison of signals," *American Economic Review*, 114, 2981–3006.

DE OLIVEIRA, H. (2018): "Blackwell's informativeness theorem using diagrams," *Games and Economic Behavior*, 109, 126–131.

DEJARNETTE, P., D. DILLENBERGER, D. GOTTLIEB, AND P. ORTOLEVA (2020): "Time lotteries and stochastic impatience," *Econometrica*, 88, 619–656.

GREENSHTEIN, E. (1996): "Comparison of sequential experiments," *The Annals of Statistics*, 24, 436–448.

JOVANOVIC, B. (1979): "Job Matching and the Theory of Turnover," *Journal of Political Economy*, 87, 972–990.

LE CAM, L. (1996): "Comparison of experiments: A short review," *Lecture Notes-Monograph Series*, 127–138.

LEHRER, E. AND T. WANG (2024): "The value of information in stopping problems," *Economic Theory*, 619–648.

LIANG, A., X. MU, AND V. SYRGKANIS (2018): "Optimal and myopic information acquisition," in *Proceedings of the 2018 ACM Conference on Economics and Computation*, 45–46.

——— (2022): "Dynamically aggregating diverse information," *Econometrica*, 90, 47–80.

MILLER, R. A. (1984): "Job Matching and Occupational Choice," *Journal of Political Economy*, 92, 1086–1120.

MOSCARINI, G. AND L. SMITH (2002): "The law of large demand for information," *Econometrica*, 70, 2351–2366.

MU, X., L. POMATTO, P. STRACK, AND O. TAMUZ (2021): "From Blackwell dominance in large samples to Rényi divergences and back again," *Econometrica*, 89, 475–506.

ROBERTS, K. AND M. L. WEITZMAN (1981): "Funding Criteria for Research, Development, and Exploration Projects," *Econometrica*, 49, 1261–1288.

ROTHSCHILD, M. (1974): "A Two-Armed Bandit Theory of Market Pricing," *Journal of Economic Theory*, 9, 185–202.

SCHMITT-GROHE, S. AND M. URIBE (2003): "Closing small open economy models," *Journal of International Economics*, 61, 163–185.

SHAKED, M. AND J. G. SHANTHIKUMAR (2007): *Stochastic orders*, Springer New York.

SHERMAN, S. (1951): "On a theorem of Hardy, Littlewood, Polya, and Blackwell," *Proceedings of the National Academy of Sciences*, 37, 826–831.

STEIN, C. (1951): "Notes on the comparison of experiments," University of Chicago.

STRASSER, H. (2011): *Mathematical theory of statistics: statistical experiments and asymptotic decision theory*, vol. 7, Walter de Gruyter.

TORGERSEN, E. (1991): *Comparison of statistical experiments*, vol. 36, Cambridge University Press.

TORGERSEN, E. N. (1970): "Comparison of experiments when the parameter space is finite," *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 16, 219–249.

UZAWA, H. (1968): *Time preference, the consumption function and the optimum asset holdings*, Wolfe, J.N. (Ed.), Value, Capital and Growth: Papers in Honor of Sir John Hicks. The University of Edinburgh Press, Edinburgh, pp. 485–504.

YANG, K. H. AND A. K. ZENTEFIS (2023): "Extreme points and first-order stochastic dominance: theory and applications," Https://arxiv.org/pdf/2302.03135.pdf.

B Online Appendix

B.1 The Precision-Timing Tradeoff

This appendix expands on the interpretation of Corollary 1. Suppose that $T < +\infty$, and let $C(t) \coloneqq \sum_{s=t}^{T} \delta_s$ be the continuation weight from date t onwards. Since the information arrives by period T, every arrival time has continuation weight at least $C(T) = \delta_T$.

We first record a simple construction. Fix $\lambda \in [0,1]$ and a random arrival time $\boldsymbol{T}$. A delay $\widehat{\boldsymbol{T}}^{\lambda}$ that still arrives by period T can satisfy

$$\mathbb{E}\left[C(\widehat{\boldsymbol{T}}^{\lambda})\right] = \lambda \mathbb{E}\left[C(\boldsymbol{T})\right]$$

if, and only if, $\lambda\mathbb{E}[C(\boldsymbol{T})] \geq \delta_T$. Necessity follows from the lower bound above. To see sufficiency, if $\mathbb{E}[C(\boldsymbol{T})] = \delta_T$, set $\widehat{\boldsymbol{T}}^{\lambda} = \boldsymbol{T}$. Otherwise, let

$$\widehat{\boldsymbol{T}}^{\lambda} = \begin{cases} \boldsymbol{T}, & U \leq \dfrac{\lambda\mathbb{E}[C(\boldsymbol{T})] - \delta_T}{\mathbb{E}[C(\boldsymbol{T})] - \delta_T}, \\ T, & \text{otherwise}, \end{cases}$$

where U is uniformly distributed over $[0,1]$ and independent of $\boldsymbol{T}$. By construction, $\boldsymbol{T} \leq \widehat{\boldsymbol{T}}^{\lambda} \leq T$ and the equality holds. In a nutshell, we either keep the original arrival time or delay the information by $T - \boldsymbol{T}$ periods, until period T. The shift equalizes receiving the payoff λV from period $\boldsymbol{T}$ onwards with receiving the payoff V from period $\widehat{\boldsymbol{T}}^{\lambda}$ onwards, in expected discounted terms.

We now apply this construction to the precision-timing comparison. First, suppose that $r^* \leq 1$. If $r^*\mathbb{E}[C(\boldsymbol{T}_F)] < \delta_T$, the experiment f cannot be δ-sufficient for g, since $\mathbb{E}[C(\boldsymbol{T}_G)] \geq \delta_T$. Otherwise, let $\widehat{\boldsymbol{T}}_F \coloneqq \widehat{\boldsymbol{T}}_F^{r^*}$ be the adjusted arrival time obtained from the construction above with $\lambda = r^*$ and $\boldsymbol{T} = \boldsymbol{T}_F$. Then $\mathbb{E}[C(\widehat{\boldsymbol{T}}_F)] = r^*\mathbb{E}[C(\boldsymbol{T}_F)]$, and the condition in Corollary 1 is equivalent to

$$\mathbb{E}\left[C(\widehat{\boldsymbol{T}}_F)\right] \geq \mathbb{E}\left[C(\boldsymbol{T}_G)\right].$$

Thus, the precision adjustment in the comparison can be represented as a delay of $\boldsymbol{T}_F$. A sufficient condition is for $\boldsymbol{T}_G$ to first-order stochastically dominate $\widehat{\boldsymbol{T}}_F$.

Second, suppose that $r^* > 1$. It is tempting to interpret the factor r^* as shifting $\boldsymbol{T}_F$ earlier, but this cannot always be done literally: if $\boldsymbol{T}_F = 1$, then $r^*C(\boldsymbol{T}_F) > 1$ is not

attainable by any earlier arrival time. Instead, we shift $\boldsymbol{T}_G$ later. If $\frac{1}{r^*}\mathbb{E}[C(\boldsymbol{T}_G)] < \delta_T$, the experiment f is δ-sufficient for g regardless of its arrival time, since $\mathbb{E}[C(\boldsymbol{T}_F)] \geq \delta_T$. Otherwise, let $\widehat{\boldsymbol{T}}_G \coloneqq \widehat{\boldsymbol{T}}_G^{1/r^*}$ be the adjusted arrival time obtained from the construction above with $\lambda = 1/r^*$ and $\boldsymbol{T} = \boldsymbol{T}_G$. Then $\mathbb{E}[C(\widehat{\boldsymbol{T}}_G)] = \frac{1}{r^*}\mathbb{E}[C(\boldsymbol{T}_G)]$, and the condition in Corollary 1 is equivalent to

$$\mathbb{E}\left[C(\boldsymbol{T}_F)\right] \geq \mathbb{E}\left[C(\widehat{\boldsymbol{T}}_G)\right].$$

Thus, if h_X is more informative than h_Y, the lower precision of h_Y can be represented as a delay of the arrival time of the information from h_Y. A sufficient condition is for $\widehat{\boldsymbol{T}}_G$ to first-order stochastically dominate $\boldsymbol{T}_F$.

If $T = \infty$, the construction is always feasible since delays are always possible.

B.2 Known Flow-Problem Comparisons

Throughout this appendix, assume $T < \infty$. The comparisons in Theorem 1, Theorem 3, and Corollary 2 vary the decision problem. We now fix (A, u) and vary the discount factor. We consider two cases: all discount factors $\delta \in \Delta$, and all decreasing discount factors $\delta \in \Delta^{\downarrow}$.

For a static experiment $h\colon \Theta \to \Delta(Z)$, let $W_p^u(h) \coloneqq \sum_{z\in Z} \max_{a\in A} \sum_{\theta\in\Theta} p(\theta)h(z \mid \theta)u(a,\theta)$. For a dynamic experiment f, set $w_t^f(p,u) \coloneqq W_p^u(f^t)$, i.e., $w_t^f(p,u)$ is the value of the fixed problem (A,u), at prior p, using the signals observed by date t. By Theorem 1,

$$V_{p,\delta}^u(f) \coloneqq \sup_{\sigma} E_{p,f,\sigma}\left[\sum_{t=1}^{T} \delta_t u(a_t,\theta)\right] = \sum_{t=1}^{T} \delta_t w_t^f(p,u). \tag{B.1}$$

Definition 3 *Fix $p \in \Delta(\Theta)$ and (A,u). We write $f \succeq_{p,u}^{\Delta} g$ if $V_{p,\delta}^u(f) \geq V_{p,\delta}^u(g)$ for every $\delta \in \Delta$. We write $f \succeq_{p,u}^{\Delta^{\downarrow}} g$ if the same inequality holds for every $\delta \in \Delta^{\downarrow}$.*

Lemma 1 *Fix $p \in \Delta(\Theta)$ and (A,u). Then*

$$f \succeq_{p,u}^{\Delta} g \quad\Longleftrightarrow\quad w_t^f(p,u) \geq w_t^g(p,u) \text{ for every } t. \tag{B.2}$$

Moreover,

$$f \succeq_{p,u}^{\Delta^{\downarrow}} g \quad\Longleftrightarrow\quad \sum_{s=1}^{t} w_s^f(p,u) \geq \sum_{s=1}^{t} w_s^g(p,u) \text{ for every } t. \tag{B.3}$$

Proof By (B.1), the comparison is $\sum_t \delta_t(w_t^f(p,u) - w_t^g(p,u)) \geq 0$, which we apply to the extreme points to get (B.2) and (B.3). ∎

Equation (B.2) is the fixed-(p,u) counterpart of Theorem 3: it requires higher value date by date, rather than sufficiency date by date. Equation (B.3) is the fixed-(p,u) counterpart of Corollary 2: it requires higher value for each running average, rather than sufficiency of each running-average experiment.

We now allow the prior to vary. For a static experiment $h\colon \Theta \to \Delta(Z)$, define

$$R_h^u := \left\{ r \in \mathbb{R}^\Theta : r_\theta = \sum_{z\in Z} h(z \mid \theta) \sum_{a\in A} \alpha(a \mid z) u(a,\theta) \text{ for some } \alpha : Z \to \Delta(A) \right\}.$$

Thus R_h^u is the set of state-contingent payoff vectors attainable in the fixed problem (A,u) after observing h, and $W_p^u(h) = \max_{r\in R_h^u} p\cdot r$. For sets $R_1,\dots,R_t \subseteq \mathbb{R}^\Theta$, write $t^{-1}\sum_{s=1}^t R_s := \{t^{-1}\sum_{s=1}^t r_s\colon r_s \in R_s \text{ for every } s\}$, and let $\mathbb{R}^\Theta_+ := \{z \in \mathbb{R}^\Theta\colon z_\theta \geq 0 \text{ for every } \theta\}$.

Definition 4 *Fix (A,u). We write $f \succeq_u^\Delta g$ if $V^u_{p,\delta}(f) \geq V^u_{p,\delta}(g)$ for every $p \in \Delta(\Theta)$ and every $\delta \in \Delta$. We write $f \succeq_u^{\Delta^\downarrow} g$ if the same inequality holds for every $p \in \Delta(\Theta)$ and every $\delta \in \Delta^\downarrow$.*

Proposition 8 *Fix (A,u). Then*

$$f \succeq_u^\Delta g \iff R^u_{g^t} \subseteq R^u_{f^t} - \mathbb{R}^\Theta_+ \text{ for every } t. \tag{B.4}$$

Moreover,

$$f \succeq_u^{\Delta^\downarrow} g \iff \frac{1}{t}\sum_{s=1}^t R^u_{g^s} \subseteq \frac{1}{t}\sum_{s=1}^t R^u_{f^s} - \mathbb{R}^\Theta_+ \text{ for every } t. \tag{B.5}$$

Proof We first record the static step. For static experiments h and k, $W_p^u(h) \geq W_p^u(k)$ for every $p \in \Delta(\Theta)$ if and only if $R_k^u \subseteq R_h^u - \mathbb{R}^\Theta_+$. If the inclusion holds, every payoff vector attainable under k is weakly dominated, state by state, by some payoff vector attainable under h, so the value comparison holds for every prior. Conversely, if the inclusion fails, by the strict separating hyperplane theorem, there is a nonzero $q \in \mathbb{R}^\Theta_+$ such that $\max_{r\in R_k^u} q\cdot r > \max_{r\in R_h^u} q\cdot r$. Normalizing q produces a prior at which k has strictly higher value than h.

For Δ, (B.2) says that $W_p^u(f^t) \geq W_p^u(g^t)$ for every prior p and every date t. Applying the static step date by date yields (B.4). Conversely, from (B.4), $W_p^u(f^t) \geq W_p^u(g^t)$ for every prior and every date, so (B.2) delivers $f \succeq_u^\Delta g$.

For $\Delta^\downarrow$, (B.3) says that, for every prior p and every t, the value of the experiment that draws s uniformly from $\{1,\dots,t\}$ and then draws a signal from f^s is higher than the

corresponding value for g. The payoff set of this experiment is $t^{-1}\sum_{s=1}^{t} R^u_{f^s}$, and similarly for g. Applying the static step to these two experiments delivers (B.5). Conversely, assume (B.5) then use (B.3) prior-by-prior to get $f \succeq_u^{\Delta^\downarrow} g$. ■

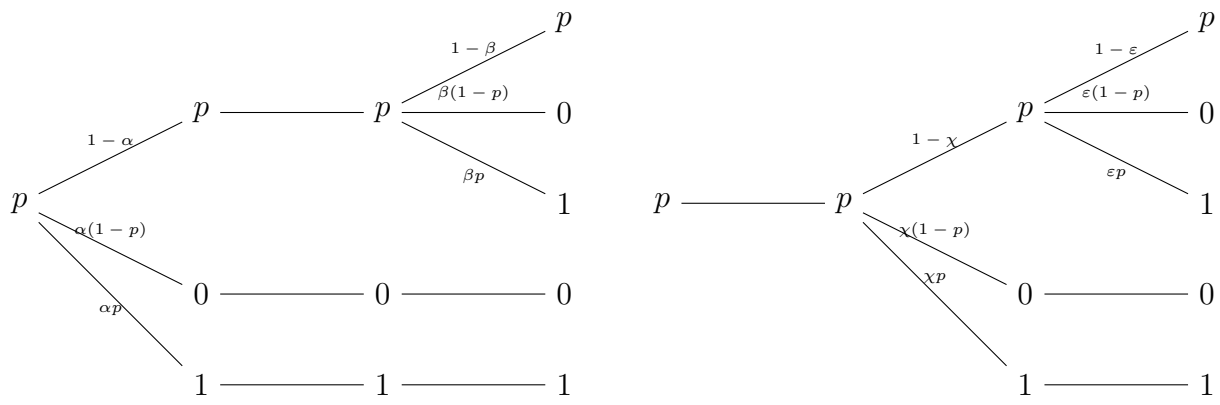
B.3 Non-monotonicity in Patience

We now illustrate how the ranking of experiments may change multiple times as the decision-maker becomes more impatient. Let $\underline{\delta}$ and $\overline{\delta}$ be two discount factors such that $\overline{\delta}$ first-order stochastically dominates $\underline{\delta}$. We interpret first-order stochastic shifts in discount factors as the decision-maker becoming more patient. For instance, if we restrict attention to geometric discounting, first-order stochastic shifts correspond to higher discount rates. If f is $\underline{\delta}$ and $\overline{\delta}$-sufficient for g, it is natural to query whether f is δ-sufficient for g for all δ, which first-order stochastically dominate $\underline{\delta}$ and are dominated by $\overline{\delta}$. An immediate consequence of Proposition 2 is that the statement is true when $T = 2$ since $\underline{\delta}$ and $\overline{\delta}$ are the extreme points of the first-order stochastic dominance intervals. With more periods, it is no longer true that $\underline{\delta}$ and $\overline{\delta}$ are the extreme points. For instance, if $\underline{\delta} = (1, 0, 0)$ and $\overline{\delta} = (0, 0, 1)$, then the discount factor $\delta = (0, 1, 0)$ first-order stochastically dominates $\underline{\delta}$, is dominated by $\overline{\delta}$, and cannot be written as convex combination of the two. (See Theorem 1 in Yang and Zentefis (2023) for a characterization of the extreme points of first-order stochastic dominance intervals.) Yet, this does not immediately imply that f is not δ-sufficient for g. The following example, however, shows this is the case.

There are two states θ_0 and θ_1, three periods, and two experiments f and g. Let $p \in (0, 1)$ be the prior belief of $[\boldsymbol{\theta} = \theta_1]$. We describe the evolution of the beliefs induced by f and g in Figure 1. (It is straightforward to construct the associated signals X_t (resp., Y_t) and kernels f_t (resp., g_t).)

The experiment f is informative at the first and third periods, while g is informative at the second and third periods. Let $v : [0, 1] \to \mathbb{R}$ be a convex function and denote $V(p) := (1 - p)v(0) + pv(1) \geq v(p)$. As already explained, δ-sufficiency is equivalent to checking whether the discounted expected value $V_f(\delta, v)$ is greater than $V_g(\delta, v)$ for all convex functions v, where

$$V_f(\delta, v) := v(p)\left[(\delta_1 + \delta_2)(1 - \alpha) + \delta_3(1 - \alpha)(1 - \beta)\right] + V(p)\left[(\delta_1 + \delta_2)\alpha + \delta_3(\alpha + (1 - \alpha)\beta)\right],$$
$$V_g(\delta, v) := v(p)\left[\delta_1 + \delta_2(1 - \chi) + \delta_3(1 - \chi)(1 - \varepsilon)\right] + V(p)\left[(\delta_2 + \delta_3)\chi + \delta_3(1 - \chi)\varepsilon\right].$$



(a) Experiment f (b) Experiment g

Figure 1: Discounting, first-order stochastic dominance, and sufficiency

Clearly, if $v(p) = V(p)$, $V_f(\delta, v) = V_g(\delta, v)$ for all δ. Assume that $V(p) > v(p)$. Comparing $V_f(\delta, v)$ and $V_g(\delta, v)$ is thus equivalent to comparing the coefficients in front of $V(p)$ in the above expression. Consider the following (geometric) discount factors: $\underline{\delta} = (1, 0, 0)$, $\overline{\delta} = (1/3, 1/3, 1/3)$ and $\delta^* = (4/7, 2/7, 1/7)$ and observe that $V_f(\delta, v) > V_g(\delta, v)$ is equivalent to

$$(\delta_1 + \delta_2)\alpha + \delta_3(\alpha + (1 - \alpha)\beta) > (\delta_2 + \delta_3)\chi + \delta_3(1 - \chi)\varepsilon.$$

We can verify that with the parametrization $\varepsilon = 0$, $\beta = 1$, $\alpha = 0$, $\chi \in (1/3, 1/2)$, $V_f(\underline{\delta}, v) = V_g(\underline{\delta}, v)$, $V_f(\delta^*, v) < V_g(\delta^*, v)$, and $V_f(\overline{\delta}, v) > V_g(\overline{\delta}, v)$ for all convex functions v such that $V(p) > v(p)$. With an arbitrarily small, but positive, α, we have that $V_f(\underline{\delta}, v) > V_g(\underline{\delta}, v)$, without affecting the other comparisons. So, the experiment f is more informative than g for a very patient or a very impatient decision-maker, and the converse holds for a moderately patient decision-maker. We have ambiguous comparative statics with respect to impatience.

B.4 Sufficiently Patient Decision-Makers

Assume throughout this appendix that $T < +\infty$, and remain in the uncontrolled case. For $\hat{\delta}, \delta \in \Delta$, write $\hat{\delta} \succeq_{\text{FOSD}} \delta$ if $\sum_{s=1}^{t} \hat{\delta}_s \leq \sum_{s=1}^{t} \delta_s$ for all $t < T$. For a fixed δ, denote the set of discount weights more patient than δ by $\Delta^{\succeq \delta} \coloneqq \{\hat{\delta} \in \Delta \colon \hat{\delta} \succeq_{\text{FOSD}} \delta\}$.

For $R \subseteq \{1, \ldots, T\}$ with $T \in R$, let $\rho_R \colon \{1, \ldots, T\} \to R$ be given by $\rho_R(t) \coloneqq \min\{r \in R \colon r \geq$

$t\}$, and let $\delta^R \in \Delta$ be given by

$$\delta^R_s \coloneqq \sum_{\{t\colon \rho_R(t)=s\}} \delta_t, \quad \text{for } s = 1, \dots, T. \tag{B.6}$$

ρ_R moves each date to the next date in R, and δ^R is the induced discount factor.

Lemma 2 *The extreme points of $\Delta^{\succeq\delta}$ are the distinct discount factors δ^R in (B.6), where R ranges over all subsets of $\{1, \dots, T\}$ with $T \in R$.*

Proof Let $D_0 \coloneqq 0$, $D_t \coloneqq \sum_{s=1}^t \delta_s$, $\hat{D}_0 \coloneqq 0$, and $\hat{D}_t \coloneqq \sum_{s=1}^t \hat{\delta}_s$. Then $\hat{\delta} \in \Delta^{\succeq\delta}$ if and only if $0 = \hat{D}_0 \leq \hat{D}_1 \leq \cdots \leq \hat{D}_T = 1$ and $\hat{D}_t \leq D_t$ for every $t < T$.

We first show that, at an extreme point, for every $t < T$, either $\hat{D}_t = \hat{D}_{t-1}$ or $\hat{D}_t = D_t$. Suppose instead that $\hat{D}_{t-1} < \hat{D}_t < D_t$ for some $t < T$. Let u be the largest date in $\{t, \dots, T-1\}$ such that $\hat{D}_t = \cdots = \hat{D}_u$. Then $\hat{D}_u < \hat{D}_{u+1}$, where $\hat{D}_T = 1$. Choose

$$0 < \varepsilon < \min\left\{\hat{D}_t - \hat{D}_{t-1}, \hat{D}_{u+1} - \hat{D}_u, \min_{s=t,\dots,u}(D_s - \hat{D}_s)\right\}.$$

Let $\hat{D}^+$ and $\hat{D}^-$ be obtained from $\hat{D}$ by replacing $\hat{D}_s$ with $\hat{D}_s + \varepsilon$ and $\hat{D}_s - \varepsilon$, respectively, for $s = t, \dots, u$, and leaving all other coordinates unchanged. Both $\hat{D}^+$ and $\hat{D}^-$ are feasible, distinct, and have average $\hat{D}$, so $\hat{\delta}$ is not extreme.

Thus, every extreme point satisfies the following alternative at each $t < T$: either $\hat{D}_t = \hat{D}_{t-1}$ or $\hat{D}_t = D_t$. Let $R \coloneqq \{t < T\colon \hat{D}_t = D_t\} \cup \{T\}$. If $R = \{r_1 < \cdots < r_m = T\}$ and $r_0 \coloneqq 0$, the alternative implies that $\hat{D}$ is flat between r_{j-1} and $r_j - 1$, and then reaches D_{r_j} at r_j. Therefore, $\hat{\delta}_{r_j} = D_{r_j} - D_{r_{j-1}} = \sum_{s=r_{j-1}+1}^{r_j} \delta_s$, and $\hat{\delta}_t = 0$ for $t \notin R$. This is exactly δ^R.

Conversely, each distinct δ^R is extreme. Indeed, suppose $\delta^R = \alpha\eta + (1-\alpha)\eta'$, where $\alpha \in (0,1)$ and $\eta, \eta' \in \Delta^{\succeq\delta}$. The partial sum of a discount sequence δ is $(\sum_{s=1}^t \delta_s)_t$. Let H, H' be the partial sums of η, η'. If the partial sum of δ^R is flat between $t-1$ and t, then $H_t = H_{t-1}$ and $H'_t = H'_{t-1}$, because both H and H' are nondecreasing. If the partial sum of δ^R equals D_t, then $H_t = H'_t = D_t$, because both H_t and H'_t are weakly below D_t. Applying these two implications date by date yields $H = H'$, and so $\eta = \eta' = \delta^R$. ■

We say that f is $\Delta^{\succeq\delta}$-sufficient for g if, for every $R \subseteq \{1, \dots, T\}$ with $T \in R$,

$$\sum_{t=1}^T \delta^R_t f^t \quad \text{is sufficient for} \quad \sum_{t=1}^T \delta^R_t g^t.$$

Corollary 6 *The following statements are equivalent:*

1. *The experiment f is more valuable than the experiment g in $\hat{\delta}$-discounted problems, for all $\hat{\delta} \in \Delta^{\succeq\delta}$.*
2. *The experiment f is $\Delta^{\succeq\delta}$-sufficient for g.*
3. *For every prior $p \in \Delta(\Theta)$ and every $R \subseteq \{1,\dots,T\}$ with $T \in R$, $\mu_{p,f}^{\delta^R}$ dominates $\mu_{p,g}^{\delta^R}$ in the convex order.*

Proof Let $K := \{\eta \in \Delta\colon f \text{ is } \eta\text{-sufficient for } g\}$. By Proposition 2, K is convex. If Item 2 holds, then K contains every δ^R in (B.6). By Lemma 2, these are exactly the extreme points of $\Delta^{\succeq\delta}$, so $\Delta^{\succeq\delta} \subseteq K$, and, thus, Item 1 via Theorem 1. Conversely, Item 1 and Theorem 1 imply that every $\eta \in \Delta^{\succeq\delta}$ belongs to K, so in particular every δ^R belongs to K, which is Item 2. The equivalence with Item 3 is the usual one. ■

Now assume that $\delta \in \Delta^{\downarrow}$, and let $\Delta^{\succeq\delta,\downarrow} := \Delta^{\succeq\delta} \cap \Delta^{\downarrow}$ denote the class of discount factors that are more patient than δ, but still decreasing.

Let $\mathcal{P} = \{B_1,\dots,B_m\}$ be a partition of $\{1,\dots,T\}$ into consecutive blocks. Write $B_j = \{a_{j-1}+1,\dots,a_j\}$, where $0 = a_0 < a_1 < \cdots < a_m = T$. Let $\delta^{\mathcal{P}} \in \Delta$ be obtained by averaging δ on each block:

$$\delta_t^{\mathcal{P}} := \frac{1}{|B_j|}\sum_{s\in B_j}\delta_s \quad \text{for } t \in B_j. \tag{B.7}$$

Lemma 3 *Assume $\delta \in \Delta^{\downarrow}$. The extreme points of $\Delta^{\succeq\delta,\downarrow}$ are the distinct discount factors $\delta^{\mathcal{P}}$ in (B.7), where $\mathcal{P}$ ranges over all partitions of $\{1,\dots,T\}$ into consecutive blocks.*

Proof Let $D_0 := 0$, $D_t := \sum_{s=1}^{t}\delta_s$, $\hat{D}_0 := 0$, and $\hat{D}_t := \sum_{s=1}^{t}\hat{\delta}_s$. Since $\delta \in \Delta^{\downarrow}$, the sequence D is concave. A discount factor $\hat{\delta}$ belongs to $\Delta^{\succeq\delta,\downarrow}$ if and only if $\hat{D}_0 = D_0 = 0$, $\hat{D}_T = D_T = 1$, $\hat{D}_t \le D_t$ for every $t < T$, and the increments $\hat{D}_t - \hat{D}_{t-1}$ are decreasing and nonnegative.

Fix a consecutive partition $\mathcal{P}$ with endpoints $0 = a_0 < a_1 < \cdots < a_m = T$, and let $D^{\mathcal{P}}$ be the partial sum of $\delta^{\mathcal{P}}$. Then $D^{\mathcal{P}}_{a_j} = D_{a_j}$ for every j, and $D^{\mathcal{P}}$ is linear between consecutive endpoints. Since D is concave, each such line segment lies weakly below D. The slopes of these line segments are the block averages in (B.7). These slopes are decreasing across blocks because δ is decreasing. Consequently, $\delta^{\mathcal{P}} \in \Delta^{\succeq\delta,\downarrow}$.

Each distinct $\delta^{\mathcal{P}}$ is extreme. Suppose that $D^{\mathcal{P}} = \alpha E + (1-\alpha)F$, where $\alpha \in (0,1)$, and where E and F are partial sums of discount factors in $\Delta^{\succeq\delta,\downarrow}$. At each endpoint a_j, both E_{a_j} and

F_{a_j} are weakly below $D_{a_j} = D^{\mathcal{P}}_{a_j}$, and their average equals $D^{\mathcal{P}}_{a_j}$. Hence, $E_{a_j} = F_{a_j} = D_{a_j}$. On each block, E and F are concave sequences with the same endpoints as the line segment $D^{\mathcal{P}}$, so both lie weakly above that line segment. Since their average is exactly the line segment, both must equal $D^{\mathcal{P}}$ on the whole block. Therefore, $E = F = D^{\mathcal{P}}$.

It remains to show that there are no other extreme points. Let $\hat{D}$ be feasible, and let $0 = a_0 < a_1 < \cdots < a_m = T$ be the dates at which $\hat{D}$ touches D (i.e., for which $\hat{D}_t = D_t$), together with 0 and T. Call these dates *touching dates*. Consider two consecutive such dates $a < b$. Suppose that $\hat{D}$ is not linear on $\{a, \ldots, b\}$. Let L be the line segment joining (a, D_a) and (b, D_b). Since $\hat{D}$ is concave and has the same endpoints as L, $\hat{D}_t \geq L_t$ for $a < t < b$, with strict inequality for at least one such t. Also, since a and b are consecutive touching dates, $\hat{D}_t < D_t$ for every $a < t < b$.

Write $z_s := \hat{D}_s - \hat{D}_{s-1}$ for $s = a+1, \ldots, b$, and let $\bar{z} := (D_b - D_a)/(b-a)$. For $a \leq t \leq b$, set $H_t := \hat{D}_t - L_t$, and set $H_t := 0$ outside this interval. For small $\varepsilon > 0$, define $D^+ := \hat{D} + \varepsilon H$ and $D^- := \hat{D} - \varepsilon H$. The two sequences have the same endpoints as $\hat{D}$, and their increments on $s = a+1, \ldots, b$ are $z_s + \varepsilon(z_s - \bar{z})$ and $z_s - \varepsilon(z_s - \bar{z})$, respectively. These increments are decreasing for small ε, because $(z_s)_s$ is decreasing. The upper bounds remain satisfied for small ε, because the bounds are strict at every interior date $a < t < b$.

We only need to check the two neighboring inequalities. At the left boundary, the first modified increment of D^+ can increase. If $a > 0$ and the inequality before the block were binding, then $z_a = z_{a+1}$. Since $\hat{D}_a = D_a$ and $\hat{D}_{a-1} \leq D_{a-1}$, this implies $z_{a+1} = z_a \geq D_a - D_{a-1}$. Since δ is decreasing, $D_a - D_{a-1} \geq D_{a+1} - D_a$. Thus, $\hat{D}_{a+1} = D_a + z_{a+1} \geq D_{a+1}$, contradicting the fact that $a+1$ is not a touching date. Thus, the left neighboring inequality is slack whenever it is needed. At the right boundary, the last modified increment of D^+ can decrease. If $b < T$ and the inequality after the block were binding, then $z_b = z_{b+1}$. Since $\hat{D}_b = D_b$, $\hat{D}_{b-1} \leq D_{b-1}$, and $\hat{D}_{b+1} \leq D_{b+1}$, we have $z_b \geq D_b - D_{b-1}$ and $z_{b+1} \leq D_{b+1} - D_b$. Since δ is decreasing, equality $z_b = z_{b+1}$ forces $\hat{D}_{b-1} = D_{b-1}$, contradicting the fact that $b-1$ is not a touching date. Thus, the right neighboring inequality is also slack whenever it is needed. Finally, if $b = T$, choose ε small enough to keep the last increment nonnegative.

Therefore, D^+ and D^- are distinct feasible partial sums whose average is $\hat{D}$. Hence, $\hat{D}$ is not extreme. It follows that, at an extreme point, $\hat{D}$ must be linear between consecutive touching dates. The associated increments are exactly the block averages in (B.7). ■

We say that f is $\Delta^{\succeq\delta,\downarrow}$-sufficient for g if, for every consecutive partition $\mathcal{P}$,

$$\sum_{t=1}^{T} \delta_t^{\mathcal{P}} f^t \quad \text{is sufficient for} \quad \sum_{t=1}^{T} \delta_t^{\mathcal{P}} g^t.$$

Equivalently, if $\mathcal{P} = \{B_1, \dots, B_m\}$,

$$\sum_{j=1}^{m} \left(\sum_{s \in B_j} \delta_s \right) \left(\frac{1}{|B_j|} \sum_{t \in B_j} f^t \right) \quad \text{is sufficient for} \quad \sum_{j=1}^{m} \left(\sum_{s \in B_j} \delta_s \right) \left(\frac{1}{|B_j|} \sum_{t \in B_j} g^t \right).$$

Corollary 7 *Assume $\delta \in \Delta^{\downarrow}$. The following statements are equivalent:*

1. *The experiment f is more valuable than the experiment g in $\hat{\delta}$-discounted problems, for all $\hat{\delta} \in \Delta^{\succeq\delta,\downarrow}$.*
2. *The experiment f is $\Delta^{\succeq\delta,\downarrow}$-sufficient for g.*
3. *For every prior $p \in \Delta(\Theta)$ and every consecutive partition $\mathcal{P}$, $\mu_{p,f}^{\delta^{\mathcal{P}}}$ dominates $\mu_{p,g}^{\delta^{\mathcal{P}}}$ in the convex order.*

Proof Let $K := \{\eta \in \Delta \colon f \text{ is } \eta\text{-sufficient for } g\}$. By Proposition 2, K is convex. If Item 2 holds, then K contains every $\delta^{\mathcal{P}}$ in (B.7). By Lemma 3, these are exactly the extreme points of $\Delta^{\succeq\delta,\downarrow}$. Since $T < +\infty$, every element of $\Delta^{\succeq\delta,\downarrow}$ is a convex combination of its extreme points. Hence, $\Delta^{\succeq\delta,\downarrow} \subseteq K$, and Theorem 1 gives Item 1. Conversely, Item 1 and Theorem 1 imply that every $\eta \in \Delta^{\succeq\delta,\downarrow}$ belongs to K, so in particular every $\delta^{\mathcal{P}}$ belongs to K, which is Item 2. The equivalence with Item 3 is the usual one. ■

B.5 Sequential Comparison

The comparisons we have made so far are from the ex-ante viewpoint, that is, the decision-maker chooses whether to observe the signals from either f or g at the first stage and cannot revise his choice at a later date. We want to strengthen our comparison to give the decision-maker the opportunity to *permanently* switch from observing the X-signals to observing the Y-signals after any history x^t of X-signals. In other words, we want the decision-maker to prefer observing signals from f not only at the ex-ante stage, but also at all (on-path) histories. The following example illustrates the issues we have to deal with.

Example. There are two periods, two states L(ow) and H(high), and two signals ℓ(ow) and h(igh). The experiments f and g are given by:

$$f_1(\ell|L) = f_1(h|H) = 3/4, f_2(h|h,\theta) = f_2(\ell|\ell,\theta) = 1,$$
$$g_1(\ell|L) = g_1(h|H) = 3/4, g_2(h|h,\theta) = g_2(\ell|\ell,\theta) = 1,$$

with $\theta = L, H$. The two experiments have the same probability law and simply repeat the first-period signal at the second period. As random processes, they may differ, however. If they are independent, after observing the first-period signal from either of them, the decision-maker always has an incentive to switch. By doing so, the decision-maker obtains two independent signals drawn from the same distribution instead of a single one. If the two processes are perfectly correlated, the decision-maker never has an incentive to switch. The example illustrates the need to model the interdependence between the experiments if we want to formalize the incentives of the decision-maker to switch from one experiment to another.

The approach we follow is to consider a joint process (a coupling) over the X- and Y-signals such that marginal over the X- (resp., Y-) signals coincides with the experiment f (resp., g). The joint process allows us to define the "continuation experiments," i.e., the experiments the decision-maker faces when he considers switching from one experiment to another after having observed the history of signals x^t. More precisely, the joint process is given by the collection of kernels $(h_t)_t$, where

$$h_t : \Theta \times X^{t-1} \times Y^{t-1} \to \Delta(X_t \times Y_t).$$

For all $x^t \in X^t$, let $\Theta[x^t] \coloneqq \{\theta \in \Theta : h^t(x^t \mid \theta) > 0\}$ and define the experiments

$$f[x^t] : \Theta[x^t] \to \Delta(X_{t+1} \times \cdots \times X_T),$$
$$g[x^t] : \Theta[x^t] \to \Delta(Y_{t+1} \times \cdots \times Y_T),$$

where

$$f[x^t](x_{t+1}, \ldots, x_T|\theta) \coloneqq \sum_{y^t} h(x_{t+1}, \ldots, x_T|x^t, y^t, \theta) h^t(y^t|x^t, \theta),$$
$$g[x^t](y_{t+1}, \ldots, y_T|\theta) \coloneqq \sum_{y^t} h(y_{t+1}, \ldots, y_T|x^t, y^t, \theta) h^t(y^t|x^t, \theta),$$

Thus, after observing the sequence of signals x^t, the decision-maker can either continue observing the signals from $f[x^t]$ or switch to observing the signals from $g[x^t]$. Note that f and g correspond to $f[\emptyset]$ and $g[\emptyset]$, respectively. We want the decision-maker to find $f[x^t]$ more valuable than $g[x^t]$ at all x^t. In what follows, we denote $\tau[x^t]$ and $\sigma[x^t]$ the strategies of the decision problem starting at period $t+1$, adapted to $g[x^t]$ and $f[x^t]$, respectively.

Definition 5 *Let δ be the discount factor. The experiment f is* sequentially most valuable *for h in discounted problems, if for all decision problems (A,u), for all histories of signals x^t (including the empty history), for all strategies $\tau[x^t]$, there exists a strategy $\sigma[x^t]$ such that*

$$\mathbb{E}_{\theta,g[x^t],\tau[x^t]}\Big[\sum_{t'\geq t+1}\delta_{t'}u(\boldsymbol{a}_{t'},\theta)\Big]\leq\mathbb{E}_{\theta,f[x^t],\sigma[x^t]}\Big[\sum_{t'\geq t+1}\delta_{t'}u(\boldsymbol{a}_{t'},\theta)\Big]$$

for all θ such that $h^t(x^t|\theta)>0$.

The definition states that the experiment f is sequentially most valuable if the decision-maker not only prefers observing the X-signals at the empty history, but also at any (on-path) history x^t. Several modifications of this definition are possible. For instance, we may want f to be sequentially most valuable not only for a fixed h, but also for all joint processes h' such that the marginal over the X-signals (resp., Y-signals) is f (resp., g). As another instance, we may want to compare the experiments not only at on-path histories, but also at off-path histories. As yet another, we may want to give the decision-maker the opportunity to freely choose the signal to observe at each history. We view our definition as a natural starting point for the sequential comparison of experiments.

Given the similarity between the ex-ante comparison of experiments (Definition 2) and their sequential comparison (Definition 5), it is natural to expect their characterizations to be similar. As Theorem 8 demonstrates, this is indeed the case. To state the theorem, we need to introduce some notation.

For all x^t, we let

$$\delta[t]:=\left(\frac{\delta_{t+1}}{\sum_{t'\geq t+1}\delta_{t'}},\frac{\delta_{t+2}}{\sum_{t'\geq t+1}\delta_{t'}},\dots\right),$$

whenever $\sum_{t'=t+1}^{T}\delta_{t'}>0$. We say that $f[x^t]$ is $\delta[t]$-sufficient for $g[x^t]$ when the mixture experiment $\sum_{t'=t+1}^{T}\delta_{t'}[t]f^{t'}[x^t]$ is Blackwell sufficient for the mixture experiment

$\sum_{t'=t+1}^{T} \delta_{t'}[t] g^{t'}[x^t]$. (By convention, if $\sum_{t'=t+1}^{T} \delta_{t'} = 0$, then the sufficiency condition is automatically satisfied.)

In words, $\Theta[x^t]$ is the set of states, which have not been excluded yet, and we compare the experiments $f[x^t]$ and $g[x^t]$ only for these states. The discount factor $\delta[t]$ is the normalization of δ to the periods $t+1$ to T.

Theorem 8 *The experiment f is sequentially most valuable for h if, and only if, $f[x^t]$ is $\delta[t]$-sufficient for $g[x^t]$ for all x^t such that $\Theta[x^t] \neq \emptyset$.*

We conclude with a discussion of the recent work of Brooks et al. (2024) on the comparison of experiments. These authors study when the decision-maker prefers observing the X-signals over the Y-signals when he has already observed some Z-signals. They adopt a different notion of comparison than ours in that they require the decision-maker to prefer observing the X-signals, regardless of the Z-signals (and, therefore, of their correlations with the X- and Y-signals).[36] We can view their model as a two-period problem, where $h_1 : \Theta \to \Delta(X_1 \times Y_1)$, $X_1 = Y_1 = Z$, $h_1(x_1, y_1|\theta) > 0$ only if $x_1 = y_1$, $h_2 : \Theta \times X_1 \times Y_1 \to \Delta(X_2 \times Y_2)$, and $\delta = (0, 1)$. In this two-period problem, the experiment f is sequentially most valuable for h if the decision-maker prefers observing the signal x_2 to y_2 in the second period, regardless of the realized signal $z = x_1 = y_1$ in the first period. Since we fix h, this is obviously weaker than the version of Brooks et al. (2024), but the spirit is similar.

C Dominance in controlled δ-discounted problems

We now establish by example that dominance in δ-discounted problems for every $\delta \in \Delta$ does not imply the equivalent conditions of Theorem 7.

Let

$$T = 2 \quad \text{and} \quad \Theta = X_1 = X_2 = Y_1 = Y_2 = K_1 = \{0, 1\},$$

and let K_0 be a singleton. The first-period experiments are

	0	1
$f_1(\cdot \mid 0)$	$\frac{3}{4}$	$\frac{1}{4}$
$f_1(\cdot \mid 1)$	$\frac{1}{4}$	$\frac{3}{4}$

and

	0	1
$g_1(\cdot \mid 0)$	$\frac{1}{2}$	$\frac{1}{2}$
$g_1(\cdot \mid 1)$	$\frac{1}{2}$	$\frac{1}{2}$

(C.1)

[36] Brooks et al. (2024) model experiments differently, but this is not essential for our discussion.

The second-period experiments are

(x_1, k_1)	$f_2(0 \mid x_1, k_1, 0)$	$f_2(0 \mid x_1, k_1, 1)$
$(0,0)$	$\frac{3}{4}$	1
$(0,1)$	$\frac{1}{2}$	$\frac{3}{4}$
$(1,0)$	1	$\frac{1}{2}$
$(1,1)$	1	$\frac{1}{4}$

(y_1, k_1)	$g_2(0 \mid y_1, k_1, 0)$	$g_2(0 \mid y_1, k_1, 1)$
$(0,0)$	$\frac{1}{2}$	$\frac{3}{4}$
$(0,1)$	0	$\frac{3}{4}$
$(1,0)$	$\frac{1}{2}$	$\frac{1}{2}$
$(1,1)$	0	0

(C.2)

Proposition 9 *For the controlled experiments f and g in (C.1) and (C.2):*

(i) for every discount factor $\delta = (d, 1-d)$, with $d \in [0,1]$, f is more valuable than g in every δ-discounted problem;

(ii) there is a separable problem and a strategy under g for which no strategy under f yields at least as high an expected payoff in each state.

Consequently, dominance in δ-discounted problems for every $\delta \in \Delta$ does not imply the equivalent conditions of Theorem 7.

First we construct a separable problem that favors g.

Proof of part (ii) of Proposition 9 Let $A_1 = A_2 = \{0,1\}$, and let the first control rule be $k_1 = 1 - a_1$. The payoffs are

	$u_1(a,0)$	$u_1(a,1)$	$u_2(a,0)$	$u_2(a,1)$
$a=0$	$-\frac{9}{32}$	$\frac{3}{32}$	0	0
$a=1$	0	0	-4	9

The decision-maker's (DM's) payoff criterion is $u_1(a_1,\theta)+u_2(a_2,\theta)$. Each state occurs with probability $1/2$.

Under f, each signal occurs with probability $1/2$, with posterior probabilities of state 1 equal to $1/4$ after $x_1 = 0$ and $3/4$ after $x_1 = 1$. Let r_x be the probability the DM chooses $a_1 = 0$ after $x_1 = x$. The expected first-period payoff is $-3r_0/32$. After $x_1 = 0$, the highest expected payoff the DM can earn in period 2 is $3/16$ under control 1 and zero under control 0. After $x_1 = 1$, action 1 is optimal in period 2 under either control, regardless of the second signal. Thus, before observing the second signal, the DM expects to earn $-4(1/4) + 9(3/4) = 23/4$. Since each first signal occurs with probability $1/2$, the highest

expected payoff the DM can earn in period 2, given the first-period strategy, is

$$\frac{1}{2}\left(r_0\frac{3}{16}+(1-r_0)\cdot 0\right)+\frac{1}{2}\left(\frac{23}{4}\right)=\frac{23}{8}+\frac{3r_0}{32}.$$

The optimal expected total payoff under f is, therefore,

$$-\frac{3r_0}{32}+\frac{23}{8}+\frac{3r_0}{32}=\frac{184}{64}.$$

Under g, let τ choose $a_1=y_1$ and choose $a_2=1$ exactly when $y_1=1$ or $y_2=0$. Its expected first-period payoff is $-3/64$. Its probabilities of choosing $a_2=1$ are $1/2$ in state 0 and $7/8$ in state 1. Hence, its expected total payoff is

$$-\frac{3}{64}+\frac{1}{2}\left[-4\left(\frac{1}{2}\right)+9\left(\frac{7}{8}\right)\right]=\frac{185}{64}>\frac{184}{64}. \tag{C.3}$$

A strategy under f that attained at least the payoff of τ in each state would also do so under the uniform prior, contradicting (C.3). ■

Next, consider the ten-action set

$$B:=\{b_0,b_1,c_{000},c_{001},c_{010},c_{011},c_{100},c_{101},c_{110},c_{111}\}.$$

There is one action b_y for each first signal y, and one action c_{yzk} for each combination of first signal y, second signal z, and control k. All ten actions are available in both periods. For $a\in B$, let $\mathbf{e}_a$ denote the point mass on a. Fix $d\in[0,1]$ and a control rule $\xi_1\colon B\to\Delta(K_1)$. Under g, choose $a_1=b_{y_1}$ and $a_2=c_{y_1y_2k_1}$, and let G_t denote the resulting state-dependent action distribution at date t.

Because K_0 is a singleton, B identifies the union of the date-one and date-two signal-control histories in the fixed-discounting criterion of Theorem 6. That criterion requires kernels $\gamma_1\colon X_1\to\Delta(B)$ and $\gamma_2\colon X_1\times K_1\times X_2\to\Delta(B)$ such that

$$d\mathbb{P}_{\theta,f,\xi_1,\gamma}(\boldsymbol{a}_1=\ell)+(1-d)\mathbb{P}_{\theta,f,\xi_1,\gamma}(\boldsymbol{a}_2=\ell)=dG_{1,\theta}(\ell)+(1-d)G_{2,\theta}(\ell) \tag{C.4}$$

for every $\theta\in\Theta$ and $\ell\in B$.

We construct two strategies under f. Write L_t and E_t for their respective state-dependent action distributions at date t.

Lemma 4 *For every ξ_1, the two strategies can be chosen so that*

$$L_2 = G_2, \qquad E_1 = G_1, \qquad \textit{and} \qquad E_2 - G_2 = -\frac{1}{4}(L_1 - G_1). \tag{C.5}$$

Proof Set $\alpha \coloneqq \xi_1(1 \mid b_0)$ and $\beta \coloneqq \xi_1(1 \mid b_1)$, and define the distributions over actions

$$D \coloneqq (1-\alpha)\mathbf{e}_{c_{010}} + \alpha \mathbf{e}_{c_{011}} \qquad \text{and} \qquad H \coloneqq \frac{1-\beta}{2}(\mathbf{e}_{c_{100}} + \mathbf{e}_{c_{110}}) + \beta \mathbf{e}_{c_{111}}.$$

The first strategy uses the date-one rule

$$s^{\mathrm{L}}(\cdot \mid 0) \coloneqq \frac{2}{3}\mathbf{e}_{b_0} + \frac{1}{3}D \qquad \text{and} \qquad s^{\mathrm{L}}(\cdot \mid 1) \coloneqq H. \tag{C.6}$$

The second strategy uses $s^{\mathrm{E}}(\cdot \mid x) \coloneqq \frac{1}{2}\mathbf{e}_{b_0} + \frac{1}{2}\mathbf{e}_{b_1}$ for both $x \in X_1$. Since $G_{1,0} = G_{1,1} = \frac{1}{2}\mathbf{e}_{b_0} + \frac{1}{2}\mathbf{e}_{b_1}$, this ensures $E_1 = G_1$. Define the desired date-two distribution for the second strategy by

$$R \coloneqq G_2 + \frac{1}{4}(G_1 - L_1). \tag{C.7}$$

Substituting (C.1), (C.2), and (C.6) into (C.7) produces

$$\begin{aligned} R_0 &= \frac{1}{8}\mathbf{e}_{b_1} + \frac{1-\alpha}{4}\mathbf{e}_{c_{000}} + \frac{3(1-\alpha)}{16}\mathbf{e}_{c_{010}} + \frac{7\alpha}{16}\mathbf{e}_{c_{011}} + \frac{7}{16}H, \quad \text{and} \\ R_1 &= \frac{1}{12}\mathbf{e}_{b_0} + \frac{1}{8}\mathbf{e}_{b_1} + \frac{3}{8}\big((1-\alpha)\mathbf{e}_{c_{000}} + \alpha \mathbf{e}_{c_{001}}\big) + \frac{5}{48}D + \frac{5}{16}H. \end{aligned} \tag{C.8}$$

Note that all coefficients are nonnegative, and each row sums to one so R is a valid state-dependent probability distribution.

Let $F(\rho_0, \rho_1)$ denote the state-dependent distribution of (x_1, k_1, x_2) under f when control 1 is chosen with probability ρ_x after $x_1 = x$. Under s^{L},

$$\rho_0 = \frac{2}{3}\alpha + \frac{1}{3}\big((1-\alpha)\xi_1(1 \mid c_{010}) + \alpha \xi_1(1 \mid c_{011})\big) \geq \frac{2}{3}\alpha. \tag{C.9}$$

Under s^{E}, the DM chooses b_0 and b_1 with probability $1/2$ each, regardless of the first signal. Accordingly, after either first signal, the realized control is 1 with probability $r \coloneqq (\alpha + \beta)/2$.

For a probability matrix M, with one row for each state and columns indexed by z, define

$$C_M(c) \coloneqq \frac{1}{2}\sum_z \big((1-c)M_1(z) - cM_0(z)\big)_+, \qquad \text{for } c \in [0,1],$$

where $v_+ := \max\{v, 0\}$. This is the "call function" of the posterior probability of state 1 under the uniform prior. The binary-state form of Blackwell's theorem states that N is a garbling of M if and only if $C_M(c) \geq C_N(c)$ for every $c \in [0, 1]$.

To verify $C_M(c) \geq C_N(c)$ for every $c \in [0, 1]$, it is enough to check the posterior values that occur with positive probability under M, together with 0 and 1. To see this, let $a < b$ be consecutive points in this set. Each term $(q - c)_+$ in the definition of C_M changes slope only at $c = q$. Since no posterior value under M lies strictly between a and b, the graph of C_M is a straight line on this interval. The convexity of C_N, meanwhile, ensures that its graph lies at or below the straight line joining its values at the endpoints. Thus, if $C_N(a) \leq C_M(a)$ and $C_N(b) \leq C_M(b)$, then, for every $t \in [0, 1]$,

$$C_N((1-t)a + tb) \leq (1-t)C_N(a) + tC_N(b) \leq (1-t)C_M(a) + tC_M(b) = C_M((1-t)a + tb).$$

Hence, the desired inequality holds throughout $[a, b]$. Including 0 and 1 ensures that these intervals cover all of $[0, 1]$.

Every posterior under $F(\rho_0, \rho_1)$ belongs to $\{0, 1/7, 4/13, 1/3, 3/7, 3/5, 1\}$. Using (C.1), (C.2), and (C.8), we obtain

c	$C_{F(\rho_0,\rho_1)}(c) - C_{G_2}(c)$	$C_{F(r,r)}(c) - C_R(c)$
$\frac{1}{7}$	$\frac{3(1-\rho_0)}{224}$	$\frac{3(2-\alpha-\beta)}{448}$
$\frac{4}{13}$	$\frac{12+3\rho_0-14\alpha}{416}$	$\frac{24-23\alpha+3\beta}{832}$
$\frac{1}{3}$	$\frac{1-\alpha}{24}$	$\frac{12-11\alpha}{288}$
$\frac{3}{7}$	$\frac{4-3\alpha}{56}$	$\frac{25-18\alpha}{336}$
$\frac{3}{5}$	$\frac{3(2+\rho_1-2\alpha)}{80}$	$\frac{28-27\alpha+9\beta}{480}$

For the middle entry at $c = 4/13$, (C.9) implies

$$12 + 3\rho_0 - 14\alpha \geq 12(1 - \alpha) \geq 0.$$

Every other entry is nonnegative because $\alpha, \beta, \rho_0, \rho_1 \in [0, 1]$. All the call functions equal $1/2$ at $c = 0$ and zero at $c = 1$.

Accordingly, there are kernels $\phi^{\mathrm{L}}, \phi^{\mathrm{E}} \colon X_1 \times K_1 \times X_2 \to \Delta(B)$ that transform $F(\rho_0, \rho_1)$ into G_2 and $F(r, r)$ into R, respectively. Use these kernels at date 2. Then $L_2 = G_2$ and $E_2 = R$, so (C.7) establishes (C.5). ■

Proof of part (i) of Proposition 9 Fix the two strategies in Lemma 4, and let

$$\omega_{\mathrm{L}} := \frac{1-d}{1+3d}, \qquad \text{and} \qquad \omega_{\mathrm{E}} := \frac{4d}{1+3d}.$$

These probabilities sum to one and satisfy $d\omega_{\mathrm{L}} - (1-d)\omega_{\mathrm{E}}/4 = 0$. We combine the strategies into the kernels required by (C.4).

At date 1, use $\gamma_1 := \omega_{\mathrm{L}} s^{\mathrm{L}} + \omega_{\mathrm{E}} s^{\mathrm{E}}$. For $i \in \{\mathrm{L}, \mathrm{E}\}$, let $m_i(x,k) := \sum_{a \in B} s^i(a \mid x)\xi_1(k \mid a)$, and let $m(x,k) := \omega_{\mathrm{L}} m_{\mathrm{L}}(x,k) + \omega_{\mathrm{E}} m_{\mathrm{E}}(x,k)$. At a date-two history (x,k,z) with $m(x,k) > 0$, define

$$\gamma_2(\ell \mid x,k,z) := \frac{\sum_{i \in \{\mathrm{L},\mathrm{E}\}} \omega_i m_i(x,k)\phi^i(\ell \mid x,k,z)}{m(x,k)}. \tag{C.10}$$

For each $i \in \{\mathrm{L}, \mathrm{E}\}$, the probabilities $\phi^i(\ell \mid x,k,z)$ sum to one over $\ell \in B$. Thus, whenever $m(x,k) > 0$,

$$\sum_{\ell \in B} \gamma_2(\ell \mid x,k,z) = \frac{\omega_{\mathrm{L}} m_{\mathrm{L}}(x,k) + \omega_{\mathrm{E}} m_{\mathrm{E}}(x,k)}{m(x,k)} = 1.$$

Each $\gamma_2(\ell \mid x,k,z)$ is also nonnegative. Hence, $\gamma_2(\cdot \mid x,k,z)$ is a probability distribution over the second-period actions. If $m(x,k) = 0$, control k has probability zero after first signal x, in either state and so the DM never observes a history (x,k,z). At each such history, we may define $\gamma_2(\cdot \mid x,k,z)$ to be any probability distribution over B, using the same distribution in both states, which choice does not affect the distribution of second-period actions, because the history has probability zero!

Under γ_1, the probability of observing (x,k,z) in state θ is $f_1(x \mid \theta)m(x,k)f_2(z \mid x,k,\theta)$. To obtain the probability of observing this history and then choosing action ℓ, we multiply by $\gamma_2(\ell \mid x,k,z)$. For $m(x,k) > 0$, substituting (C.10) for this action probability, the product is

$$\begin{aligned} &f_1(x \mid \theta)m(x,k)f_2(z \mid x,k,\theta)\frac{\sum_{i \in \{\mathrm{L},\mathrm{E}\}} \omega_i m_i(x,k)\phi^i(\ell \mid x,k,z)}{m(x,k)} \\ &\quad = f_1(x \mid \theta)f_2(z \mid x,k,\theta) \sum_{i \in \{\mathrm{L},\mathrm{E}\}} \omega_i m_i(x,k)\phi^i(\ell \mid x,k,z). \end{aligned}$$

Summing over (x,k,z), and using the definition of γ_1 at date 1, we obtain

$$\mathbb{P}_{\theta,f,\xi_1,\gamma}(\boldsymbol{a}_1 = \ell) = \omega_{\mathrm{L}} L_{1,\theta}(\ell) + \omega_{\mathrm{E}} E_{1,\theta}(\ell), \quad \text{and} \quad \mathbb{P}_{\theta,f,\xi_1,\gamma}(\boldsymbol{a}_2 = \ell) = \omega_{\mathrm{L}} L_{2,\theta}(\ell) + \omega_{\mathrm{E}} E_{2,\theta}(\ell).$$

Applying (C.5) to these distributions, their discounted sum is

$$\begin{aligned} &d(\omega_{\mathrm{L}} L_1 + \omega_{\mathrm{E}} E_1) + (1-d)(\omega_{\mathrm{L}} L_2 + \omega_{\mathrm{E}} E_2) \\ &\quad = dG_1 + (1-d)G_2 + \left(d\omega_{\mathrm{L}} - \frac{1-d}{4}\omega_{\mathrm{E}} \right)(L_1 - G_1) = dG_1 + (1-d)G_2. \end{aligned}$$

Thus, (C.4) holds. Since d and ξ_1 were arbitrary, the fixed-discounting criterion establishes part (i). ■