

Total Recall

Jakub Steiner*

Colin Stewart[†]

October 15, 2026

Abstract

A decision-maker observes a signal, forgets it, and acts on its garbled record. Because her information sets are not nested, delegating to her future self can be unattractive: for generic memory channels, some interim selves are better off with a blank memory in some decision problems, despite stable preferences. We characterize channels immune to such deviations in every decision problem. Every memory state must either recall the payoff-relevant information perfectly or induce one and the same belief. When interim selves can manipulate the record rather than merely suppress it, only perfect recall and total amnesia survive.

Keywords: Imperfect recall, incentive compatibility, information transmission, commitment, information suppression, value of information.

JEL codes: D83, D82, D91

1 Introduction

Memory—be it institutional or individual—is subject to frictions. Records can typically provide only a coarse summary of available information. Individual memory is subject to random lapses and distortions. In the face of these frictions, the person keeping the record may want to misrepresent information, even when her interests coincide with the final decision-maker’s. In this paper, we identify the memory mechanisms immune to such manipulation in every decision problem.

Consider a hospital radiologist reading a scan for a disease. Besides evidence on the presence and severity of the disease, she observes the scan’s quality. The record-keeping system allows her to

*University of Zurich, CERGE-EI, and CTS.

[†]University of Toronto.

report only that the scan was positive, negative, or nondiagnostic, without distinguishing between high- and low-quality scans. Based on this report and the patient’s comorbidities, a physician at the same hospital then decides whether to intervene or order a follow-up test. Both doctors observe the patient’s comorbidities and share the hospital’s objective for patient care. Depending on the patient’s comorbidity status, the radiologist may prefer to misreport a low-quality negative scan as nondiagnostic. Doing so may trigger a follow-up test that she considers warranted but that the physician would not order after a negative report. What kinds of record-keeping systems would avoid incentives for such misreporting regardless of the patient’s circumstances—that is, in every decision problem?

We propose a model in which a sender and a receiver have identical preferences. The sender observes a signal and reports to the receiver through a process subject to frictions. The receiver updates by Bayes’ rule and chooses an action. A memory channel describes these frictions, mapping each signal to a distribution over messages. Channels can capture deterministic coarsening (as in the example), stochastic noise, or both. For instance, the radiologist may write detailed notes that the referring physician reads with stochastic attention, leading to a nontrivial distribution of his interpretations of them.

We consider three forms of manipulation by the sender: erasure, input manipulation, and output manipulation. Under erasure, the sender chooses whether to let the memory channel operate or suppress the record; under input manipulation, she chooses which signal to report; and under output manipulation, she directly chooses the distribution of channel outputs the receiver observes. In each case, we ask which memory channels leave no type of the sender with an incentive to deviate, whatever the decision problem.

Few channels qualify: when erasure is the sender’s only tool for manipulation, she has an incentive to misreport in some decision problems unless the procedure either perfectly recalls the interim belief or induces the same belief as erasure. Under a regularity condition, robustness to input manipulation requires the procedure either to perfectly recall the interim belief or to return to the prior, with the probability of recall constant across all interim signals. Robustness to output manipulation is more restrictive still: only perfect recall and total amnesia remain.

Our key innovation is to require robustness across all decision problems. This captures a designer who cannot foresee every use of the records generated by the memory channel and therefore wants the sender to have no incentive to deviate, whatever the sender believes about how the records will be used.

Similar issues arise in individual memory, where the sender and receiver are the same person at different times. Our results show that unless the memory channel takes a special form, the individual is better off encoding some information falsely in some decision problem. Behavior that may look

like self-deception to an outsider arises not from a conflict of preferences between the two selves, but from the initial self’s attempt to steer the final self’s action toward what her own information warrants.

2 Literature

A number of papers have studied comparisons of information or beliefs robustly across all decision problems in the presence of frictions. Lehrer (1998) ranks two subjective beliefs evaluated from the perspective of a third belief. As an intermediate step, we extend Lehrer’s result to distributions of beliefs: the sender in our model robustly compares the receiver’s posteriors resulting from truth-telling to those resulting from some deviation. A combination of the robust ranking of the posteriors with Bayes plausibility delivers our sharp characterization of robust memory channels.¹

Morris and Shin (1997) characterize when a misspecified experiment is (robustly) valuable. Whitmeyer (2026) characterizes updating rules that preserve the Blackwell order of experiments (and ensure in particular that an informative experiment weakly improves on no information). Frick, Iijima, and Ishii (2024) compare the welfare impact of updating biases or misspecified beliefs about experiments in static and dynamic learning problems. In our model, beliefs differ from perfect-recall Bayesian posteriors due to imperfect memory instead of misspecification or non-Bayesian updating. Brooks, Frankel, and Kamenica (2024) strengthen the Blackwell comparison by requiring it to hold regardless of background information. Whereas they focus on robust information accumulation, we examine robust bounded recall.

Our study of imperfect memory relates to lossy information compression in rate–distortion theory (Shannon, 1959). Rather than optimize compression under a given distortion criterion, we characterize fixed memory channels that resist interim manipulation in every decision problem. Within a single decision problem, Piccione and Rubinstein (1997) highlight problems of time inconsistency associated with imperfect memory.² Wilson (2014) shows that interim selves (the informed senders in our terminology) do not want to deviate from the ex ante optimal use of a finite set of memory states in a fixed decision problem.³ Dow (1991) studies optimal memory in a search problem. Li

¹Our focus on a Bayesian receiver distinguishes our analysis from memory models in which the decision-maker does not account for selective recall, such as in Mullainathan (2002), Bordalo, Gennaioli, and Shleifer (2020) and Fudenberg, Lanzani, and Strack (2024).

²Kline (2002) studies conditions on the memory structure that ensure time consistency across decision problems. While he focuses on ex ante optimal strategies in extensive games, we study the robustness of a fixed memory channel.

³Sims (2003), read as a model of forgetting, and Azeredo da Silveira, Sung, and Woodford (2024) optimize the memory structure at the ex ante stage for a given environment. They do not examine interim incentive compatibility. Wilson’s argument, which rests on a cap on the number of memory states, does not apply directly to their information costs, which are defined on the joint distribution of state and record.

and Majumdar (2005) show that selective erasure of memory can do better than always recalling the last signal in a two-armed bandit problem.

Manipulation of memory has been shown to be beneficial in models of decision-making with self-control problems or belief-dependent utility. Bénabou and Tirole (2002) and Bénabou and Tirole (2004) study manipulation and design of memory to manage self-control problems. Gottlieb (2014) and Hagenbach and Koessler (2022) consider self-deception by a decision-maker with motivated beliefs using erasure of signals along the lines of our erasure incentive compatibility constraint. Unlike in these papers, the decision-maker in our model has stable preferences and information has purely instrumental value.

Our model can be viewed as one of communication through an imperfect channel between a sender and receiver who share common interests. Jäger, Metzger, and Riedel (2011) consider coarse cheap talk messages, which is analogous to input incentive compatibility with a deterministic memory channel in our framework. Crémer, Garicano, and Prat (2007) apply coarse cheap talk to the study of organizational structure. In Blume and Board (2013), the receiver is uncertain about the set of messages available to the sender. While these papers focus on equilibrium in a given decision problem, we ask which communication channels are immune to various kinds of manipulation across decision problems.

Perez-Richet and Skreta (2022) share with our paper the feature that the sender can manipulate the input into a noisy experiment. Unlike us, they focus on the design of the experiment in a particular environment with partially conflicting interests and costly falsification.

3 Imperfect Memory

We consider a decision-maker—an organization or an individual—who may forget some of their initial information before choice.

The state ω takes values in a fixed finite set Ω . A *decision problem* $d = (\mathcal{A}, u)$ specifies a finite action set $\mathcal{A}$ and a utility function $u : \mathcal{A} \times \Omega \rightarrow \mathbb{R}$; the collection of all decision problems is $\mathcal{D}$. We extend the utility function to allow beliefs $\tau \in \Delta(\Omega)$ as arguments: we write $u(a, \tau) = \mathbb{E}_{\omega \sim \tau} u(a, \omega)$.

There are two periods. In period 1, the decision-maker observes the decision problem d and a realization x of a random signal X with a fixed finite support $\mathcal{X}$, but not the state ω . The joint distribution π of the state-signal pairs (ω, x) has full support over the signals. We write π_Ω for the state marginal (the *prior*) and π_X for the signal marginal. We denote the Bayesian belief the decision-maker forms upon observing x by ι_x ; we call it the *interim* belief. We refer to the period-1 decision-maker who has observed x as the *x-sender*.

The decision-maker may forget some of her interim information between the periods. She per-

fectly recalls the decision problem d , but instead of recalling the signal realization x , she observes in period 2 a realization m of a random *memory state* M drawn from the conditional distribution $h(m \mid x)$. Memory has no access to the state ω except through the period-1 observation: conditional on x , the memory state is independent of ω . We call the system of conditional distributions $h = (h(\cdot \mid x))_{x \in \mathcal{X}}$ the *memory channel*, and take $\mathcal{M}$ to be the union of its supports across all x ; $\mathcal{M}$ is finite.

We refer to the period-2 decision-maker as the *receiver*. After observing memory state m , she forms the Bayesian *posterior*

$$\tau_m := \sum_{x \in \mathcal{X}} \frac{\pi_X(x) h(m \mid x)}{\sum_{x'} \pi_X(x') h(m \mid x')} \iota_x,$$

a convex combination of the interim beliefs. Finally, the receiver chooses an a posteriori optimal action

$$a_d^*(\tau_m) \in \arg \max_{a \in \mathcal{A}} u(a, \tau_m),$$

breaking ties by a rule that, for each decision problem, depends only on the belief.

4 Erasure

In our baseline model, we endow the senders with an outside option. Besides the memory states in $\mathcal{M}$, there is a blank memory state $m_\emptyset \notin \mathcal{M}$ that arises only off path: each x -sender may choose erasure, and the receiver, observing $m_\emptyset$, forms an *erasure* posterior $q \in \Delta(\Omega)$. Erasure is therefore detected: the receiver learns that the record was suppressed. The erasure posterior is exogenous in this section; in later sections, the beliefs induced by deviations are endogenous.

4.1 Erasure Incentive Compatibility

The x -sender chooses whether to erase before the memory state is drawn. Proceeding with the memory channel h in a decision problem d yields the value

$$V_x(h; d) := \mathbb{E}_{m \sim h(\cdot \mid x)} u(a_d^*(\tau_m), \iota_x)$$

for the x -sender: the receiver chooses at the posterior τ_m , but the x -sender evaluates that choice at her own belief ι_x . Erasure incentive compatibility requires each x -sender, in every decision problem, to prefer proceeding with h to erasing; the latter would leave the receiver choosing at q .

The robustness requirement varies the action set and utility function while holding π , h , and q fixed.

Definition 1. A memory channel h is erasure incentive compatible (EIC) if

$$V_x(h; d) \geq u(a_d^*(q), \iota_x) \quad \text{for all } x \in \mathcal{X} \text{ and all } d \in \mathcal{D}.$$

EIC imposes an individual rationality constraint for all interim signals. One motivation comes from institutional record-keeping. An institution with two workers fixes a record-generating procedure h in advance. A decision problem d arrives and both workers observe it. An upstream worker (the sender) observes x and is instructed to file a record $m \sim h(\cdot | x)$; she may instead suppress it, leaving a blank record. A downstream worker (the receiver) knows the procedure, sees the record but not x , and acts on τ_m after an on-path record and on q after a blank one. Both workers share the institution's objective, so any incentive to suppress stems from information, not from a conflict of interest. We study an institution that seeks a procedure under which no x -sender prefers to suppress the record, whatever the decision problem.

EIC has no bite under sequential learning. Imperfect memory contrasts sharply with standard sequential learning. If the decision-maker retains x and subsequently observes an additional signal m , her later information refines her interim information. Conditional on any x , observing m is therefore a standard Blackwell experiment, so in every decision problem the x -sender weakly prefers to proceed with the experiment rather than act on any fixed belief. Under imperfect memory, by contrast, m replaces rather than supplements x , and this interim value-of-information argument no longer applies. The distinction is entirely interim. Ex ante, any memory channel h is itself a standard experiment about the payoff state, so every channel is weakly preferred to acting on any fixed belief. Thus, ex ante, failures of the EIC inequality for some x are offset by slack for others. EIC captures precisely the additional requirement that the memory process be desirable at every interim information set.

4.2 Recall or Erase

Erasure incentive compatibility holds only for a small class of memory channels. Let

$$\mathcal{X}_m := \{x \in \mathcal{X} : h(m | x) > 0\}$$

be the set of signal realizations that can generate the memory state m . We call m a *perfect-recall* state if $\tau_m = \iota_x$ for every $x \in \mathcal{X}_m$, and an *erasure* state if $\tau_m = q$. A perfect-recall state pools

only realizations that share an interim belief: it retains all payoff-relevant information, though not necessarily the signal realization itself. An erasure state is an on-path memory state inducing the same belief q that a blank memory would induce off path.⁴

Definition 2. A memory channel h is recall-or-erase if each memory state $m \in \mathcal{M}$ is either a perfect-recall state or an erasure state.

Recall-or-erase constrains partial recall. A recall-or-erase channel may randomize between perfect recall of the interim belief and returning the erasure posterior, but never lands in between.

Proposition 1. A memory channel is EIC if and only if it is recall-or-erase.

The proof of necessity consists of two steps. First, EIC requires every posterior τ_m generated from x to lie on the line segment connecting q and ι_x ; otherwise, the separating-hyperplane theorem yields a binary-action decision problem in which the x -sender strictly prefers erasure. Second, combined with Bayes plausibility, these geometric restrictions imply the recall-or-erase property.

Proof. Sufficiency is immediate: Under a recall-or-erase channel, every posterior reached from x is ι_x or q ; erasing only shifts probability from the former to the latter.

For necessity we first show that

$$h \text{ is EIC} \Rightarrow \tau_m \in \text{co}\{q, \iota_x\} \text{ for all } x \in \mathcal{X} \text{ and all } m \in \text{supp } h(\cdot | x).$$

To see this, suppose some x reaches, with positive probability, a posterior $\tau_m \notin \text{co}\{q, \iota_x\}$. A hyperplane then strictly separates τ_m from $\text{co}\{q, \iota_x\}$, so we can construct a binary-action decision problem with a safe action a^0 that pays zero, and a risky action a^1 such that

$$u(a^1, \tau_m) > 0, \quad u(a^1, q) < 0, \quad \text{and } u(a^1, \iota_x) < 0.$$

The receiver takes the safe action at q and the risky action at τ_m , which generates a negative payoff under the sender's belief ι_x . Erasure therefore gives the x -sender zero, while proceeding with the channel gives her a strictly negative expected payoff.

Second, fix a memory state m . If $\tau_m = q$, then m is an erasure state. Otherwise, notice that the preceding characterization applies to all $x \in \mathcal{X}_m$. Thus, for each $x \in \mathcal{X}_m$, ι_x lies on the ray from q through τ_m , at or beyond τ_m . By Bayes' rule, τ_m is a convex combination of these interim beliefs

⁴The perfect-recall and erasure classes need not be disjoint: m belongs to both when $\iota_x = q$ for every $x \in \mathcal{X}_m$.

with strictly positive weights. Since all of them lie at or beyond τ_m on the same ray, their average can equal τ_m only if every one of them equals τ_m . Thus m is a perfect-recall state.⁵ $\square$

The characterization holds for any erasure posterior q : an EIC channel induces only interim beliefs and, possibly, q . The location of q nonetheless matters for how large the EIC class is. Every on-path posterior must be a convex combination of interim beliefs, so for q outside their hull, every memory state of an EIC channel must be a perfect-recall state. Within the hull, the prior is the natural focal value for the erasure belief: if blank memories arose from trembles independent of x , the posterior after a blank memory would be the prior. Propositions 2 and 3 confirm, from the other direction, the special role of the prior as the erasure posterior: once deviations are endogenized, memory processes in which amnesia triggers a return to the prior belief are more robust to deviations than others.

5 Input Manipulation

Section 4 endowed each sender with a single exogenous outside option: a deterministic erasure posterior. We now consider two alternative forms of memory manipulation. Under *input manipulation*, an x -sender can privately replace the true channel input x with another input x' . Under the stronger *output manipulation*, explored in the next section, the x -sender can privately choose any distribution over the outputs of the memory channel. In both cases, the deviation posteriors are endogenously determined by the channel rather than specified exogenously. Unlike erasure, these deviations are unobserved: the receiver continues to use the posteriors τ_m computed from the original channel h .

The model of Section 4 can be read as verifiable disclosure with noise added before the receiver sees the disclosure: the sender chooses only whether to disclose. Input and output manipulation correspond instead to noisy cheap talk. Under input manipulation, the sender chooses a message and the channel adds noise; under output manipulation, she chooses the distribution of messages directly.

These manipulation opportunities sharply restrict the set of robust memory channels. Subject to a non-collinearity condition on the interim beliefs, robustness to input manipulation requires uniform recall-or-forget: the decision-maker either recalls the payoff-relevant interim information perfectly, with a probability independent of x , or forgets it and returns to the prior.

⁵The argument uses Bayes' rule only through the fact that τ_m is a strictly positive-weight convex combination of the interim beliefs ι_x , $x \in \mathcal{X}_m$. Proposition 1 therefore continues to hold under any generalized updating rule that satisfies this property.

Definition 3. A memory channel h is input incentive compatible (IIC) if

$$V_x(h; d) \geq \mathbb{E}_{m \sim h(\cdot | x')} u(a_d^*(\tau_m), \iota_x) \quad \text{for all } x, x' \in \mathcal{X} \text{ and all decision problems } d \in \mathcal{D}.$$

Input incentive compatibility treats memory as a stochastic channel with inputs from $\mathcal{X}$ and outputs from $\mathcal{M}$. For every decision problem, Definition 3 requires each x -sender to weakly prefer the truthful input to any private deviation to another input.

Definition 4. A memory channel h is a uniform recall-or-forget channel if each m is either a perfect-recall memory state or a state inducing $\tau_m = \pi_\Omega$, and the probability of perfect recall is constant across all x such that $\iota_x \neq \pi_\Omega$.

This definition differs from recall-or-erase in two respects. First, the erasure posterior must be the prior; we refer to this special case as *forgetting*. Second, the probability of perfect recall must now be independent of x .⁶

Proposition 2. Suppose the interim beliefs $\{\iota_x\}_{x \in \mathcal{X}}$ are not all collinear. A memory channel is IIC if and only if it is a uniform recall-or-forget channel.

Non-collinearity is a richness requirement, and it cannot be dropped. To see why, consider a binary interim signal space $\mathcal{X} = \{0, 1\}$, for which the two interim beliefs are necessarily collinear. Every memory channel is then IIC, regardless of the number of payoff or memory states. Let $r_m = \Pr(X = 1 \mid M = m)$ denote the receiver's posterior probability of the high signal. The distribution of r_M induced by communicating the high signal first-order stochastically dominates that induced by communicating the low signal.⁷ Since τ_m is an affine function of r_m , the receiver's chosen action yields a payoff nondecreasing in r_m for the high-signal sender and nonincreasing in r_m for the low-signal sender. Thus, both x -senders weakly prefer truthful communication in every decision problem.

We defer the proof of the proposition to Section 7, where we develop tools that apply to both input and output manipulation. The argument follows the same basic strategy as the proof of Proposition 1: incentive compatibility constrains the geometry of the induced posteriors, and Bayes plausibility turns these constraints into restrictions on the channel.

⁶The definition does not restrict the recall probability when $\iota_x = \pi_\Omega$, because recall and forgetting are then undifferentiated.

⁷Since $\mathcal{X}$ is binary, the distribution of r_M conditional on $X = 1$ dominates that conditional on $X = 0$ in the likelihood-ratio order, which implies first-order stochastic dominance.

6 Output Manipulation

Output incentive compatibility strengthens IIC by letting each x -sender privately deviate to any distribution over $\mathcal{M}$, not only to those that arise from some $x' \in \mathcal{X}$.

Definition 5. A memory channel h is output incentive compatible (OIC) if

$$V_x(h; d) \geq \mathbb{E}_{m \sim \sigma} u(a_d^*(\tau_m), \iota_x) \quad \text{for all } x \in \mathcal{X}, \text{ all } \sigma \in \Delta(\mathcal{M}), \text{ and all } d \in \mathcal{D}.$$

Output incentive compatibility is extremely restrictive: the memory process either perfectly preserves all payoff-relevant interim information or always produces total amnesia and returns to the prior.

Proposition 3. A memory channel is OIC if and only if either every memory state is a perfect-recall state, or every memory state induces the posterior π_Ω .

The three propositions reveal a hierarchy across deviation opportunities. Robustness to a single exogenous opt-out permits only recall-or-erase channels. Under the non-collinearity condition, robustness to input manipulation additionally requires the recall probability to be the same for all informative signals and forgetting to lead back to the prior. Robustness to output manipulation eliminates intermediate channels altogether: only perfect recall and total amnesia remain.

7 Incentive Compatibility and Radial Contraction

This section develops the necessary tools and then proves Propositions 2 and 3. A key step is to extend the analysis from Section 4 to allow the posterior following erasure to be stochastic and its distribution to depend on x . We provide a simple geometric characterization of the memory channels for which every x -sender weakly prefers following the channel to deviating to a specified posterior lottery, in every decision problem. We then apply this characterization to IIC and OIC by considering the posterior lotteries that the sender can induce through input and output manipulation, respectively.

7.1 Stochastic Erasure Posterior

For each $x \in \mathcal{X}$, let $o_x \in \Delta(\Delta(\Omega))$ be an exogenous finite-support distribution over posteriors, and write $o = (o_x)_{x \in \mathcal{X}}$. We generalize erasure incentive compatibility by allowing each x -sender either to proceed with the memory channel h or to deviate to the outside option: the posterior lottery o_x .

Definition 6. A memory channel h is o -EIC if

$$V_x(h; d) \geq \mathbb{E}_{q \sim o_x} u(a_d^*(q), \iota_x) \quad \text{for all } x \in \mathcal{X} \text{ and } d \in \mathcal{D}.$$

Both input and output manipulation problems can be written in terms of o -EIC: Let $\mu_x^h \in \Delta(\Delta(\Omega))$ denote the posterior lottery that h induces given x . A channel h is IIC if and only if it is o^y -EIC for every $y \in \mathcal{X}$, where

$$o_x^y = \mu_y^h \quad \text{for every } x \in \mathcal{X}.$$

For output manipulation, let

$$\mathcal{T}_h := \{\tau_m : m \in \mathcal{M}\}$$

be the set of posteriors generated by h . A channel h is OIC if and only if it is o^τ -EIC for every $\tau \in \mathcal{T}_h$, where, for every x , o_x^τ is the degenerate distribution assigning probability one to τ .⁸

We characterize o -EIC using a condition that we call radial contraction. For two finite-support posterior lotteries $\mu, \nu \in \Delta(\Delta(\Omega))$, we say that μ is a *radial contraction of ν toward ι_x* if there exists a coupling⁹ γ of μ and ν such that

$$\gamma(\tau, \tau') > 0 \implies \tau \in \text{co}\{\iota_x, \tau'\}.$$

Equivalently, μ can be generated by drawing a posterior τ' from ν and then stochastically moving it radially toward ι_x ; see Figure 1.

Proposition 4. A memory channel h is o -EIC if and only if μ_x^h is a radial contraction of o_x toward ι_x , for each $x \in \mathcal{X}$.

The proof in the Appendix extends the separation argument in the proof of Proposition 1. Sufficiency follows because moving the receiver's belief toward ι_x cannot lower the x -sender's payoff when the receiver reoptimizes. For necessity, if radial contraction fails, then along some ray from ι_x , the posterior lottery μ_x^h places more probability sufficiently far from ι_x than does o_x . Separating hyperplanes then yield a decision problem in which the receiver's choices give the sender payoff zero at these distant posteriors and one at every other supported posterior. The sender therefore strictly prefers the deviation lottery.

⁸Because payoffs are linear in the distribution over memory states, checking these deviations to deterministic posteriors is equivalent to checking deviations to all posterior lotteries on $\mathcal{T}_h$.

⁹Recall that a joint distribution γ over the pairs of posteriors (τ, τ') is a coupling of μ and ν if its two marginal distributions coincide with μ and ν , respectively.




Figure 1: Comparison of radial contraction with the Blackwell order.

Comparison with the Blackwell order. In addition to its use in understanding IIC and OIC, *o*-EIC answers a second question: when does every x -sender prefer channel h to channel h' in every decision problem? This is the interim analogue, for memory channels, of Blackwell's comparison of experiments. In this comparison, the receiver knows which channel is used and forms the corresponding Bayesian posteriors.

Corollary 1. *Every x -sender weakly prefers memory channel h to h' in every decision problem if and only if, for each x , μ_x^h is a radial contraction of $\mu_x^{h'}$ toward ι_x .*

To see this, let $o_x := \mu_x^{h'}$ for each $x \in \mathcal{X}$. Then h is *o*-EIC if and only if

$$V_x(h; d) \geq V_x(h'; d) \quad \text{for all } x \in \mathcal{X} \text{ and } d \in \mathcal{D},$$

that is, if and only if every x -sender weakly prefers h to h' in every decision problem; Proposition 4 then gives the corollary. As is well known, Blackwell dominance can be characterized by a coupling condition that requires the posterior lottery generated by the preferred experiment to be a mean-preserving spread of the other. Radial contraction instead requires that, for each x -sender, the posterior lottery induced by h can be obtained by moving the dominated stochastic posterior radially toward the interim belief ι_x . As with the Blackwell comparisons, this is formalized by a coupling requirement: the posterior lottery induced at each x by following h must be obtained by first generating a posterior via the channel h' , and then stochastically shifting it toward ι_x .

7.2 Proofs of Propositions 2 and 3

Proof of Proposition 2. Sufficiency is immediate. For necessity, suppose h is IIC and consider $x, y \in \mathcal{X}$ such that $\iota_x \neq \iota_y$. By Proposition 4, μ_x^h is a radial contraction of μ_y^h toward ι_x and, vice versa, μ_y^h is a radial contraction of μ_x^h toward ι_y . We prove that this implies that the two posterior lotteries μ_x^h and μ_y^h coincide at all posteriors outside the line ℓ that passes through ι_x and ι_y . Additionally, when $\iota_x = \iota_y$, mutual radial contraction toward the same point ι_x directly implies that $\mu_x^h = \mu_y^h$.

Notice that for any posterior τ not on the line ℓ and any $\tau' \neq \tau$ radially farther from ι_x , the Euclidean distance from ℓ to τ' is larger than that from ℓ to τ . The fact that μ_x^h is a radial contraction of μ_y^h toward ι_x therefore implies that the distribution of distances of the posteriors from ℓ under μ_y^h weakly first-order stochastically dominates that under μ_x^h . Since the same is true when μ_x^h and μ_y^h are reversed, the two distributions of distances must coincide. Since radial movement away from ι_x increases the distance from ℓ , this in turn implies that μ_x^h and μ_y^h coincide at posteriors not on ℓ .

Consider a supported posterior τ such that $\tau \neq \iota_x$ for all $x \in \mathcal{X}$. We argue that then $\tau = \pi_\Omega$. For this, we show that $\mu_x^h(\tau) = \mu_y^h(\tau)$ for all $x, y \in \mathcal{X}$. This equality follows immediately from the previous paragraph when τ is not on the line $\ell_{x,y}$ passing through ι_x and ι_y . The equality also holds when $\tau \in \ell_{x,y}$ because, by the non-collinearity assumption, we can find $z \in \mathcal{X}$ such that $\tau \notin \ell_{x,z}$ and $\tau \notin \ell_{y,z}$, and then $\mu_x^h(\tau) = \mu_z^h(\tau) = \mu_y^h(\tau)$. Since $\mu_x^h(\tau)$ is constant across all $x \in \mathcal{X}$, Bayes' rationality implies that $\tau = \pi_\Omega$, as needed.

Next consider a supported posterior τ such that $\tau = \iota_x$ for some $x \in \mathcal{X}$. The same argument implies that $\mu_y^h(\tau)$ is constant across all y such that $\iota_y \neq \iota_x$. When $\iota_x \neq \pi_\Omega$, non-collinearity and Bayes' rationality imply that the posterior $\tau = \iota_x$ is reached with zero probability from any y such that $\iota_y \neq \iota_x$.

Thus, at most two posteriors, π_Ω and ι_x , are reached from each x such that $\iota_x \neq \pi_\Omega$. Since π_Ω is reached with a constant probability across all such x , the probability of perfect recall is also constant. (When $\iota_x = \pi_\Omega$, the posterior equals π_Ω almost surely.) $\square$

Proof of Proposition 3. Sufficiency is immediate. For necessity, suppose that h is OIC. Recall that $\mathcal{T}_h$ is the set of posteriors induced by h . For each $\tau \in \mathcal{T}_h$, OIC implies that h is o^τ -EIC.

If $\mathcal{T}_h$ is a singleton, its unique posterior equals π_Ω by Bayes plausibility, and the result follows. Otherwise, fix a memory state m and choose $\tau \in \mathcal{T}_h$ with $\tau \neq \tau_m$. By Proposition 4,

$$\tau_m \in \text{co}\{\iota_x, \tau\} \quad \text{for every } x \in \mathcal{X}_m.$$

Thus, every ι_x , $x \in \mathcal{X}_m$, lies on the ray from τ through τ_m , weakly beyond τ_m . By Bayes' rule, τ_m is

a convex combination of these interim beliefs with strictly positive weight on each. This is possible only if $\tau_m = \iota_x$ for every $x \in \mathcal{X}_m$, so m is a perfect-recall state. Since m was arbitrary, every memory state is a perfect-recall state. $\square$

8 Summary

Our results sharply restrict the memory channels and record-keeping procedures that are robust to interim manipulation across all decision problems. For generic channels and procedures, profitable interim deviations exist in some decision problems, despite common interests, stable preferences, and Bayesian updating. Thus, common interests and Bayesian rationality do not generally suffice to sustain the prescribed use of an imperfect memory channel across all decision problems. A natural next step, which we leave for future research, is to study the design of memory channels and record-keeping procedures while accounting for their task-dependent equilibrium use.

A Proof of Proposition 4

There is a useful equivalent geometric formulation of radial contraction. Fix $x \in \mathcal{X}$ and define the radial order of beliefs, centered at ι_x , by

$$\tau \leq_x \tau' \iff \tau \in \text{co}\{\iota_x, \tau'\}.$$

Thus, $\tau \leq_x \tau'$ means that τ is a radial movement of τ' toward ι_x . For each belief τ , define the *outward radial tail* beginning at τ by

$$R_x(\tau) := \{\tau' \in \Delta(\Omega) : \tau \leq_x \tau'\}.$$

We say that the posterior lottery μ_x^h is *raywise dominated* by o_x if

$$\mu_x^h(R_x(\tau)) \leq o_x(R_x(\tau)) \quad \text{for every } \tau \neq \iota_x. \quad (\text{RD})$$

Condition (RD) says that, on each ray originating at ι_x , the deviation lottery o_x places weakly more mass than the posterior lottery induced by h in every outward tail.

Lemma 1. *The posterior lottery μ_x^h is a radial contraction of o_x toward ι_x if and only if it is raywise dominated by o_x .*

Proof. Note that a set $S \subset \Delta(\Omega)$ of beliefs is increasing with respect to $\preceq_x$ if and only if it is a union of outward radial tails.¹⁰ Since any two outward radial tails are either disjoint or nested, it follows that (RD) holds if and only if $\mu_x^h(S) \leq o_x(S)$ for all increasing sets S . Strassen's theorem therefore implies the equivalence of (RD) and radial contraction. $\square$

Proof of Proposition 4. Sufficiency. Fix x, d and a coupling γ witnessing that μ_x^h is a radial contraction of o_x . For any coupled pair (τ, τ') , let $a = a_d^*(\tau)$ and $a' = a_d^*(\tau')$. If $\tau = \tau'$, the tie-breaking rule gives $a = a'$. Otherwise, write $\tau = (1 - \lambda)\tau' + \lambda\iota_x$ for some $\lambda \in (0, 1]$. Since a' is optimal at τ' , strict preference for a' over a at ι_x would imply strict preference at τ , contradicting optimality of a . Thus $u(a, \iota_x) \geq u(a', \iota_x)$. Averaging over γ , whose marginals are μ_x^h and o_x , yields the o -EIC inequality.

Necessity. By Lemma 1, it suffices to show that o -EIC implies (RD) for every $x \in \mathcal{X}$. Fix $x \in \mathcal{X}$ and suppose that (RD) fails. Then there is an outward radial tail $R_x(\tau)$ to which the posterior lottery μ_x^h assigns strictly more probability than the posterior lottery o_x . We construct a decision problem that makes landing in this tail strictly undesirable for the x -sender.

Consider the finitely many posterior beliefs in the supports of μ_x^h and o_x . For any posterior s outside this tail $R_x(\tau)$, the segment joining ι_x to s is disjoint from the tail $R_x(\tau)$. Let a^0 be a safe action that pays zero in every state. The separating-hyperplane theorem allows us, for each posterior s outside the tail, to construct an action a^s that is preferred to a^0 at s and ι_x , while a^0 is preferred at every posterior inside the tail. We normalize these actions a^s so that, when evaluated at ι_x , they all yield the payoff 1.¹¹ Combine all actions a^s and the action a^0 into one decision problem. In this problem, the receiver chooses a^0 at every posterior in the tail, which pays 0 to the x -sender. At every supported posterior outside the tail, the receiver chooses one of the actions a^s , all of which pay 1 to the x -sender. Because μ_x^h places strictly more probability on the tail than o_x , the x -sender strictly prefers the deviation lottery. This contradicts o -EIC as needed. $\square$

References

- Azeredo da Silveira, R., Y. Sung, and M. Woodford (2024). Optimally imprecise memory and biased forecasts. *American Economic Review* 114(10), 3075–3118.
- Bénabou, R. and J. Tirole (2002). Self-confidence and personal motivation. *Quarterly Journal of Economics* 117(3), 871–915.

¹⁰Recall that a set is *increasing* with respect to $\preceq_x$ if $\tau \in S$ implies $\tau' \in S$ whenever $\tau \preceq_x \tau'$.

¹¹For this normalization, take any such action a^s and divide by the payoff it yields at ι_x . This is a division by a positive constant, so it does not affect the separation.

- Bénabou, R. and J. Tirole (2004). Willpower and personal rules. *Journal of Political Economy* 112(4), 848–886.
- Blume, A. and O. Board (2013). Language barriers. *Econometrica* 81(2), 781–812.
- Bordalo, P., N. Gennaioli, and A. Shleifer (2020). Memory, attention, and choice. *Quarterly Journal of Economics* 135(3), 1399–1442.
- Brooks, B., A. Frankel, and E. Kamenica (2024). Comparisons of signals. *American Economic Review* 114(9), 2981–3006.
- Crémer, J., L. Garicano, and A. Prat (2007). Language and the theory of the firm. *Quarterly Journal of Economics* 122(1), 373–407.
- Dow, J. (1991). Search decisions with limited memory. *Review of Economic Studies* 58(1), 1–14.
- Frick, M., R. Iijima, and Y. Ishii (2024). Welfare comparisons for biased learning. *American Economic Review* 114(6), 1612–1649.
- Fudenberg, D., G. Lanzani, and P. Strack (2024). Selective-memory equilibrium. *Journal of Political Economy* 132(12), 3978–4020.
- Gottlieb, D. (2014). Imperfect memory and choice under risk. *Games and Economic Behavior* 85, 127–158.
- Hagenbach, J. and F. Koessler (2022). Selective memory of a psychological agent. *European Economic Review* 142, 104012.
- Jäger, G., L. P. Metzger, and F. Riedel (2011). Voronoi languages: Equilibria in cheap-talk games with high-dimensional types and few signals. *Games and Economic Behavior* 73(2), 517–537.
- Kline, J. J. (2002). Minimum memory for equivalence between ex ante optimality and time-consistency. *Games and Economic Behavior* 38(2), 278–305.
- Lehrer, E. (1998). Comparison of experts. *Journal of Mathematical Economics* 30(2), 207–214.
- Li, H. and S. Majumdar (2005). Dynamic decisions with short-term memories. Technical report, Working paper.
- Morris, S. and H. S. Shin (1997). The rationality and efficacy of decisions under uncertainty and the value of an experiment. *Economic Theory* 9(2), 309–324.

- Mullainathan, S. (2002). A memory-based model of bounded rationality. *The Quarterly Journal of Economics* 117(3), 735–774.
- Perez-Richet, E. and V. Skreta (2022). Test design under falsification. *Econometrica* 90(3), 1109–1142.
- Piccione, M. and A. Rubinstein (1997). On the interpretation of decision problems with imperfect recall. *Games and Economic Behavior* 20(1), 3–24.
- Shannon, C. E. (1959). Coding theorems for a discrete source with a fidelity criterion. *IRE National Convention Record* 7(Part 4), 142–163.
- Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics* 50(3), 665–690.
- Whitmeyer, M. (2026). Blackwell-monotone updating rules. *Journal of Political Economy* 134(8), 000–000.
- Wilson, A. (2014). Bounded memory and biases in information processing. *Econometrica* 82(6), 2257–2294.