

Rational Inattention via Ignorance Equivalence

Michèle Müller-Itten*, Roc Armenter[†], Zachary R. Stangebye[‡]

February 20, 2023

Abstract

We introduce the concept of the *ignorance equivalent* to effectively summarize the payoff possibilities of an agent who is facing a finite Rational Inattention (RI) problem. The ignorance equivalent is a unique fictitious action that does not distort learning incentives when added to the agent's menu but also makes ignorance optimal. In doing so, it allows us to restate the RI problem as a choice over a richer menu without learning. The approach provides new insights for menu expansion, the formation of consideration sets, the trading of information, and belief elicitation. We fully characterize the relationship between the ignorance equivalent and the optimal choice in the original RI problem. When multiple RI agents with different learning costs interact, the ignorance equivalent emerges in equilibrium, facilitating trade that allows agents to emulate the first-best learning strategy.

Keywords: Rational inattention, information acquisition, learning.

JEL: D81, D83, C63

*University of Notre Dame: michele.muller@nd.edu

[†]Federal Reserve Bank of Philadelphia: roc.armenter@phil.frb.org

[‡]University of Notre Dame: zstangeb@nd.edu

The views expressed in this paper are solely those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Philadelphia or the Federal Reserve System. The authors would like to thank Nageeb Ali, Isaac Baley, Alex Bloedel, Tommaso Denti, Thomas Gresik, Maciej Kotowski, John Leahy, Filip Matějka, Jun Nie, editor Bart Lipman, three anonymous referees and participants at various conferences and seminars for constructive feedback and comments. Financial support from the Notre Dame Institute for Scholarship in the Liberal Arts is gratefully acknowledged. Philadelphia Fed working papers are free to download at <https://philadelphiafed.org/research-and-data/publications/working-papers>.

1 Introduction

In the face of risk and incomplete information, agents typically seek and acquire information and thus effectively shape the uncertainty that they take on. Yet, the exact details of this information acquisition process are usually unobservable. Rational Inattention (RI) posits that agents can condition their choice on any state-dependent signal but face an additive cost from generating, accessing or processing (“learning”) that information [Sims, 2003]. The result is an endogenous information structure that responds to incentives and changes in the environment. RI has been documented to reproduce empirical regularities in a variety of contexts, from portfolio design to price setting.

The agent’s rich learning possibilities are at once a strength and a hindrance to the integration of the RI framework in broader economic models. Learning introduces complementarities between actions, since a diverse menu allows the agent to better tailor her action choice to the realized state. This gives rise to interesting behavioral predictions: For example, even actions that are unattractive by themselves can open up new learning opportunities when used in combination with the existing actions and can considerably reshape behavior. Similarly, even a small information shock may lead the agent to reassess her learning strategy entirely and rely on different actions. However, the ability to generate such complex behavior comes at a cost. Outside a handful of special cases that admit a closed-form solution, the sheer size of the information structure — each learning strategy is a joint probability distribution of actions and states — can make it hard to identify and communicate the key insights.

We introduce the concept of an *ignorance equivalent* to summarize the most pertinent features of a finite RI problem for a broad class of information costs. The ignorance equivalent is a fictitious action with state-dependent payoffs that makes the agent no better or worse off whether added to or in place of the original menu of choices. That is, the ignorance equivalent is just attractive enough for the agent to forgo all existing learning opportunities and yet generates no new profitable learning opportunities when added to the menu. We show that the ignorance equivalent exists and is unique, and argue that it is a parsimonious way to compare across learning strategies, menus, and beliefs.

The ignorance equivalent is reminiscent of the certainty equivalent for choice problems under risk. Both concepts preserve key properties of the original problem but

reduce its complexity by abstracting from learning and risk, respectively. Both concepts incorporate all salient aspects of beliefs: Regardless of their (possibly subjective) priors, two agents with the same ignorance equivalent use the same learning strategies, and two agents with the same certainty equivalent purchase the same lotteries. And just as the certainty equivalent naturally emerges when a risk-neutral company designs the most profitable insurance contract, so the ignorance equivalent emerges when the menu itself is strategically designed.

The primary challenge to characterizing the ignorance equivalent is to ensure that no new profitable learning opportunities arise when it is added to the menu. Starting with an optimal strategy over the original menu, we show that it is possible to design a fictitious action that qualifies as an ignorance equivalent and is no more attractive than the original strategy at all possible beliefs about the state. Quite naturally, this is sufficient to rule out any new profitable learning opportunities and, more surprisingly, it is also necessary because of the rich learning possibilities in the RI framework. Together with the agent’s indifference at the prior, this dominance ‘across beliefs’ proves instrumental to derive the key properties of the ignorance equivalent. The first and foremost property is quite immediate now: Agents with different priors self-select into their appropriate ignorance equivalent.

We show that one can recast any finite RI problem as a choice over a richer menu without learning. We construct what we term the *learning-proof menu* from the collection of ignorance equivalents across all priors. The key observation is that, by the self-selection property, adding an ignorance equivalent to the menu generates no profitable learning opportunities regardless of the agent’s prior. Any agent is thus indifferent between the original and the learning-proof menu — and since unconditional choice is always optimal in the latter, learning becomes mute and the agent’s problem effectively turns into a ‘fixed-information’ expected utility maximization over the learning-proof menu.

The learning-proof representation succinctly summarizes how an RI agent’s learning and action choices respond to changes in the decision problem. For instance, a researcher may want to know whether an agent will adjust her learning strategy if a new action becomes available. A priori, the researcher would have to analyze the welfare impact of complex learning strategies that involve all available actions. As it turns out, it is sufficient to consider only those that rely on the original ignorance equivalent and the new action — resulting in a much more tractable condition. Re-

latedly, a newly available action may increase the appeal of a previously unchosen action because of the learning complementarities. The learning-proof menu allows us to identify which existing, but currently not chosen, actions may ever be used in response to a menu expansion. Quite simply, actions that belong to the learning-proof menu will be chosen if the right new action is made available; actions that do not will remain unchosen in all expanded menus.

The learning-proof menu is also helpful when determining the value that an agent assigns to a specific signal structure. This value depends not just on the precision of the signal but also on whether the new information complements or replaces the agent's endogenously acquired information, which is determined endogenously by the menu at the agent's disposal. For instance, even an informative signal may not induce any changes in the ignorance equivalent and thus may have no effect on her optimal behavior. In that sense, the local geometry of the learning proof menu characterizes which information is 'actionable' for the RI agent.

In addition, ignorance equivalence has relevance for experimental design. For instance, it informs a principal on how to design a menu of state-dependent transfers that incentivize an RI agent to acquire and disclose an arbitrary signal. The principal's expected cost for this 'indirect' learning differs from the agent's cost by only a constant that reflects outside options. And if the principal does *not* want the agent to learn (for example, because he wants to elicit the agent's unknown prior without generating belief-distorting learning incentives), he should design transfers such that the resulting menu is learning-proof.

In multiplayer contracting, the ignorance equivalent takes on an important strategic role: It allows RI agents with different learning technologies to coordinate on the socially optimal (joint) learning strategy. We here provide a simple example to illustrate how the ignorance equivalent emerges naturally in equilibrium.

Example. An investor (she) is looking to place her wealth in one of several different assets with state-dependent returns. Being a rationally inattentive agent, the investor will typically exert some effort into learning more about the state before making her investment choices. The ignorance equivalent would be a fund with state-dependent returns that, if available, the agent would be willing to unconditionally purchase – yet the agent would be no worse off in its absence. But what if no such fund is available?

As it turns out, this fund is exactly what an asset manager (he) with access to cheaper information would offer to attract the investor's business. The manager does

not want to reveal any free information through his fund design, and he also wants to avoid investor learning because it leads to missed business opportunities and adverse selection issues. Offering the ignorance equivalent enables the manager to extract the maximal information rents. It also implements the social optimum because all information acquisition is done by the more astute learner (the manager). \diamond

Very broadly, we show that agents can emulate the socially optimal allocation, including any associated learning, by trading the opportunity among themselves using terms dictated by an appropriately defined ignorance equivalent. We also show that this may require contracts with state-dependent transfers, and that therefore the first best can be unachievable if states remain non-verifiable ex-post. This echoes an important insight from [Cr mer and Khalil \[1992\]](#), because the possibility of learning can cause trades to fail even when they yield positive surplus to all involved agents. Our analysis goes beyond this simple case: We study the optimal allocation of an opportunity among any number of agents with different menus and learning costs. In the example, the opportunity is generated by the investor’s wealth and a menu specifies which financial assets can be purchased by each agent. As an industry insider, the manager may have lower cost and a bigger menu than a retail investor. Importantly, the ignorance equivalent is not only relevant to *avoid* learning: It can also be used to induce the right kind of learning by the right agents. In a team setting, for instance, the opportunity may represent a purchase decision of a new technology and involve agents across many units, resulting in comparative learning advantages and heterogeneous payoffs from each technology. The optimal teams solution may involve a sequence of cost-benefit investigations and involve multiple agents before settling on a final choice.

Paper structure. [Section 2](#) clarifies notation and recaps the standard Rational Inattention problem. [Section 3](#) introduces the ignorance equivalent and its key properties, and [Section 4](#) defines the learning-proof menu. [Section 5](#) and [Section 6](#) show how these tools are useful to analyze behavior in situations with a single or multiple rationally inattentive agents, respectively. [Section 7](#) discusses the related literature and [Section 8](#) concludes. All proofs are in the Appendix.

2 Rational Inattention Problem

The rationally inattentive decision maker has to implement an action from the finite menu \mathcal{A} .¹ Payoffs of each action depend on an unknown state of the world $i \in \mathcal{I} := \{1, \dots, I\}$. The ex-ante likelihood of all states are captured by a full support prior $\boldsymbol{\pi} \in \text{int}(\Delta\mathcal{I})$. No two actions are payoff equivalent, and we identify an action $\mathbf{a} \in \mathcal{A}$ by its state-dependent payoffs $(a_1, \dots, a_I) \in \mathbb{R}^I$.

The agent can condition her choice on the outcome of a costly signal $\mathcal{S} = \langle S, \mathbf{q} \rangle$, where S refers to a finite signal realization space and $\mathbf{q} \in (\Delta S)^{\mathcal{I}}$ denotes the conditional probabilities $q_i(s)$ of realization s in state i . Upon observing a signal realization s , the agent updates her belief to $\boldsymbol{\pi}^s$ according to Bayes' rule and selects a utility-maximizing action in $\arg \max_{\mathbf{a} \in \mathcal{A}} \boldsymbol{\pi}^s \cdot \mathbf{a}$. We denote the cost of signal \mathcal{S} under prior belief $\boldsymbol{\rho} \in \Delta\mathcal{I}$ by $c(\mathcal{S}, \boldsymbol{\rho}) \in [0, \infty)$, and discuss admissible cost functions in [Section 2.1](#).

The welfare of a choice problem $(\mathcal{A}, \boldsymbol{\pi}, c)$ is obtained by optimizing over expected consumption utility net of learning costs,

$$W(\mathcal{A}, \boldsymbol{\pi}, c) = \sup_{\mathcal{S}=\langle S, \mathbf{q} \rangle} \sum_{s \in S} (\boldsymbol{\pi} \cdot \mathbf{q}(s)) \left(\max_{\mathbf{a} \in \mathcal{A}} \boldsymbol{\pi}^s \cdot \mathbf{a} \right) - c(\mathcal{S}, \boldsymbol{\pi}). \quad (\text{RI})$$

Since additional actions weakly increase achievable utility $\max_{\mathbf{a}} \boldsymbol{\pi}^s \cdot \mathbf{a}$, welfare is menu-monotone, $W(\mathcal{A}, \boldsymbol{\pi}, c) \leq W(\mathcal{A}', \boldsymbol{\pi}, c)$ whenever $\mathcal{A} \subseteq \mathcal{A}'$, a property also known as ‘preference for flexibility’ [[Kreps, 1979](#)]. And since consumption utility and learning costs interact additively, optimal learning is unaffected when all payoffs are shifted by the same vector $\mathbf{v} \in \mathbb{R}^I$ ([Lemma A.3](#)).

To simplify language, we say that an agent “follows learning strategy $\mathcal{S} = \langle \mathcal{A}, \mathbf{q} \rangle$ ” if each signal realization identifies a recommended action that the agent implements. By [Lemma A.4](#), restricting the agent to such learning strategies incurs no loss of generality. We say that “unconditional implementation of \mathbf{a} is optimal” if the degenerate learning strategy with $\mathbf{q}(\mathbf{a}) = \mathbf{1}$ is optimal, and that “it is optimal to implement \mathbf{a} with positive probability” if at least one learning strategy with $\mathbf{q}(\mathbf{a}) > \mathbf{0}$ is optimal.²

¹Our notation carries over to compact menus, as long as learning is restricted to finite signals. For the cost functions that we consider, this is without loss of generality ([Lemma A.1](#)).

²Throughout, we use the convention that $\mathbf{v} \geq \mathbf{w}$ if and only if $v_i \geq w_i \forall i$, that $\mathbf{v} > \mathbf{w}$ if and only if $\mathbf{v} \geq \mathbf{w}$ and $\mathbf{v} \neq \mathbf{w}$, and that $\mathbf{v} \gg \mathbf{w}$ if and only if $v_i > w_i \forall i$.

2.1 Admissible learning costs

The agent’s cost function is uniformly posterior separable [Caplin et al., 2022] and prior-concave. Specifically, there exists a differentiable, convex potential function $\phi : \Delta\mathcal{I} \rightarrow \mathbb{R}$ such that

$$c(\langle S, \mathbf{q} \rangle, \boldsymbol{\pi}) = \sum_{s \in S} (\boldsymbol{\pi} \cdot \mathbf{q}(s)) \phi(\boldsymbol{\pi}^s) - \phi(\boldsymbol{\pi}) \quad (\text{UPS+})$$

is concave in the prior belief $\boldsymbol{\pi}$.³ Within that class, leading examples are Mutual Information [Sims, 2003], some variants of the Tsallis costs [Caplin et al., 2022] and Total Information [Bloedel and Zhong, 2020], which subsumes the Wald cost by Morris and Strack [2019] and the Fisher Information of Hébert and Woodford [2020].

While mathematically concise, the potential function in (UPS+) does not identify which specific features of the agent’s learning technology are driving our results. We instead base our proofs and discussions on five economically meaningful properties of the learning cost. As we show in Lemma A.2, they are jointly equivalent to (UPS+); but most results hold even without the last property.

(C1) Cost is continuous: For any signal space S and any cutoff $\hat{c} \in \mathbb{R}$, the pre-images $\{(\mathbf{q}, \boldsymbol{\pi}) \in (\Delta S)^{\mathcal{I}} \times \Delta\mathcal{I} \mid c(\langle S, \mathbf{q} \rangle, \boldsymbol{\pi}) \geq \hat{c}\}$ are open.

Continuity is instrumental in ensuring that the choice problem (RI) admits an optimal solution, and that welfare is locally continuous in the prior belief (Lemma A.7).

(C2) Free information disposal: Cost function $c(\cdot, \boldsymbol{\pi})$ is non-decreasing in the Blackwell order and $c(\mathcal{S}, \cdot)$ is weakly concave in the prior for all \mathcal{S} .

Condition (C2) is best explained by introducing an ‘assistant’ to the decision maker. For Blackwell monotonicity,⁴ suppose the assistant draws a signal and then, depending on its realization, communicates a garbled message to the agent. Since the garbling is uncorrelated with the state, the agent is weakly less informed than the assistant, and thus should incur a weakly lower cost. Blackwell monotonicity gives rise to the well known ‘obedience principle,’ which allows us to restrict attention to learning

³For twice-differentiable potentials, Bloedel and Zhong [2020] provide a test for prior-concavity that is based on the Hessian of ϕ .

⁴Signal $\mathcal{S} = \langle S, \mathbf{q} \rangle$ is Blackwell more informative than signal $\tilde{\mathcal{S}} = \langle \tilde{S}, \tilde{\mathbf{q}} \rangle$ if for each $s \in S$, there exists a lottery $m^s \in \Delta\tilde{S}$ such that $\tilde{\mathbf{q}} = \sum_{s \in S} m^s \mathbf{q}(s)$ [Blackwell, 1953].

strategies (Lemma A.4). For prior-concavity, suppose the assistant privately observes a free signal about the state. The concavity of $c(\mathcal{S}, \cdot)$ then simply states that it should be no more expensive (in expectation) for the agent to implement a particular learning strategy \mathcal{S} with access to the assistant’s information than it is without.⁵ Prior-concavity implies in particular that welfare W is convex in the prior belief (Lemma A.5).

The next property considers an RI agent who faces a binary choice between payoff vector \mathbf{a} and an outside option $\mathbf{0}$. Both options are equally attractive ex ante, but the choice is payoff-relevant, $\boldsymbol{\pi} \cdot |\mathbf{a}| > 0$.

(C3) Ties are broken through learning: For any belief $\boldsymbol{\pi} \in \Delta\mathcal{I}$ and any payoff vector $\mathbf{a} \in \mathbb{R}^I$ that has zero expected utility $\boldsymbol{\pi} \cdot \mathbf{a} = 0$ but is nonzero with positive probability, $\boldsymbol{\pi} \cdot |\mathbf{a}| > 0$, there exists a binary signal $\mathcal{S} = \langle \{0, 1\}, \mathbf{q} \rangle$ whose cost is below the expected benefit, $c(\mathcal{S}, \boldsymbol{\pi}) < \sum_{i \in \mathcal{I}} \pi_i q_i(1) a_i$.

The property asserts that some learning is optimal, no matter how small the stakes \mathbf{a} or how extreme the belief $\boldsymbol{\pi}$. This rules out fixed costs for learning and ensures that the agent can randomize across actions at no cost (Lemma A.6). Since the choice of \mathbf{a} is arbitrary, the condition also guarantees that there exists a noisy enough signal whose expected benefit outweighs the cost by any arbitrary factor (Lemma A.9).

The final two properties compare costs across sequential learning strategies, where the agent first draws a signal $\mathcal{S} = \langle S, \mathbf{q} \rangle$ and, upon observing $s \in S$, draws a second signal $\mathcal{S}^s = \langle S^s, \mathbf{q}^s \rangle$. We denote this contingency plan as $(\mathcal{S}, \{\mathcal{S}^s\})$, and write $\mathcal{S} \times \{\mathcal{S}^s\}$ for the one-shot implementation of the same process.⁶ It is appealing to assume that the one-shot implementation costs no more than the contingency plan, as a patient agent would otherwise exploit any possible cost savings. By assuming that $c(\mathcal{S}, \boldsymbol{\pi})$ is already the cost-minimizing envelope over all sequential information strategies [Bloedel and Zhong, 2020], one can avoid more cumbersome notation.

(C4) Sequential information acquisition brings no cost savings: In expectation, any contingency plan $(\mathcal{S}, \{\mathcal{S}^s\})$ costs no less than the one-shot signal $\mathcal{S} \times \{\mathcal{S}^s\}$, $c(\mathcal{S}, \boldsymbol{\pi}) + \sum_{s \in S} (\boldsymbol{\pi} \cdot \mathbf{q}(s)) c(\mathcal{S}^s, \boldsymbol{\pi}^s) \geq c(\mathcal{S} \times \{\mathcal{S}^s\}, \boldsymbol{\pi})$.

⁵The link between prior-concavity and the agent’s value of information has previously been discussed by Denti et al. [2022] and Bloedel and Zhong [2020].

⁶Formally, $\mathcal{S} \times \{\mathcal{S}^s\}$ is the signal $\langle S \times \bigcup_{s \in S} S^s, \tilde{\mathbf{q}} \rangle$ with $\tilde{q}_i((s, \tilde{s})) = q_i(s) q_i^s(\tilde{s})$.

The property yields an optimal stopping result [Hébert and Woodford, 2019, Zhong, 2019]: If the agent implements $\mathbf{a} \in \mathcal{A}$ at some posterior $\boldsymbol{\rho}$, then she would do so unconditionally if her prior were $\boldsymbol{\rho}$ (Lemma A.8).

The opposite inequality feels more restrictive.

(C5) Sequential information acquisition incurs no extra costs: In expectation, any contingency plan $(\mathcal{S}, \{\mathcal{S}^s\})$ costs no more than the one-shot signal $\mathcal{S} \times \{\mathcal{S}^s\}$
 $c(\mathcal{S}, \boldsymbol{\pi}) + \sum_{s \in \mathcal{S}} (\boldsymbol{\pi} \cdot \mathbf{q}(s)) c(\mathcal{S}^s, \boldsymbol{\pi}^s) \leq c(\mathcal{S} \times \mathcal{S}^s, \boldsymbol{\pi})$.

We do not mention this property anywhere in the proofs for Sections 3 and 4. This matters in particular when we analyze joint learning between heterogeneous agents in Section 6, as the social cost from their optimal joint learning sequence is only guaranteed to satisfy properties (C1) to (C4).

3 Ignorance Equivalent

The central concept of our paper is the notion of the ignorance equivalent. The ignorance equivalent of an RI problem $(\mathcal{A}, \boldsymbol{\pi}, c)$ is a payoff vector $\boldsymbol{\alpha} \in \mathbb{R}^I$ that, as a fictitious action, leaves the agent no worse as a replacement of menu \mathcal{A} and yet delivers no welfare gains as an addition to the menu \mathcal{A} .

Definition 1. The payoff vector $\boldsymbol{\alpha} \in \mathbb{R}^I$ is an **ignorance equivalent** of the RI problem $(\mathcal{A}, \boldsymbol{\pi}, c)$ if and only if

$$W(\{\boldsymbol{\alpha}\}, \boldsymbol{\pi}, c) \geq W(\mathcal{A}, \boldsymbol{\pi}, c) \quad \text{and} \quad W(\mathcal{A}, \boldsymbol{\pi}, c) \geq W(\mathcal{A} \cup \{\boldsymbol{\alpha}\}, \boldsymbol{\pi}, c).$$

Intuitively, the first condition means the agent would be willing to commit to *always* implement $\boldsymbol{\alpha}$, forgoing any learning opportunities that are present in \mathcal{A} . The second condition means that she would also commit to *never* implement $\boldsymbol{\alpha}$, forgoing any learning opportunities that arise when $\boldsymbol{\alpha}$ is added to the original menu. Together, the two conditions imply that $W(\{\boldsymbol{\alpha}\}, \boldsymbol{\pi}, c) \geq W(\mathcal{A} \cup \{\boldsymbol{\alpha}\}, \boldsymbol{\pi}, c)$, but since larger menus weakly raise welfare, all inequalities are binding. In this sense, the payoff vector $\boldsymbol{\alpha}$ is such that the agent can forgo all learning opportunities, old or new, without loss or gain. It is thus appropriately called an ‘ignorance equivalent.’

The ignorance equivalent is reminiscent of the certainty equivalent for lotteries. Neither is typically available to the agent, unless we are considering degenerate lotteries or a decision problem where no learning is optimal. Yet, both concepts allow




Figure 1: Dominated payoff vectors and construction of the ignorance equivalent.

us to abstract from the crux of the underlying economic problem (risk, learning) to reduce its complexity, all while preserving its key properties to enable comparative statics. The similarity between the two concepts extends to their construction: Just like the certainty equivalent is equal to the highest payment that is dominated by the lottery, we now show that the ignorance equivalent is equal to the payoff vector with the highest expected utility that is dominated by an optimal learning strategy.

Formally, this is what we mean with dominance by a learning strategy: An agent with belief $\rho \in \Delta\mathcal{I}$ obtains expected utility $\rho \cdot x$ from implementing action x unconditionally and incurs no learning costs. If instead the agent follows the learning strategy $\mathcal{S} = \langle \mathcal{A}, \mathbf{q} \rangle$, she achieves expected consumption utility $\mathbf{a}_i^{\mathcal{S}} = \sum_{\mathbf{a} \in \mathcal{A}} q_i(\mathbf{a}) a_i$ in each state i . Welfare is obtained by weighing these state-wise expectations by the agent's belief and subtracting the signal cost. Payoff vector x is dominated by learning strategy \mathcal{S} if the latter yields weakly larger welfare for any belief ρ . As we shall see later, the reference to other beliefs is what rules out profitable new learning opportunities that may arise from the addition of x to the menu.

Definition 2. Payoff vector x is **dominated** by a learning strategy $\mathcal{S} = \langle \mathcal{A}, \mathbf{q} \rangle$, denoted $x \preceq \mathcal{S}$, if and only if

$$\rho \cdot x \leq \rho \cdot \mathbf{a}^{\mathcal{S}} - c(\mathcal{S}, \rho) \quad \forall \rho \in \Delta\mathcal{I}.$$

Figure 1 sketches a sample RI problem (\mathcal{A}, π, c) with two states (on either axis)

and two actions (\mathbf{a}^1 and \mathbf{a}^2) to illustrate the concept of dominance. The learning strategy \mathcal{S} implements action \mathbf{a}^2 with probability 1/4 in state 1 and with probability 1 in state 2, leading to expected consumption utility $a_1^{\mathcal{S}} = \frac{1}{4}a_1^2 + \frac{3}{4}a_1^1$ and $a_2^{\mathcal{S}} = a_2^2$. For each belief, the signal cost determines the maximal expected utility for any dominated payoff vector. We indicate this upper bound for prior $\boldsymbol{\pi}$ as a dashed line and for belief $\boldsymbol{\rho}$ as a dotted line. The intersection of all such lower half-spaces forms the set of dominated payoff vectors, which is thus naturally closed, convex and unbounded below. The concavity of $c(\mathcal{S}, \cdot)$ determines the curvature of the boundary, and the cost $c(\mathcal{S}, \cdot)$ determines its distance to $\mathbf{a}^{\mathcal{S}}$. As a special case of interest, Total Information [Bloedel and Zhong, 2020] describes (UPS+) costs that are linear in the prior, $c(\mathcal{S}, \boldsymbol{\pi}) = \boldsymbol{\pi} \cdot \boldsymbol{\Delta}^{\mathcal{S}}$ for an appropriately defined $\boldsymbol{\Delta}^{\mathcal{S}} \in [0, \infty)^I$. In this case, dominance $\mathbf{x} \preceq \mathcal{S}$ is simply equivalent to the vector inequality $\mathbf{x} \leq \mathbf{a}^{\mathcal{S}} - \boldsymbol{\Delta}^{\mathcal{S}}$ and the boundary is ‘maximally curved’ around a single point, so much so that it is actually nondifferentiable.

Our first result shows how to obtain the ignorance equivalent of $(\mathcal{A}, \boldsymbol{\pi}, c)$ from any optimal signal using the dominance relationship, and establishes the existence and uniqueness of the ignorance equivalent.⁷

Theorem 1. *Each RI problem $(\mathcal{A}, \boldsymbol{\pi}, c)$ admits a unique ignorance equivalent $\boldsymbol{\alpha} \in \mathbb{R}^I$. It is obtained from any optimal learning strategy \mathcal{S} as the payoff vector that maximizes expected utility over all \mathcal{S} -dominated payoff vectors,*

$$\boldsymbol{\alpha} = \arg \max_{\mathbf{x} \preceq \mathcal{S}} \boldsymbol{\pi} \cdot \mathbf{x}. \quad (1)$$

Three arguments are key to the result:

First, unconditional implementation of the best \mathcal{S} -dominated payoff vector $\boldsymbol{\alpha}$ is as attractive as following strategy \mathcal{S} . Using Fenchel’s strong duality theorem, Lemma A.11 shows that the optimization problems

$$\bar{w}^{\mathcal{S}} = \sup_{\mathbf{x} \in \mathbb{R}^I} \{ \boldsymbol{\pi} \cdot \mathbf{a} \mid \mathbf{x} \preceq \mathcal{S} \} \quad \text{and} \quad \underline{w}^{\mathcal{S}} = \inf_{\substack{\nu \in \Delta \Delta \mathcal{I} \\ \mathbb{E}[\nu] = \boldsymbol{\pi}}} \mathbb{E}[\boldsymbol{\rho} \cdot \mathbf{a}^{\mathcal{S}} - c(\mathcal{S}, \boldsymbol{\rho}) \mid \boldsymbol{\rho} \sim \nu] \quad (2)$$

⁷Existence and uniqueness of the ignorance equivalent can be established for the larger class of posterior separable costs and derived from the Lagrangian Lemma in Caplin et al. [2022], as explained in Online Appendix B.1.2. In the main paper, our restriction to (UPS+) costs guarantees that the ignorance equivalent does not just exist but also has the properties that make it an attractive summary of the RI problem (see also Online Appendix B.2.2).

are each other’s dual. The former ($\bar{w}^{\mathcal{S}}$) maximizes the agent’s welfare, assuming she implements a dominated payoff vector $\mathbf{x} \preceq \mathcal{S}$ unconditionally. The latter ($\underline{w}^{\mathcal{S}}$) computes the agent’s worst-case welfare from following strategy \mathcal{S} , when an adversarial nature can reveal free information before the signal is drawn – or, equivalently, nature triggers a belief update to priors $\boldsymbol{\rho}$ drawn from any Bayes-plausible distribution ν . This belief update may change expected signal costs, but only to the agent’s benefit because of prior-concavity (C2). The worst case occurs under no free information, and hence $\underline{w}^{\mathcal{S}} = \boldsymbol{\pi} \cdot \mathbf{a}^{\mathcal{S}} - c(\mathcal{S}, \boldsymbol{\pi})$.

Strong duality asserts that both problems attain the same welfare, $\bar{w}^{\mathcal{S}} = \underline{w}^{\mathcal{S}}$. In the case of an optimal signal — which exists by continuity of the cost (C1) — we obtain $\underline{w}^{\mathcal{S}} = W(\mathcal{A}, \boldsymbol{\pi}, c)$. Since this value is finite, duality also guarantees that the supremum is attained, ensuring that $\boldsymbol{\alpha} = \arg \max_{\mathbf{x} \preceq \mathcal{S}} \boldsymbol{\pi} \cdot \mathbf{x}$ is well-defined with

$$W(\{\boldsymbol{\alpha}\}, \boldsymbol{\pi}, c) = \bar{w}^{\mathcal{S}} = \underline{w}^{\mathcal{S}} = W(\mathcal{A}, \boldsymbol{\pi}, c). \quad (3)$$

Visually, this means that the set of dominated payoff vectors $\{\mathbf{x} \mid \mathbf{x} \preceq \mathcal{S}\}$ in Figure 1 touches the dashed line that describes the net utility from implementing \mathcal{S} under prior $\boldsymbol{\pi}$.

Second, dominance $\boldsymbol{\alpha} \preceq \mathcal{S}$ is sufficient to rule out new learning opportunities that arise from adding $\boldsymbol{\alpha}$ to the original menu \mathcal{A} : Whenever a potential strategy over menu $\mathcal{A} \cup \{\boldsymbol{\alpha}\}$ recommends implementation of $\boldsymbol{\alpha}$ at some posterior $\boldsymbol{\rho}$, the agent is just as well off by instead continuing with strategy \mathcal{S} and relying only on actions in menu \mathcal{A} . By sequential optimality (C4), the agent can thus achieve at least as much welfare by restricting attention to menu \mathcal{A} , and $W(\mathcal{A} \cup \{\boldsymbol{\alpha}\}, \boldsymbol{\pi}, c) = W(\mathcal{A}, \boldsymbol{\pi}, c)$. Together, these three steps construct a payoff vector that satisfies both conditions of Definition 1 and thus establish existence of the ignorance equivalent.

The final step establishes that the ignorance equivalent is unique, even if the (RI) problem admits multiple optimal signals. The proof relies heavily on the flexibility of learning, which drives the agent to break all ties through learning (C3). This property allows us to construct a welfare-enhancing sequential strategy from any two distinct candidate ignorance equivalents. And by sequential optimality (C4), these welfare gains could be realized even in a one-shot implementation.

Necessity of dominance. [Theorem 1](#) implies that imposing dominance is not just sufficient to rule out learning opportunities – it is also necessary.⁸ It may surprise that there is no loss in requiring the optimal signal to be preferable to the unconditional implementation of α for *all* beliefs, including those occurring with zero probability. There is indeed no welfare gain in relaxing the optimization in [Equation \(1\)](#) to a larger domain that requires \mathcal{S} only to be preferable to payoff vector \mathbf{x} ‘locally,’ that is, for beliefs within an open neighborhood of the prior π . Prior-concavity of the cost [\(C2\)](#) implies that if \mathbf{x} achieves the same welfare as \mathcal{S} under prior π and is preferable to \mathcal{S} at *some* posterior ρ , then it also delivers welfare gains at the local perturbation $\rho^\varepsilon := (1 - \varepsilon)\pi + \varepsilon\rho$ for arbitrarily small $\varepsilon > 0$.⁹ In turn, property [\(C3\)](#) ensures that the agent could build a profitable learning strategy from *any* local welfare gains, even if they occur in a direction in which the agent does not, currently, update her beliefs.

Connection with optimal strategies. While the ignorance equivalent is always unique for any RI problem (\mathcal{A}, π, c) , there may exist multiple optimal learning strategies. Nevertheless, the construction in [Theorem 1](#) establishes a one-to-one relationship between the agent’s ignorance equivalent α and the set of all optimal strategies, which are exactly those that dominate α . This relationship can be stated implicitly, as in [Equation \(1\)](#), or explicitly, as in [Lemma A.13](#) in the appendix. In particular, for Shannon entropy costs with potential $\phi(\rho) = \lambda \sum_{i \in \mathcal{I}} \rho_i \ln(\rho_i)$, the expression simplifies to

$$\alpha_i = a_i^{\mathcal{S}} - \lambda \sum_{\substack{\mathbf{a} \in \mathcal{A} \\ q(\mathbf{a}) > 0}} q_i(\mathbf{a}) \ln(\pi_i^{\mathbf{a}}/\pi_i)$$

for any optimal learning strategy $\mathcal{S} = \langle \mathcal{A}, \mathbf{q} \rangle$.

Since the dominance relationship does not depend on the prior π , any two RI problems (\mathcal{A}, π, c) and (\mathcal{A}, π', c) that share the same ignorance equivalent also share all optimal learning strategies.

Corollary 1. *If the ignorance equivalent of RI problems (\mathcal{A}, π, c) and (\mathcal{A}, π', c) is the same, then so is the set of optimal learning strategies.*

Even though their ignorance equivalents coincide, the posterior beliefs of the

⁸In situations with multiple optimal signals, [Theorem 1](#) requires that *each* of them dominates α .

⁹Formally, $\rho^\varepsilon \cdot \mathbf{x} = \varepsilon\rho \cdot \mathbf{x} + (1 - \varepsilon)\pi \cdot \mathbf{x}$. If \mathbf{x} is strictly preferable to \mathcal{S} under prior ρ and as good as \mathcal{S} under π , the expected payoff under ρ^ε is strictly larger than $\varepsilon[\rho \cdot \mathbf{a}^{\mathcal{S}} - c(\mathcal{S}, \rho)] + (1 - \varepsilon)[\pi \cdot \mathbf{a}^{\mathcal{S}} - c(\mathcal{S}, \pi)]$, which by prior-concavity [\(C2\)](#) is at least $\rho^\varepsilon \cdot \mathbf{a}^{\mathcal{S}} - c(\mathcal{S}, \rho^\varepsilon)$.

agents will still diverge in general – but they seek out the same information sources and display the same state-dependent stochastic behavior.

Continuity of the ignorance equivalent is another consequence of this one-to-one correspondence linking the ignorance equivalent and the set of optimal learning strategies. Since the optimal RI signals are upper hemicontinuous by Berge’s Theorem (Lemma A.7), the ignorance equivalent is continuous in the prior.

Corollary 2. *The mapping $\pi \mapsto \alpha^{(\mathcal{A}, \pi, c)}$ is continuous at any prior $\pi \in \text{int}(\Delta\mathcal{I})$.*

Self-selection. By definition, the ignorance equivalent generates no additional learning under the agent’s prior. Theorem 1 further implies that the ignorance equivalent must be dominated, and thus it generates no additional learning opportunities under any prior, regardless of whether it is the one it was designed for. We refer to this result as the ‘self-selection’ property of the ignorance equivalent for the following reason: Suppose two RI agents with different priors face the same menu and cost function. Adding both ignorance equivalents to the menu would not be welfare-enhancing for either agent, yet both would now be willing to forgo learning by unconditionally implementing their respective ignorance equivalent.

Corollary 3. *Let α denote the ignorance equivalent of RI problem (\mathcal{A}, π, c) . For any belief $\rho \in \Delta\mathcal{I}$, the ignorance equivalent of (\mathcal{A}, ρ, c) is equal to that of $(\mathcal{A} \cup \{\alpha\}, \rho, c)$.*

Self-selection is a direct consequence of dominance, $\alpha \succsim \mathcal{S}$. Dominance implies that no matter the agent’s interim belief, continuing with strategy \mathcal{S} is always at least as attractive as implementing α . An agent who follows the optimal strategy for $(\mathcal{A} \cup \{\alpha\}, \rho, c)$ can achieve the same welfare in (\mathcal{A}, ρ, c) by simply continuing with \mathcal{S} whenever her strategy calls for α . For that reason, the set of optimal learning strategies grows (under the set-inclusion ordering) as ignorance equivalents are added to the menu. In that sense, adding an ignorance equivalent may give rise to additional optimal strategies, but it does not ‘distort’ the agent’s previous choice.

On a technical note, our uniqueness result relies on the full-support assumption on the prior. For beliefs on the boundary $\rho \in \Delta\mathcal{I} \setminus \text{int}(\Delta\mathcal{I})$, the ignorance equivalent is unique with probability one but can have arbitrary payoffs in the zero-probability

states,¹⁰ and [Corollary 3](#) should be understood to say that the *set* of ignorance equivalents is unchanged when α is added to the menu.

4 Learning-Proof Menu

By design, the agent’s learning ability becomes obsolete once her ignorance equivalent is added to the menu. The agent is no better off than she was before the menu expansion — but she is now also no worse off if she were to lose access to her learning technology. The collection of ignorance equivalents across all priors thus forms a fictitious menu that captures all the payoff possibilities of the RI agent regardless of her belief.

Definition 3. Letting $\alpha^{(\mathcal{A}, \pi, c)}$ denote the ignorance equivalent of RI problem (\mathcal{A}, π, c) , the **learning-proof menu** for menu \mathcal{A} under cost c is given by

$$\bar{\mathcal{A}} := \{ \alpha^{(\mathcal{A}, \pi, c)} \mid \pi \in \text{int}(\Delta\mathcal{I}) \}.$$

By studying the agent’s choice over this fictitious menu $\bar{\mathcal{A}}$ rather than the original menu \mathcal{A} , one can speak to situations with flexible learning without having to account for state-dependent stochastic choices. The learning-proof menu owes much of its appeal to properties it inherits from its various equivalent characterizations.

Alternative characterizations. The menu $\bar{\mathcal{A}}$ is ‘learning-proof’ because regardless of the agent’s prior, learning brings no welfare gains over the best unconditional action choice. At the same time, $\bar{\mathcal{A}}$ does not distort the agent’s choice because it maintains all optimal learning strategies over the original menu \mathcal{A} by [Corollary 3](#).

Theorem 2. *For any menu \mathcal{A} and cost function c , the learning-proof menu $\bar{\mathcal{A}}$ is the smallest set such that for any prior $\pi \in \text{int}(\Delta\mathcal{I})$,*

(a) *Ignorance is an optimal strategy in $(\bar{\mathcal{A}}, \pi, c)$, $W(\bar{\mathcal{A}}, \pi, c) = \max_{\mathbf{a} \in \bar{\mathcal{A}}} \pi \cdot \mathbf{a}$.*

(b) *Any strategy that is optimal in (\mathcal{A}, π, c) is also optimal in $(\bar{\mathcal{A}}, \pi, c)$.*

¹⁰Formally, one can recast the problem by dropping the dimensions with zero probability and redefining costs as the infimum over all signals that are Blackwell-equivalent over the remaining states.

Since the learning strategies of the original menu remain optimal, the problems $(\mathcal{A}, \boldsymbol{\pi}, c)$ and $(\bar{\mathcal{A}}, \boldsymbol{\pi}, c)$ also share the same ignorance equivalent by [Theorem 1](#).

The learning-proof menu accurately captures the agent’s payoff possibilities because it can be constructed from the agent’s welfare function. Indeed, because each ignorance equivalent is uniquely maximal in some direction $\boldsymbol{\pi} \in \Delta\mathcal{I}$, the learning-proof menu represents the upper boundary¹¹ of the strictly convex set

$$\bigcap_{\boldsymbol{\pi} \in \Delta\mathcal{I}} \{ \boldsymbol{x} \in \mathbb{R}^I \mid \boldsymbol{\pi} \cdot \boldsymbol{x} \leq W(\mathcal{A}, \boldsymbol{\pi}, c) \}. \quad (4)$$

This characterization has direct implications for the comparative statics of $\bar{\mathcal{A}}$: Any change in the problem parameters that raises agent welfare across all priors will move this boundary outward. Examples of such changes include the addition of new actions to the menu \mathcal{A} or reductions in the learning cost c .

Characterization (4) also yields an analogy to standard producer theory: Hotelling’s lemma states that if a firm maximizes expected profit over a strictly convex production set, its optimal production vector is equal to the price-gradient of its profit function. In our case, the agent maximizes expected utility over (the upper boundary of) a strictly convex set of ‘feasible’ payoffs. The optimum identifies the ignorance equivalent, which is thus equal to the gradient of welfare with respect to the prior.

Similarly, the learning-proof menu accounts for all learning opportunities of the agent because it can be constructed from the set of learning strategies. [Theorem 1](#) shows that each ignorance equivalent is obtained by maximizing expected utility across all dominated payoff vectors and across all learning strategies. The upper boundary of the union over all dominated payoff vectors,

$$\bigcup_{\boldsymbol{q} \in (\Delta\mathcal{A})^{\mathcal{I}}} \{ \boldsymbol{x} \in \mathbb{R}^I \mid \boldsymbol{x} \succeq \langle \mathcal{A}, \boldsymbol{q} \rangle \} \quad (5)$$

is thus also equal to the learning-proof menu $\bar{\mathcal{A}}$. [Figure 2](#) illustrates this construction of the learning-proof menu by plotting $\bar{\mathcal{A}}$ in a simple RI problem with two states and three actions, along with three sample strategies \mathcal{S}^k . For each learning strategy, the solid labeled line indicates the maximal dominated payoff vectors $\boldsymbol{x} \succeq \mathcal{S}^k$. Strategies \mathcal{S}^1 and \mathcal{S}^2 rely only on actions in the original menu and are feasible in RI problem

¹¹Point \boldsymbol{x} is part of the upper boundary of $X \subseteq \mathbb{R}^I$ if and only if $\boldsymbol{x} \in Y$ for any closed superset $Y \supseteq X$ and $\boldsymbol{x} + \boldsymbol{y} \notin X$ for any $\boldsymbol{y} > \mathbf{0}$.




Figure 2: Construction of the Learning-Proof Menu $\bar{\mathcal{A}}$ (thick solid line) for an RI problem with two states and three actions (action payoffs are indicated as black dots).

(\mathcal{A}, π, c) . They are therefore included in the union of Equation (5). Strategy \mathcal{S}^1 is optimal at π since it is tangent to $\bar{\mathcal{A}}$ at that prior. Strategy \mathcal{S}^2 is optimal at a prior that leans more toward state 1. There are other learning strategies that are not optimal at any prior: For those, all dominated payoff vectors lie strictly below $\bar{\mathcal{A}}$. Strategy \mathcal{S}^0 implements α unconditionally at no cost. It is an example of a learning strategy that becomes newly feasible in menu $\bar{\mathcal{A}}$. Strategy \mathcal{S}^0 , too, is optimal at prior π . Yet by the self-selection property (Corollary 3), none of the newly feasible learning strategies dominate additional payoffs vectors.

Reduced-Form approach. If the learning-proof menu $\bar{\mathcal{A}}$ is known, one can identify the solution to the original (RI) problem in two steps by first locating the ignorance equivalent α through the expected utility maximization problem $\max_{\alpha \in \bar{\mathcal{A}}} \pi \cdot \alpha$. Theorem 1 then identifies the full set of optimal learning strategies as exactly those that dominate α . The convenience of this approach is that the first step does not have to account for learning at all, and the second step is independent of the prior belief.

This logic may at first sound circular, as we define the learning-proof menu from the ignorance equivalent, which in turn is obtained from the optimal learning strategy. In Sections 5 and 6, we however show that the learning-proof menu often emerges

endogenously in economic applications. And so to the extent that $\bar{\mathcal{A}}$ can be observed through the equilibrium contract terms, it allows inferences on the agent’s optimal learning or problem fundamentals such as π .

Anchor Actions. The intersection $A = \mathcal{A} \cap \bar{\mathcal{A}}$ contains all actions that are implemented unconditionally under at least one prior. We call these actions *anchors*, because they connect the ‘fictitious’ learning-proof menu $\bar{\mathcal{A}}$ to the ‘physical’ menu \mathcal{A} . There are several equivalent characterizations for these actions.

Corollary 4. *Fix a menu \mathcal{A} and a cost function c . For any available action $\mathbf{a} \in \mathcal{A}$, the following are equivalent:*

- (a) *Action \mathbf{a} is an anchor of the learning-proof menu, $\mathbf{a} \in \bar{\mathcal{A}}$.*
- (b) *It is optimal to implement \mathbf{a} unconditionally for some prior $\pi \in \Delta\mathcal{I}$.*
- (c) *It is optimal to implement \mathbf{a} with positive probability for some prior $\pi \in \Delta\mathcal{I}$.*
- (d) *There exists no prior $\rho \in \Delta\mathcal{I}$ such that the ignorance equivalent of (\mathcal{A}, ρ, c) dominates \mathbf{a} statewise, $\alpha^{(\mathcal{A}, \rho, c)} > \mathbf{a}$.*

Characterization (c) points out that the RI agent always restricts her attention to anchor actions, even at priors where ignorance is not an optimal strategy. The literature uses the term *consideration set* [Caplin et al., 2018] to refer to the (typically small) submenu of actions that are implemented with positive probability. Corollary 4 implies that the union of consideration sets across priors yields exactly the set of anchors.¹²

Characterization (d) describes in what sense non-anchor actions $\mathbf{a} \in \mathcal{A} \setminus \bar{\mathcal{A}}$ are suboptimal: It is not just that each RI agent, depending on her prior, finds some other learning strategy more attractive. It is also true that the learning-proof menu contains a payoff vector $\alpha^{(\mathcal{A}, \rho, c)}$ that dominates \mathbf{a} statewise. By Theorem 1, this also implies that there exists a specific learning strategy \mathcal{S} (any one that is optimal under prior ρ) which all agents, irrespective of their prior, strictly prefer to \mathbf{a} .

¹²The equivalence between (b) and (c) follows directly from the optimal stopping rule (Lemma A.8) that has previously been derived by Hébert and Woodford [2019], Zhong [2019]. We include both characterizations for clarity.

5 Applications with a Single RI Agent

Recasting access to learning as access to a better menu yields new theoretical insights on the comparative statics of choice. In this section, we focus on situations in which a researcher or principal either seeks to learn from or about a single rationally inattentive agent.

5.1 Menu Expansion

When new actions are added to the menu, a rationally inattentive agent re-calibrates her entire learning strategy. As a result, the comparative statics of the consideration set depend in complex ways on the full menu \mathcal{A} and the cost function. Fortunately, the ignorance equivalent and the learning-proof menu bring structure to menu expansion. To determine whether a new action is implemented with positive probability, it is without loss of generality to replace the original menu with its ignorance equivalent.¹³ And since the ignorance equivalent can be derived from the optimal learning strategy (Theorem 1), the result also means that unchosen actions do not affect *regardless of whether* a new action is attractive.

Theorem 3. *Let α denote the ignorance equivalent of RI problem (\mathcal{A}, π, c) . The following hold for any payoff vector $\mathbf{a}^+ \in \mathbb{R}^I$:*

- (a) $W(\mathcal{A} \cup \{\mathbf{a}^+\}, \pi, c) > W(\mathcal{A}, \pi, c) \iff W(\{\alpha, \mathbf{a}^+\}, \pi, c) > W(\{\alpha\}, \pi, c)$.
- (b) $W(\mathcal{A} \cup \{\mathbf{a}^+\}, \pi, c) \geq W(\{\alpha, \mathbf{a}^+\}, \pi, c)$.

The absolute welfare gains from the new addition can be larger in the full menu (claim b), since the diversity of action payoffs in the menu \mathcal{A} presents more opportunities for profitable learning. As such, previously unchosen actions may affect *how often* and in which contingencies the new action is implemented.

The menu expansion result can be used in reverse to check whether ignorance is optimal: Unconditional implementation of an available action $\mathbf{a} \in \mathcal{A}$ is optimal if and only if there exist no profitable learning opportunities in any binary submenu $\{\mathbf{a}, \mathbf{a}'\} \subseteq \mathcal{A}$. The analyst does not need to worry about more complicated learning deviations that incorporate multiple other actions.

¹³In the case of Shannon entropy costs, this result is mathematically related to the ‘market entry condition’ in Caplin et al. [2018].

Corollary 5. Consider an RI problem $(\mathcal{A}, \boldsymbol{\pi}, c)$. An available payoff vector $\mathbf{a} \in \mathcal{A}$ is implemented with probability 1 if and only if

$$W(\{\mathbf{a}, \mathbf{a}'\}, \boldsymbol{\pi}, c) \leq W(\{\mathbf{a}\}, \boldsymbol{\pi}, c) \quad \forall \mathbf{a}' \in \mathcal{A}.$$

Anchor actions describe a ‘latent’ consideration set not only because appropriate changes in the prior can induce the agent to implement any anchor action with positive probability (Corollary 4(c)). The same is true if we keep the prior fixed and instead introduce a new action to the menu.

Theorem 4. Given an RI problem $(\mathcal{A}, \boldsymbol{\pi}, c)$, the following two are equivalent for any action $\mathbf{a} \in \mathcal{A}$:

- (a) Action \mathbf{a} is an anchor action $\mathbf{a} \in \mathcal{A} \cap \bar{\mathcal{A}}$.
- (b) There exists a payoff vector $\mathbf{a}^+ \in \mathbb{R}^I$ such that it is optimal to implement \mathbf{a} with positive probability in RI problem $(\mathcal{A} \cup \{\mathbf{a}^+\}, \boldsymbol{\pi}, c)$.

Matějka and McKay [2015] show by example that adding a new action to the RI agent’s menu may ‘activate’ a previously unchosen action, which now is implemented with positive probability. This activation is a distinguishing feature of Rational Inattention and is absent in fixed-information or random-utility models. Theorem 4 shows that it is exactly the anchor actions that can be activated this way.

It is worth highlighting that it is too vague to ask whether a novel action \mathbf{a}^+ would be ‘attractive’ to an agent facing RI problem $(\mathcal{A}, \boldsymbol{\pi}, c)$: The answer may be negative if only \mathbf{a}^+ is added to the menu but positive if \mathbf{a}^+ is added alongside other actions. Learning introduces complementarities between actions, and thus the attractiveness of any single action is always menu dependent. For example, suppose payoff a_j^+ is large and positive, but the action has negative payoff consequences in all other states, $a_i^+ < 0$ for all $i \neq j$. When the agent is close to certain that j is to be, \mathbf{a}^+ can bring large consumption gains relative to her existing option $\mathcal{A} = \{\mathbf{0}\}$. Yet, the cost of a sufficiently informative signal may remain prohibitive unless the agent also has access to an action that does well in the opposite contingency.

By combining Theorems 3 and 4, we nevertheless obtain a comprehensive answer for both interpretations of the question: To identify actions that are attractive if

added in isolation, one has to look no further than the ignorance equivalent. To identify actions that are attractive in some supermenu $\mathcal{A}' \supseteq \mathcal{A}$, the learning-proof menu is the answer: If \mathbf{a}^+ is located on or above $\bar{\mathcal{A}}$,¹⁴ the action becomes an anchor in $\mathcal{A} \cup \{\mathbf{a}^+\}$ and can be activated by the simultaneous addition of some complementary action. Conversely, if \mathbf{a}^+ is located below $\bar{\mathcal{A}}$, activation is impossible.

5.2 Trading Information

Since the learning-proof menu allows us to abstract away from state-dependent choices, it is particularly well suited to study situations in which a principal wants to trade information with a rationally inattentive agent. We consider three scenarios:

Selling information. A principal wants to sell an exogenous signal at the highest possible price to a rationally inattentive agent. The agent's willingness to pay depends on whether this information would replace or complement her own information acquisition, and whether it helps her discriminate between the choices at her disposal.¹⁵

Corollary 6. *An agent facing RI problem $(\mathcal{A}, \boldsymbol{\pi}, c)$ is willing to pay at most*

$$\sum_{s \in \mathcal{S}} (\boldsymbol{\pi} \cdot \mathbf{q}(s)) \max_{\mathbf{a} \in \bar{\mathcal{A}}} (\boldsymbol{\pi}^s \cdot \mathbf{a}) - \max_{\mathbf{a} \in \bar{\mathcal{A}}} \boldsymbol{\pi} \cdot \mathbf{a}$$

for access to a signal $\mathcal{S} = \langle S, \mathbf{q} \rangle$.

In particular, the local geometry of the learning-proof menu around the ignorance equivalent determines whether the agent benefits from access to noisy information, and thus whether her willingness to pay for a diffusion signal is positive. The tangent space of the surface $\bar{\mathcal{A}}$ at the ignorance equivalent, V , determines what local information is *actionable* for the agent. If a signal realization updates her belief orthogonally, $\boldsymbol{\pi}^s - \boldsymbol{\pi} \perp V$, then small enough belief updates merely cause the supporting hyperplane to rotate, but it remains tangent at $\boldsymbol{\alpha}$. Such a signal realization changes the agent's belief about the world, but not in a way that causes her to change her ignorance equivalent or, by [Corollary 1](#), her optimal learning strategy. If this is true for all signal realizations, then the information has no value to the agent.

¹⁴Formally, \mathbf{a}^+ is on or above $\bar{\mathcal{A}}$ if there exists $\mathbf{v} \gg 0$ with $\mathbf{v} \cdot \mathbf{a}^+ \geq \mathbf{v} \cdot \mathbf{a}$ for all $\mathbf{a} \in \bar{\mathcal{A}}$.

¹⁵In [Online Appendix B.2.3](#), we illustrate the interaction between exogenous information and endogenous learning with a concrete example.

Buying information. A principal wants to purchase signal $\mathcal{S} = \langle S, \mathbf{q} \rangle$ at minimal cost from a risk-neutral, rationally inattentive agent. Both share the prior $\boldsymbol{\pi}$. The principal can offer a set of state-dependent transfers to the agent but cannot independently verify the agent’s learning. Thus, he has to design the payoffs in a way that incentivizes the agent to pick payoff \mathbf{a}^s with probability $q_i(s)$ conditional on state i . To further complicate his task, any menu offer from the principal has to contain the non-empty set \mathcal{A}^0 . In the simplest case, $\mathcal{A}^0 = \{\mathbf{0}\}$ simply captures that the agent’s participation is voluntary.

Corollary 7. *The cost of incentivizing an RI agent to reveal signal \mathcal{S} is given by*

$$c(\mathcal{S}, \boldsymbol{\pi}) + W(\mathcal{A}^0, \boldsymbol{\pi}, c),$$

where \mathcal{A}^0 contains all mandatory offers, $\boldsymbol{\pi}$ is their common prior and c is the agent’s learning cost.

Since the cost function of the principal and the agent differ only by a constant, this means that the principal will either not engage with the agent or learn exactly the same amount as if he had direct access to the agent’s learning technology. Moreover, the proof in [Appendix A.3](#) is constructive in that it identifies exactly how the principal can structure the optimal transfer menu \mathcal{A} .

In mechanism design and delegation problems, ‘mandatory offers’ may appear endogenously, out of the principal’s desire to contract with multiple distinct agent types. For instance, [Yoder \[2022\]](#) considers a two-state setup where the principal is unaware of the agent’s cost. When learning is unobservable,¹⁶ information acquisition and truth-telling both have to be incentivized through the design of the transfer menu \mathcal{A} . Letting \mathcal{A}^0 denote the transfers intended for other agent types and the outside option $\mathbf{0}$, the formulation in [Corollary 7](#) implies that the resulting information rent for an agent with cost c is exactly equal to $W(\mathcal{A}^0, \boldsymbol{\pi}, c)$.

Eliciting beliefs. A principal wants to identify an RI agent’s unknown prior belief without learning distortions as in [Tsakas \[2020\]](#). For example, a marketing company may want to gauge the subjective beliefs about a new product within a representative

¹⁶This is in contrast to [Yoder \[2022\]](#), who focuses on situations with hard information and allows the principal to condition the transfer either directly on the acquired signal or on the realized posterior belief, thereby removing the need to incentivize truth-telling.

sample of consumers. Unincentivized belief elicitation is subject to all the standard pitfalls of stated preference, but too strong incentives might lead the agent to invest in learning, making her less representative of the population as whole. The solution is to tie each announcement to a state-dependent payment, and the analyst’s job is to design this menu of transfers \mathcal{A} in a way that discourages learning but incentivizes truth-telling. The analyst is successful if \mathcal{A} itself is learning-proof, $\mathcal{A} = \bar{\mathcal{A}}$, and each belief announcement $\boldsymbol{\pi}$ maps to a transfer $\mathbf{a}^\pi \in \mathcal{A}$ that maximizes expected utility, $\mathbf{a}^\pi = \arg \max_{\mathbf{a} \in \mathcal{A}} \boldsymbol{\pi} \cdot \mathbf{a}$.

If \mathcal{A} has points of non-differentiability, several announcements may map to the same transfer. In the extreme, unincentivized belief elicitation collapses all announcements to the single payoff vector $\mathcal{A} = \{0\}$.

Intuitively, the spread of \mathcal{A} governs the difference in payoff when the agent announces $\boldsymbol{\pi}$ or $\boldsymbol{\rho}$ and thus decreases the appeal of lying. At the same time, menus with a larger spread also provide incentives for learning – to counteract that, the curvature of the menu needs to be sufficiently high. The learning-proof menu tells the analyst how to transform any candidate set of transfers \mathcal{A}^0 into a menu $\bar{\mathcal{A}}^0$ that maintains the same spread but avoids learning.

6 Optimal Allocation with Multiple RI Agents

Ignorance equivalence also plays an important role in multiplayer settings. We describe a generic allocation problem and analyze whether the first-best solution is implementable when agents can learn. We study two scenarios with one common theme: State-dependent terms informed by the ignorance equivalent align agents’ learning incentives with those of the social planner.

Generic Setup. A single opportunity has to be allocated to one of finitely many RI agents $k \in \{1, \dots, K\}$. The agent who eventually executes the opportunity gets access to a menu of actions, \mathcal{A}^k , and enjoys the resulting state-dependent payoffs. All other agents receive a payoff of zero. All agents share a common prior $\boldsymbol{\pi}^0$ about the state of the world, and cost functions c^k capture their respective learning costs. The opportunity initially rests with agent 1, but it is transferable as long as it has not been executed. Agents may learn at any time, and such learning does not execute the opportunity.

To determine the social optimum, we consider a planner with direct access to the joint menu $\mathcal{A}^P = \bigcup_{k=1}^K \mathcal{A}^k$ and multiple learning technologies with costs $\{c^k\}_{k=1}^K$. The planner may save learning costs through sequential optimization, in which the realization of one signal determines the design of the next. [Bloedel and Zhong \[2020\]](#) define the planner’s *indirect* cost c^P as the one-shot equivalent of the optimization over sequential experiments with *direct* cost function $\min_k c^k$. By optimally assigning learning and opportunity ownership, the planner can generate social surplus

$$\Delta = W(\mathcal{A}^P, \boldsymbol{\pi}^0, c^P) - W(\mathcal{A}^1, \boldsymbol{\pi}^0, c^1) \quad (6)$$

relative to the autarky allocation, in which agent 1 executes the opportunity.

The goal of this section is to show when and how agents can emulate the planner’s choices through trade. To establish a trade between agents k and ℓ , both agents need to agree on *terms* $\mathbf{t} \in \mathbb{R}^I$. If they do, agent k releases the opportunity to agent ℓ , who in turn pays the former t_i once the state i realizes. If at least one agent objects, the opportunity remains with agent k . Both agents can learn and selectively participate only in those contingencies that are most favorable for them. As is standard in the mechanism design literature, we do not let the agents propose other trades.

Examples. Pre-trade learning opportunities abound in a range of economic applications. We give four concrete examples:

- *Finance.* Agent 1 is looking to invest his wealth, and the menu \mathcal{A}^1 captures the state-dependent returns of each feasible portfolio. Agents $k > 1$ are investment managers who are raising funds. Each manager could invest the wealth into portfolios \mathcal{A}^k , yielding a potentially larger return due to their improved fund access and lower research costs.
- *Procurement.* Agent 1’s business needs to produce some parts using one out of several technologies with state-dependent costs \mathcal{A}^1 . Agents $k > 1$ are suppliers with access to technologies \mathcal{A}^k , and their expertise gives them easier access to relevant information on the availability of inputs.
- *Real Estate.* Home owner 1 is contemplating a range of last-minute fixes \mathcal{A}^1 that could affect the sale price of his house. Agents $k > 1$ specialize in property redevelopment, and each has grandiose plans \mathcal{A}^k for the same lot. Although

Agent 1's fixes do not technically preclude the realization of another agent's vision, it is socially wasteful to renovate a house before demolition. Our model applies as long as the planner would select a single project in \mathcal{A}^P : The choice of the project can depend on learning, but we do not allow him to assign a sequence of projects.

- *Teams.* A firm is considering the purchase of a new technology that will benefit workers across many specialized units. Some workers are uniquely qualified to learn about specific characteristics of the technology, and the optimal teams solution may involve a sequence of cost-benefit investigations by multiple workers before settling on a final choice.

The 'opportunity' grants its owner the right to purchase any available technology and full control over access. Once the purchase is complete, the state realizes, and the owner can sell non-exclusive access rights to all his collaborators, thereby extracting the entire surplus in each state.

Absolute Advantage. We start with a situation in which agent K has weakly lower costs than everyone else and can access a weakly larger menu, so that $c^P = c^K$ and $\mathcal{A}^P = \mathcal{A}^K$. This ranking makes unconditional trade socially optimal and allows us to focus entirely on implementability constraints. Our goal is to characterize the entire set of terms T that incentivize agents 1 and K to participate unconditionally.

Agent 1 can either participate for a payoff of \mathbf{t} or decline trade in favor of any other option $\mathbf{a} \in \mathcal{A}^1$. Unconditional participation is optimal if and only if \mathbf{t} is the ignorance equivalent of the RI problem $W(\mathcal{A}^1 \cup \{\mathbf{t}\}, \boldsymbol{\pi}^0, c^1)$. By [Corollary 5](#), this is equivalent to

$$W(\{\mathbf{t}\}, \boldsymbol{\pi}^0, c^1) \geq W(\{\mathbf{a}, \mathbf{t}\}, \boldsymbol{\pi}^0, c^1) \quad \forall \mathbf{a} \in \mathcal{A}^1, \quad (7)$$

which imposes lower bounds on the terms \mathbf{t} .

Agent K can either participate and then implement any $\mathbf{a} \in \mathcal{A}^K$ for a payoff of $\mathbf{a} - \mathbf{t}$, or decline trade for a zero payoff. Unconditional participation is optimal if and only if removing the option to reject does not lower welfare,

$$W((\mathcal{A}^K - \{\mathbf{t}\}) \cup \{\mathbf{0}\}, \boldsymbol{\pi}^0, c^K) \leq W(\mathcal{A}^K - \{\mathbf{t}\}, \boldsymbol{\pi}^0, c^K).$$

Since the optimal strategy is unaffected by an unconditional shift of all payoff vectors




Figure 3: Transfers T that implement unconditional trade.

(Lemma A.3), this can equivalently be stated as $W(\mathcal{A}^K \cup \{t\}, \pi^0, c^K) \leq W(\mathcal{A}^K, \pi^0, c^K)$, or by Theorem 3 with reference to the ignorance equivalent α^K of $(\mathcal{A}^K, \pi^0, c^K)$,

$$W(\{\alpha^K, t\}, \pi^0, c^K) \leq W(\{\alpha^K\}, \pi^0, c^K), \quad (8)$$

imposing an upper bound on the terms t .

Together, Equations (7) and (8) fully characterize the set T of terms that implement unconditional trade in line with the social optimum. They do so by relying on binary learning strategies only, even for complex problems with many states and rich menus. Figure 3 illustrates the construction in a simple two-state example. Agent 1's menu comprises all payoffs marked \bullet ; agent K 's menu also contains those marked \circ . Agent K 's learning advantage manifests itself by the higher curvature of his learning-proof menu (\setminus for agent 1, \curvearrowright for agent K). Equation (7) imposes a lower bound on the transfers for each of agent 1's actions (\cdots), and Equation (8) imposes an upper bound through agent K 's ignorance equivalent (\cdots). The set is bounded because learning incentives depend on the full vector t rather than just its expected payoff. In particular, trade can fail even at terms that yield positive surplus to both agents, with $W(\mathcal{A}^1, \pi^0, c^1) < \pi^0 \cdot t < W(\mathcal{A}^K, \pi^0, c^K)$.

The selection of terms within T determines the split of the optimal surplus Δ across the two agents and thus captures their relative bargaining power. In particular, if one agent can extend a take-it-or-leave-it (TIOLI) offer, the equilibrium features unconditional trade at terms equal to his opponent's ignorance equivalent.

In this way, the agent in power secures the entire surplus Δ . He can do no better by proposing a plan with partial trade either: Incentive compatibility requires nonnegative surplus for the opponent, lest outright rejection be a profitable deviation.

The surplus can be split between more than two agents. An example of this arises when agent 2 plays the role of a platform with $\mathcal{A}^2 = \emptyset$, and her ability to form or restrict matches grants her some bargaining power. In that setting, the full surplus is realized no matter at what terms $\mathbf{t}^1 \in T$ platform 2 buys the opportunity from agent 1 and at what terms $\mathbf{t}^K \in T$ she sells it to agent K . The difference $\boldsymbol{\pi} \cdot (\mathbf{t}^K - \mathbf{t}^1)$ constitutes the expected platform profits. In particular, if the platform has TIOLI power, she can secure the entire surplus Δ by buying at terms $\boldsymbol{\alpha}^1$ and selling at terms $\boldsymbol{\alpha}^K$. She could also claim a smaller flat fee τ by choosing terms with $\mathbf{t}^K - \mathbf{t}^1 = \tau \mathbf{1}$, thereby shielding platform profits entirely from the state.

The construction also speaks to whether ‘rich’ opportunities are harder to trade. Although a menu with diverse options is no more attractive than one with just the ignorance equivalent $\boldsymbol{\alpha}$, it offers more incentives for learning. This can shrink the set of terms that implement the first best: Indeed, each additional action in Agent 1’s menu \mathcal{A}^1 imposes another lower bound through [Equation \(7\)](#). This is true even when his ignorance equivalent remains unchanged, as is the case for action \times in [Figure 3](#). In contrast, diversity of options for Agent K does not affect implementability of the first best, since his menu only enters [Equation \(8\)](#) through the ignorance equivalent $\boldsymbol{\alpha}^K$. Economically speaking, this is because the richness in Agent 1’s menu offers him potential *alternatives* to trade, whereas the richness in Agent K ’s menu matters only *conditional* on trade.

Finally, the construction of T also answers whether state-dependent payment terms are *necessary* to achieve the first-best surplus. In [Figure 3](#), the menu payoffs are heavily skewed in favor of state 1. This asymmetry generates an incentive to learn unless the suggested transfer is also higher in that state. In general, unconditional trade can be achieved with constant payment terms only when T contains a point on the diagonal $\{\tau \mathbf{1} \mid \tau \in \mathbb{R}\}$. A natural follow-up question is what happens when states are unverifiable and payment terms need to be constant across states. Would Agent 1 sometimes make an offer so generous that Agent K always accepts, or always make a less generous offer that Agent K sometimes declines? — In [Online Appendix B.2.4](#), we answer this question for the specific two-state example of [Figure 3](#) by relying on the equivalence between learning ability and access to fictitious pay-

offs. We show that when Agent 1 gets to make a take-it-or-leave-it offer, the Perfect Bayesian Equilibrium features partial trade and pre-trade learning by both agents.

Comparative Advantages. To showcase the versatility of the ignorance equivalent, we now allow for non-nested menus and non-ordered cost functions, thereby incorporating agents’ comparative advantages in terms of feasible payoffs or learning. To focus on learning incentives rather than truth-telling constraints, we assume that learning produces hard information in line with Yoder [2022]: Learning itself is non-contractible, but the realization of any signal draw is publicly observable, so that agents share the same beliefs at all times.¹⁷ In Online Appendix B.2.5, we formally define the dynamic game that results from allowing agents to acquire costly information and trade the opportunity among themselves until someone executes it. We show that the first-best allocation remains implementable: Agents can realize surplus Δ in a subgame-perfect equilibrium that features repeated trades with terms chosen from the learning proof menu $\bar{\mathcal{A}}^P$ resulting from menu \mathcal{A}^P under cost c^P .¹⁸ The key intuition is that trading at terms $\arg \max_{\mathbf{a} \in \bar{\mathcal{A}}^P} \boldsymbol{\rho} \cdot \mathbf{a}$ at all (joint) beliefs $\boldsymbol{\rho}$ aligns agents’ learning incentives with those of the planner.

Claim. *Agents emulate the planner’s strategy in a Perfect Bayesian Equilibrium by repeatedly trading at terms chosen from the planner’s learning-proof menu.*

This result speaks to the ‘teams’ literature and its focus on optimal compensation schemes that induce cooperation among workers with a common task [Bergemann and Välimäki, 2002, Holmström and Milgrom, 1990, Itoh, 1991, among others]. If the task is one of pure information acquisition, the result implies that heterogeneous workers can effectively organize themselves as long as they are able to write binding, state-contingent contracts. In equilibrium, the ‘decision power’ (viz. ‘opportunity’) changes hands repeatedly, and information is acquired sequentially by various agents working together. This simplifies for the manager: As long as she ensures that transfer promises are honored ex-post, she does not have to worry about tailoring incentives to the relative learning skills of each individual worker. Instead, she can address the team *as if* it were one worker with cost function c^P and let workers organize themselves.

¹⁷Yoder [2022] calls this results-based contracting.

¹⁸Although the cost function c^P may not belong to the (UPS+) family, we show that it satisfies properties (C1) to (C4), and hence the learning-proof menu $\bar{\mathcal{A}}^P$ is well defined.

7 Related Literature

Rational inattention models are widely used to model flexible learning in both the applied and theoretic literature, in fields ranging from finance to political economy (see [Maćkowiak et al. \[2018\]](#) for a survey). The most common approaches to RI models fall broadly into two categories: For tractability reasons, most applied research relies on special cases that admit an analytic solution (e.g., a Linear-Quadratic Gaussian setup and Shannon-entropy costs, or costs that are infinite except for Gaussian signals). Theory research typically relies on the ‘posterior-based approach’ [[Caplin et al., 2022](#)]. In this section, we discuss how ignorance equivalence stands to benefit both applied and theoretic research.

First, ignorance equivalence does not impose any functional form or domain restrictions on the agent’s learning technology. As such, it identifies robust features of rationally inattentive behavior, thereby generalizing and extending results from prior research. For instance, [Caplin et al. \[2018\]](#) derive a ‘market entry condition’ for the specific case of Shannon entropy costs to judge whether an agent stands to benefit from a newly available action. [Theorem 3](#) derives an analogous expression for any (UPS+) cost function, including those that do not satisfy the Invariance under Compression axiom of the Shannon cost function [[Caplin et al., 2022](#)]. [Matějka and McKay \[2015\]](#) show in a numeric example (also with Shannon costs) that menu expansion can raise the appeal of a previously unchosen action, a distinguishing feature that is absent in random utility models of stochastic choice. [Theorem 4](#) proves that it is exactly the set of anchor actions that can be ‘activated’ through menu expansion.

By tying our proofs to economically meaningful properties of the cost function, our work follows in the vein of [[Bloedel and Zhong, 2020](#), [Denti et al., 2022](#), [Hébert and Woodford, 2020](#), [Mensch, 2018](#), [Pomatto et al., 2018](#)], who axiomatize various cost functions. Our cost assumptions are compatible with some, but not all, of the studied cost specifications, and we rely heavily on [Bloedel and Zhong \[2020\]](#)’s characterization of costs that are sequential learning proof. Yet, our primary focus is not on characterizing which cost functions are compatible with (C1) to (C5) — rather, we argue that these intuitively appealing properties enable summary statistics that abstract away from learning. None of these other papers condense the ‘payoff possibilities’ into a single vector in the spirit of our ignorance-equivalent approach.

Second, ignorance equivalence is related to the posterior-based approach, but lends

itself to different questions. For background, the posterior-based approach expresses a signal as a Bayes-plausible distribution over posterior beliefs, and states [Equation \(RI\)](#) equivalently as

$$W(\mathcal{A}, \boldsymbol{\pi}, c) - \phi(\boldsymbol{\pi}) = \sup_{\substack{\mu \in \Delta\Delta\mathcal{I} \\ \mathbb{E}[\mu] = \boldsymbol{\pi}}} \int_{\Delta\mathcal{I}} (\max_{\mathbf{a} \in \mathcal{A}} \boldsymbol{\rho} \cdot \mathbf{a} - \phi(\boldsymbol{\rho})) d\mu(\boldsymbol{\rho}).$$

The optimal behavior is then characterized implicitly through the concave upper envelope over the ‘net utility’ functions $\boldsymbol{\rho} \mapsto \boldsymbol{\rho} \cdot \mathbf{a} - \phi(\boldsymbol{\rho})$. Ignorance equivalence can be seen as the dual of this approach. Concavification determines the least upper bound to $W(\mathcal{A}, \boldsymbol{\pi}, c) - \phi(\boldsymbol{\pi})$ for each belief — ignorance equivalence maximizes expected utility over the set of dominated payoff vectors.¹⁹ We believe that both views are useful.

From a technical perspective, the posterior-based approach is very general. It does not require prior-concavity and can even accommodate potentials that depend on the prior belief [[Caplin et al., 2022](#)]. The ignorance-equivalent approach, on the other hand, is more parsimonious because it transforms the high-dimensional search for an optimal signal (a distribution over posteriors, an element of $\Delta\Delta\mathcal{I}$) into a lower-dimensional choice (an action from the learning-proof menu, an element of $\mathbb{R}^{\mathcal{I}}$). This reduction in complexity and the geometry of the learning-proof menu can be exploited to improve numerical solution methods that stand to benefit the applied literature, as we do in [Armenter et al. \[2021\]](#) for Shannon entropy costs.

From a conceptual perspective, the posterior-based approach puts the agent’s beliefs front and center, and its obvious parallels to Bayesian Persuasion [[Gentzkow and Kamenica, 2014](#), [Kamenica and Gentzkow, 2011](#)] are helpful in information-design applications. In contrast, the ignorance-equivalence approach is particularly promising in scenarios where payoffs are determined endogenously, as in the applications that we discuss in [Sections 5](#) and [6](#). By describing signals via the induced state-dependent choices $\mathbf{q}(\mathbf{a})$ rather than the induced posteriors $\boldsymbol{\pi}^{\mathbf{a}}$, we obtain a clear distinction between prior beliefs and behavior. This separation is helpful in particular when comparing across agents with different beliefs, as in [Corollary 1](#).

There is also a large literature that still models learning much more rigidly, where the agent simply chooses whether to ‘uncover’ the state at a cost. Sometimes, this

¹⁹[Online Appendix B.1.2](#) formally establishes the connection between the ignorance equivalent and the Lagrangian Lemma in [Caplin et al. \[2022\]](#).

binary choice reflects the actual mechanics of the information acquisition process. Other times, the modeling decision is motivated more by tractability rather than realism and could fail to capture important features, like the intensive margin of learning or the agent’s endogenous focus on the most salient aspects of a multifaceted decision. By recasting the agent’s ability to learn *as if* she had access to a richer menu, ignorance equivalence opens new doors to incorporate flexible learning in multiplayer settings with transfers. Importantly, the applications in [Sections 5](#) and [6](#) highlight that ignorance equivalence is useful not only in settings where learning is, in fact, to be avoided. Indeed, a principal can use it specifically to elicit the ‘right kind’ of learning from an agent, and a team can use it to coordinate learning across members of differing expertise. The fundamental reason is that the learning-proof menu does not just make *ignorance* optimal; it also maintains optimality of the original learning strategy — as well as any ‘intermediate’ strategy that involves less informative signals. In this way, the learning-proof menu can induce a continuum of information acquisition strategies, including the one that maximizes the principal’s (or social) welfare.

8 Conclusion

The ignorance-equivalent approach simplifies the description of optimal agent behavior under costly learning. In essence, it points out that the agent’s ability to learn acts *as if* she instead had access to a fictitious action whose payoffs are given by the ignorance equivalent. In allocation problems in which multiple agents interact, this ignorance-equivalent action often emerges as part of the surplus-maximizing trading scheme. The same is true in experimental settings in which the analyst is free to design payoffs in a way that boosts his ability to, for instance, identify the agent’s prior. Yet, even in menus where the ignorance equivalent is only a conceptual shortcut, it characterizes the full set of optimal signals and allows for parsimonious comparisons across learning strategies, menus, and beliefs. Much like the certainty equivalent reduces the complexity of economic problems with exogenous uncertainty, the ignorance equivalent allows us to abstract away from learning by simply maximizing expected utility over a larger menu.

A Proofs

A.1 Properties of the cost function

We first establish that as long as the state space is finite, it is without loss of generality to assume that the optimal signal is finite, even when the menu contains infinitely many actions.²⁰ Since we are deriving this result algebraically from the (UPS+) specification, it is more convenient to represent a signal (equivalently) as a distribution over posteriors $\sigma \in \Delta\Delta\mathcal{I}$.

Lemma A.1. *For any compact menu \mathcal{A} , welfare under (UPS+) cost c ,*

$$W(\mathcal{A}, \boldsymbol{\pi}, c) = \sup_{\substack{\sigma \in \Delta\Delta\mathcal{I} \\ \mathbb{E}[\sigma] = \boldsymbol{\pi}}} \int_{\Delta\mathcal{I}} (\max_{\mathbf{a} \in \mathcal{A}} \boldsymbol{\rho} \cdot \mathbf{a} - \phi(\boldsymbol{\rho})) d\sigma(\boldsymbol{\rho}) + \phi(\boldsymbol{\pi}),$$

is maximized at a signal with finite support.

Proof. The integral is affine in σ , and thus maximized at one of the extreme points of the space $\mathcal{M} = \{\sigma \in \Delta\Delta\mathcal{I} \mid \mathbb{E}[\sigma] = \boldsymbol{\pi}\}$. By Winkler [1988], extreme points of \mathcal{M} are exactly those probability distributions whose support contains at most $|\mathcal{I}|$ posteriors. \square

Next, we show that the properties (C1) to (C5) are jointly equivalent to the (UPS+) family. The proof draws heavily on Bloedel and Zhong [2020]’s characterization of sequentially optimal cost functions, and often refers to them verbatim.

Lemma A.2. *Conditions (C1) to (C5) are jointly equivalent to the family of (UPS+) costs.*

Proof. Suppose c belongs to the (UPS+) family, which explicitly requires prior-concavity and implies Blackwell-monotonicity (C2) by convexity of ϕ . Continuity (C1) follows directly from that of the potential ϕ . Bloedel and Zhong [2020] establish that any uniformly posterior separable cost satisfies indifference to sequential learning, implying that the inequalities in (C4) and (C5) both hold. For (C3), fix any $\boldsymbol{\pi} \in \Delta\mathcal{I}$ and $\mathbf{a} \in \mathbb{R}^I$ with $\boldsymbol{\pi} \cdot \mathbf{a} = 0$ and $\boldsymbol{\pi} \cdot |\mathbf{a}| > 0$. For any $\varepsilon > 0$ small enough, consider the binary signal $\mathcal{S}^\varepsilon = \langle \{0, 1\}, \mathbf{q}^\varepsilon \rangle$ with $q_i^\varepsilon(1) = \frac{1}{2}(1 + \varepsilon a_i) \in (0, 1)$, and note

²⁰We thank Tommaso Denti for pointing us in the right direction.

that $\boldsymbol{\pi} \cdot \mathbf{q}^\varepsilon(0) = \boldsymbol{\pi} \cdot \mathbf{q}^\varepsilon(1) = 1/2$ since $\boldsymbol{\pi} \cdot \mathbf{a} = 0$. By Taylor's theorem, cost grows less than linearly in ε ,

$$\lim_{\varepsilon \rightarrow 0} \frac{c(\mathcal{S}^\varepsilon, \boldsymbol{\pi})}{\varepsilon} = \lim_{\varepsilon \rightarrow 0} \frac{c(\mathcal{S}^0, \boldsymbol{\pi})}{\varepsilon} + \left. \frac{\partial c(\mathcal{S}^\varepsilon, \boldsymbol{\pi})}{\partial \varepsilon} \right|_{\varepsilon=0} = \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} 0 + \frac{1}{2} \sum_{i \in \mathcal{I}} \nabla_i \phi(\boldsymbol{\pi})(a_i \pi_i - a_i \pi_i) = 0,$$

while the benefit grows linearly in ε , $\sum_{i \in \mathcal{I}} \pi_i q_i^\varepsilon(1) a_i = \frac{\varepsilon}{2} \sum_{i \in \mathcal{I}} \pi_i a_i^2$. In particular, benefits strictly outweigh costs for ε small enough.

Conversely, suppose c satisfies (C1) to (C5). By [Bloedel and Zhong \[2020\]](#), (C4) and (C5) jointly restrict costs to the uniformly posterior separable family with a convex potential ϕ , which means it only remains to show that ϕ is differentiable. For convenience, we extend ϕ to all of $\mathbb{R}_{\geq 0}^{\mathcal{I}}$ by imposing homogeneity of degree one $\phi(k\boldsymbol{\rho}) = k\phi(\boldsymbol{\rho})$ [as in [Hébert and Woodford, 2019](#)]. This extension is differentiable in $\mathbb{R}_{\geq 0}^{\mathcal{I}}$ if and only if ϕ itself is differentiable in $\Delta\mathcal{I}$. Fix any belief $\boldsymbol{\pi} \in \Delta\mathcal{I}$. Continuity of ϕ is implied by (C1), or an arbitrarily noisy belief perturbation would have costs bounded below by a positive constant, strictly above the free uninformative signal. To show differentiability, let G denote the set of subgradients of ϕ at $\boldsymbol{\pi}$, defined as the set of all vectors $\mathbf{g} \in \mathbb{R}^I$ such that

$$\phi(\boldsymbol{\rho}) - \phi(\boldsymbol{\pi}) \geq \mathbf{g} \cdot (\boldsymbol{\rho} - \boldsymbol{\pi}) \quad \forall \boldsymbol{\rho} \in \Delta\mathcal{I}.$$

Continuity and convexity of ϕ imply that G is nonempty, and ϕ is differentiable at $\boldsymbol{\pi}$ if and only if G is a singleton. (If $\boldsymbol{\pi}$ is on the boundary, G needs to be unique only on the positive probability states.) Note that homogeneity of degree one implies in particular that

$$\phi((1+k)\boldsymbol{\pi}) - \phi(\boldsymbol{\pi}) = k\phi(\boldsymbol{\pi}) \geq k\mathbf{g} \cdot \boldsymbol{\pi} \quad \forall k \in (-1, \infty),$$

which is binding at $k = 0$ and thus implies that $\boldsymbol{\pi} \cdot \mathbf{g} = \phi(\boldsymbol{\pi})$ for all subgradients $\mathbf{g} \in G$. Assume now by contradiction that \mathbf{g} and \mathbf{h} are both elements of G , and they differ in at least one state to which $\boldsymbol{\pi}$ assigns positive probability. Define the payoff vector $\mathbf{a} = \mathbf{g} - \mathbf{h}$, which by the above yields expected utility zero, but is nonzero with positive probability. By (C3), the agent would break her indifference between \mathbf{a} and $\mathbf{0}$ with a binary signal. For convenience, we refer to the signal using the marginal likelihood q of implementing \mathbf{a} , and the shift in posterior $\boldsymbol{\Delta} = \boldsymbol{\pi}^a - \boldsymbol{\pi}$, so we can

write the inequality in (C3) as

$$0 > q [\phi(\boldsymbol{\pi} + \boldsymbol{\Delta}) - \phi(\boldsymbol{\pi})] + (1 - q) \left[\phi \left(\boldsymbol{\pi} - \frac{q}{1 - q} \boldsymbol{\Delta} \right) - \phi(\boldsymbol{\pi}) \right] - q(\boldsymbol{\pi} + \boldsymbol{\Delta}) \cdot \mathbf{a}.$$

However, substituting both terms in square brackets with a lower bound from one of the subgradients, the right side is bounded below by

$$q\boldsymbol{\Delta} \cdot (\mathbf{g} - \mathbf{h} - \mathbf{a}) - q\boldsymbol{\pi} \cdot \mathbf{a} = 0.$$

As a result, ϕ is differentiable, and the c belongs to the (UPS+) family. \square

Lemma A.3. *Shifting the menu by a vector $\mathbf{v} \in \mathbb{R}^I$ from \mathcal{A} to $\mathcal{A} + \{\mathbf{v}\}$ increases consumption utility from each learning strategy by the same constant $\boldsymbol{\pi} \cdot \mathbf{v}$, and thus does not alter optimal learning.*

Proof. After the shift, the consumption utility of any learning strategy $\mathcal{S} = \langle \mathcal{A}, \mathbf{q} \rangle$ equals

$$\sum_{\mathbf{a} \in \mathcal{A}} (\boldsymbol{\pi} \cdot \mathbf{q}(\mathbf{a})) (\boldsymbol{\pi}^s \cdot (\mathbf{a} + \mathbf{v})) = \sum_{\mathbf{a} \in \mathcal{A}} (\boldsymbol{\pi} \cdot \mathbf{q}(\mathbf{a})) (\boldsymbol{\pi}^s \cdot \mathbf{a}) + \sum_{\mathbf{a} \in \mathcal{A}} (\boldsymbol{\pi} \cdot \mathbf{q}(\mathbf{a})) (\boldsymbol{\pi}^s \cdot \mathbf{v}).$$

Since consumption utility and learning costs interact additively, optimal learning is unaffected when all payoffs are shifted by the same vector $\mathbf{v} \in \mathbb{R}^I$ (Lemma A.3). \square

Lemma A.4. *Under (C2), it is without loss of generality in RI problem (RI) to restrict attention to learning strategies $\langle \mathcal{A}, \mathbf{q} \rangle$.*

Proof. For any signal $\mathcal{S} = \langle S, \mathbf{q} \rangle$ and conditional selection $\mathbf{a}^s \in \mathcal{A}$ for each $s \in S$, we can define the learning strategy $\hat{\mathcal{S}} = \langle \mathcal{A}, \hat{\mathbf{q}} \rangle$ with $\hat{q}_i(\mathbf{a}) = \sum_{s \in S: \mathbf{a}^s = \mathbf{a}} q_i(s)$ for each $\mathbf{a} \in \mathcal{A}$. This learning strategy achieves the same expected consumption utility and is Blackwell less informative than the original signal \mathcal{S} . By (C2), it thus achieves a weakly higher welfare. It is therefore without loss of optimality to restrict attention to learning strategies only. \square

Lemma A.5. *Under (C2), welfare W is convex in the prior belief.*

Proof. Convexity follows readily from the linearity of the consumption utility and the fact that signal costs are prior-concavity (C2). Formally, let \mathcal{S} be the optimal learning strategy for RI problem $(\mathcal{A}, t\boldsymbol{\pi} + (1 - t)\boldsymbol{\pi}')$. The welfare is bounded above by

the linear interpolation of the welfare achieved in $(\mathcal{A}, \boldsymbol{\pi})$ and $(\mathcal{A}, \boldsymbol{\pi}')$ when the same strategy is used.

$$\begin{aligned}
W(\mathcal{A}, t\boldsymbol{\pi} + (1-t)\boldsymbol{\pi}', c) &= \sum_{i=1}^I (t\pi_i + (1-t)\pi'_i) \sum_{\mathbf{a} \in \mathcal{A}} q(\mathbf{a}|i) a_i - c(\mathcal{S}, t\boldsymbol{\pi} + (1-t)\boldsymbol{\pi}') \\
&\leq t \left[\sum_{i=1}^I \pi_i \sum_{\mathbf{a} \in \mathcal{A}} q(\mathbf{a}|i) a_i - c(\mathcal{S}, \boldsymbol{\pi}) \right] + (1-t) \left[\sum_{i=1}^I \pi'_i \sum_{\mathbf{a} \in \mathcal{A}} q(\mathbf{a}|i) a_i - c(\mathcal{S}, \boldsymbol{\pi}') \right] \\
&\leq tW(\mathcal{A}, \boldsymbol{\pi}, c) + (1-t)W(\mathcal{A}, \boldsymbol{\pi}', c). \quad \square
\end{aligned}$$

Lemma A.6. *Properties (C2) and (C3) imply that any pure-noise signal $\mathcal{S} = \langle S, \mathbf{q} \rangle$ with $q_i(s) \equiv q_j(s) \forall i, j \in \mathcal{I}$ and $s \in S$ is free, and the agent can freely randomize across actions.*

Proof. All pure-noise signals are Blackwell equivalent, and thus have the same cost c^0 by (C2). Suppose by contradiction that $c^0 > 0$ and consider any two states $j, k \in \text{support}(\boldsymbol{\pi})$. Let $\mathbf{a} \in \mathbb{R}^I$ be such that

$$a_j = \frac{c^0}{2\pi_j}, a_k = -\frac{c^0}{2\pi_k} \quad \text{and} \quad a_i \equiv 0 \quad \forall i \in \mathcal{I} \setminus \{j, k\}.$$

By (C3), there exists a binary signal $\mathcal{S}' = \langle \{0, 1\}, \mathbf{q} \rangle$ with

$$c(\mathcal{S}', \boldsymbol{\pi}) < \sum_{i \in \mathcal{I}} \pi_i q_i(1) a_i = (q_j(1) - q_k(1)) \frac{c^0}{2} \leq c^0.$$

However, \mathcal{S}' is Blackwell more informative than any pure-noise signal, and so monotonicity (C2) implies that the left side is weakly bounded below by c^0 , creating a contradiction. \square

Lemma A.7. *Under (C1) and (C2), welfare W is finite and continuous in the prior belief $\boldsymbol{\pi}$. Moreover, there exists an upper hemicontinuous correspondence $Q^* : \Delta\mathcal{I} \rightarrow (\Delta\mathcal{A})^{\mathcal{I}}$ with nonempty and compact values, such that a learning strategy $\mathcal{S} = \langle \mathcal{A}, \mathbf{q} \rangle$ is optimal in RI problem $(\mathcal{A}, \boldsymbol{\pi}, c)$ if and only if $\mathbf{q} \in Q^*(\boldsymbol{\pi})$.*

Proof. Since it is without loss of generality to restrict attention to learning strategies

(Lemma A.4), we restate (RI) as a choice over conditionals

$$W(\mathcal{A}, \boldsymbol{\pi}, c) = \sup_{\mathbf{q} \in (\Delta\mathcal{A})^{\mathcal{I}}} \sum_{\mathbf{a} \in \mathcal{A}} \sum_{i \in \mathcal{I}} \pi_i q_i(\mathbf{a}) a_i - c(\langle \mathcal{A}, \mathbf{q} \rangle, \boldsymbol{\pi})$$

and define $Q^*(\boldsymbol{\pi})$ to be the set of conditionals that achieve this supremum. Since the objective is continuous over the compact domain $(\Delta\mathcal{A})^{\mathcal{I}} \times \Delta\mathcal{I}$, the claim then follows by Berge's Theorem of the Maximum. \square

Lemma A.8. *Suppose (C4) holds. Consider an optimal learning strategy $\mathcal{S} = \langle \mathcal{A}, \mathbf{q} \rangle$ to RI problem $(\mathcal{A}, \boldsymbol{\pi}, c)$. Let $\mathbf{a} \in \text{support}(\mathbf{q})$ be an action that is implemented with positive probability, and $\boldsymbol{\pi}^{\mathbf{a}}$ the associated posterior belief, $\pi_i^{\mathbf{a}} = \frac{\pi_i q_i(\mathbf{a})}{\boldsymbol{\pi} \cdot \mathbf{q}(\mathbf{a})}$. In RI problem $(\mathcal{A}, \boldsymbol{\pi}^{\mathbf{a}}, c)$, unconditional implementation of \mathbf{a} is optimal, $W(\mathcal{A}, \boldsymbol{\pi}^{\mathbf{a}}, c) = \boldsymbol{\pi}^{\mathbf{a}} \cdot \mathbf{a}$.*

Proof. Since unconditional implementation is feasible, clearly $W(\mathcal{A}, \boldsymbol{\pi}^{\mathbf{a}}, c) \geq \boldsymbol{\pi}^{\mathbf{a}} \cdot \mathbf{a}$. By contradiction, suppose $W(\mathcal{A}, \boldsymbol{\pi}^{\mathbf{a}}, c) > \boldsymbol{\pi}^{\mathbf{a}} \cdot \mathbf{a}$ and consider the sequential strategy where the agent first draws \mathcal{S} and follows its recommendation except for when it evaluates to \mathbf{a} , when she instead continues with an optimal strategy for $(\mathcal{A}, \boldsymbol{\pi}^{\mathbf{a}}, c)$. This sequential strategy achieves utility

$$\sum_{i=1}^I \pi_i \sum_{\mathbf{a}' \in \mathcal{A} \setminus \{\mathbf{a}\}} q_i(\mathbf{a}') a'_i + (\boldsymbol{\pi} \cdot \mathbf{q}(\mathbf{a})) W(\mathcal{A}, \boldsymbol{\pi}^{\mathbf{a}}, c) - c(\mathcal{S}, \boldsymbol{\pi}).$$

Since $W(\mathcal{A}, \boldsymbol{\pi}^{\mathbf{a}}, c) > \boldsymbol{\pi}^{\mathbf{a}} \cdot \mathbf{a}$, this is strictly larger than

$$\sum_{i=1}^I \pi_i \sum_{\mathbf{a} \in \mathcal{A}} q_i(\mathbf{a}) a_i - c(\mathcal{S}, \boldsymbol{\pi}),$$

which implies that this sequential strategy achieves strictly higher utility than signal \mathcal{S} alone. By (C4), the same is true for its one-shot equivalent, contradicting the optimality of \mathcal{S} in $(\mathcal{A}, \boldsymbol{\pi}, c)$. \square

Lemma A.9. *Consider any belief $\boldsymbol{\pi} \in \Delta\mathcal{I}$ and any two actions with the same expected utility, $u := \mathbf{a} \cdot \boldsymbol{\pi} = \mathbf{a}' \cdot \boldsymbol{\pi}$, that differ in at least one positive probability state. Under (C2) and (C3), and any $L > 0$, there exists a learning strategy $\langle \{\mathbf{a}, \mathbf{a}'\}, \mathbf{q} \rangle$ such that*

the change in consumption utility outweighs the signal cost by more than a factor L ,

$$\sum_{i=1}^I \pi_i (q_i(\mathbf{a})a_i + q_i(\mathbf{a}')a'_i) - u > L c(\langle \{\mathbf{a}, \mathbf{a}'\}, \mathbf{q} \rangle, \boldsymbol{\pi}).$$

Proof. The payoff vector $\tilde{\mathbf{a}} = \frac{1}{L}(\mathbf{a} - \mathbf{a}')$ has zero expected utility $\boldsymbol{\pi} \cdot \tilde{\mathbf{a}} = \frac{1}{L}(u - u) = 0$ and is non-zero with positive probability since $a_i \neq a'_i$ for some i with $\pi_i > 0$. By condition (C3), there exists a signal $\mathcal{S} = \langle \{0, 1\}, \mathbf{q} \rangle$ with $\sum_{i \in \mathcal{I}} \pi_i q_i(1) \tilde{a}_i > c(\mathcal{S}, \boldsymbol{\pi})$. By (C2), the agent incurs the same cost for the Blackwell-equivalent learning strategy where she implements \mathbf{a}' after signal 0 and \mathbf{a} after signal 1. Multiplying both sides with L , we obtain

$$\sum_{i \in \mathcal{I}} \pi_i q_i(1)(a_i - a'_i) = \sum_{i \in \mathcal{I}} \pi_i (q_i(1)a_i + q_i(0)a'_i) - u > L c(\mathcal{S}, \boldsymbol{\pi}). \quad \square$$

A.2 Ignorance equivalent

Duality. Fenchel's duality theorem allows us to apply strong duality results to optimization problems with infinitely many constraints. For our purposes, the following variant is strong enough – we formally derive it from [Borwein and Zhu \[2005, Theorem 4.4.3\]](#) in [Online Appendix B.1.1](#).

Lemma A.10 (Fenchel Duality for Linear Constraints). *For any Lebesgue measurable function $\Phi : \Delta\mathcal{I} \rightarrow \mathbb{R}$ and any belief $\boldsymbol{\pi} \in \text{int}(\Delta\mathcal{I})$, the two optimization problems*

$$p^* = \inf_{\substack{\nu \in \Delta\Delta\mathcal{I} \\ \mathbb{E}[\nu] = \boldsymbol{\pi}}} \mathbb{E}[\Phi(\boldsymbol{\rho}) \mid \boldsymbol{\rho} \sim \nu] \quad \text{and} \quad d^* = \sup_{\mathbf{a} \in \mathbb{R}^I} \{\boldsymbol{\pi} \cdot \mathbf{a} \mid \boldsymbol{\rho} \cdot \mathbf{a} \leq \Phi(\boldsymbol{\rho}) \forall \boldsymbol{\rho} \in \Delta\mathcal{I}\}$$

achieve the same bound, $p^* = d^*$. Moreover, the supremum is attained whenever that bound is finite.

When Φ is convex, the primal is trivially optimized at the degenerate distribution that puts full weight on $\boldsymbol{\pi}$, and hence the optimum in the dual attains $\Phi(\boldsymbol{\pi})$.

Lemma A.11. *Let $\Phi : \Delta\mathcal{I} \rightarrow \mathbb{R}$ denote a convex function over $\Delta\mathcal{I}$. Then for any $\boldsymbol{\pi} \in \text{int}(\Delta\mathcal{I})$, the optimization problem*

$$\sup_{\mathbf{a} \in \mathbb{R}^I} \{\boldsymbol{\pi} \cdot \mathbf{a} \mid \boldsymbol{\rho} \cdot \mathbf{a} \leq \Phi(\boldsymbol{\rho}) \forall \boldsymbol{\rho} \in \Delta\mathcal{I}\} \quad (9)$$

admits a maximum at some $\mathbf{a}^* \in \mathbb{R}^I$ with objective value $\boldsymbol{\pi} \cdot \mathbf{a}^* = \Phi(\boldsymbol{\pi})$.

Proof. Since (9) corresponds to the dual d^* in Lemma A.10, its objective value is equal to that of

$$\inf_{\substack{\nu \in \Delta \Delta \mathcal{I} \\ \mathbb{E}[\nu] = \boldsymbol{\pi}}} \mathbb{E}[\Phi(\boldsymbol{\rho}) \mid \boldsymbol{\rho} \sim \nu].$$

In turn, convexity of Φ and Jensen's inequality imply that the infimum is attained when full weight is placed on $\boldsymbol{\pi}$, implying $d^* = \Phi(\boldsymbol{\pi}) \in \mathbb{R}$. By duality, the supremum is also attained and achieves objective value $\Phi(\boldsymbol{\pi})$. \square

Existence and Uniqueness. Lemma A.11 implies that the agent is always indifferent between a signal and its most attractive dominated payoff vector.

Lemma A.12. *Under (C2), for any learning strategy \mathcal{S} and any interior belief $\boldsymbol{\pi} \in \text{int}(\Delta \mathcal{I})$, the payoff vector $\mathbf{a}^{(\mathcal{S}, \boldsymbol{\pi})} = \arg \max_{\mathbf{x} \succsim \mathcal{S}} \boldsymbol{\pi} \cdot \mathbf{x}$ exists and satisfies*

$$\boldsymbol{\pi} \cdot \mathbf{a}^{(\mathcal{S}, \boldsymbol{\pi})} = \boldsymbol{\pi} \cdot \mathbf{a}^{\mathcal{S}} - c(\mathcal{S}, \boldsymbol{\pi}). \quad (10)$$

Proof. By definition, dominance $\mathbf{x} \succsim \mathcal{S}$ is equivalent to $\boldsymbol{\rho} \cdot \mathbf{x} \leq \Phi(\boldsymbol{\rho}) \forall \boldsymbol{\rho} \in \Delta \mathcal{I}$, with an upper bound function $\Phi(\boldsymbol{\rho}) = \boldsymbol{\rho} \cdot \mathbf{a}^{\mathcal{S}} - c(\mathcal{S}, \boldsymbol{\rho})$ that is convex by prior-concavity of the cost (C2). The result then follows by Lemma A.11. \square

Of particular interest is the case where \mathcal{S} is optimal under a specific prior $\boldsymbol{\pi}$, in which case Lemma A.12 asserts that there exists a point $\boldsymbol{\alpha} \succsim \mathcal{S}$ that achieves expected utility $W(\mathcal{A}, \boldsymbol{\pi}, c)$. We now show that $\boldsymbol{\alpha}$ constitutes the ignorance equivalent.

Proof of Theorem 1: We start with existence of the ignorance equivalent and then focus on uniqueness.

Existence. Continuity of the cost function (C1) ensures that the RI problem $(\mathcal{A}, \boldsymbol{\pi}, c)$ admits an optimal learning strategy \mathcal{S} (Lemma A.7). By Lemma A.12, there exists a point $\boldsymbol{\alpha} \succsim \mathcal{S}$ such that the first inequality in Definition 1 binds, $W(\{\boldsymbol{\alpha}\}, \boldsymbol{\pi}, c) = \boldsymbol{\pi} \cdot \boldsymbol{\alpha} = W(\mathcal{A}, \boldsymbol{\pi}, c)$.

To show that the second inequality in Definition 1 also holds, let $\tilde{\mathcal{S}} = \langle \mathcal{A} \cup \{\boldsymbol{\alpha}\}, q \rangle$ denote an optimal learning strategy for RI problem $(\mathcal{A} \cup \{\boldsymbol{\alpha}\}, \boldsymbol{\pi}, c)$. If $\mathbf{q}(\boldsymbol{\alpha}) = \mathbf{0}$, then the strategy does not rely on the presence of $\boldsymbol{\alpha}$, and is thus also feasible in $(\mathcal{A}, \boldsymbol{\pi}, c)$. If $\mathbf{q}(\boldsymbol{\alpha}) > \mathbf{0}$, let $\boldsymbol{\rho}$ denote the agent's posterior belief upon observing

realization α . Consider the sequential strategy where the agent draws $\tilde{\mathcal{S}}$ and follows its recommendation except for realization α , when she instead draws and follows \mathcal{S} . The only change in the agent's payoff occurs conditional on realization α , when she achieves expected utility $\rho \cdot \alpha^{\mathcal{S}} - c(\mathcal{S}, \rho)$ rather than $\rho \cdot \alpha$. Since $\alpha \succsim \mathcal{S}$ and hence $\rho \cdot \alpha \leq \rho \cdot \alpha^{\mathcal{S}} - c(\mathcal{S}, \rho)$, the sequential strategy weakly increases welfare. A one-shot implementation of this same strategy is weakly cheaper by (C4) and thus forms a lower bound for $W(\mathcal{A}, \pi, c)$ that weakly exceeds $W(\mathcal{A} \cup \{\alpha\}, \pi, c)$.

Uniqueness. By contradiction, suppose that there exist $\alpha^1 \neq \alpha^2$ that both satisfy Definition 1. Since

$$u := W(\{\alpha^1\}, \pi, c) = W(\mathcal{A}, \pi, c) = W(\{\alpha^2\}, \pi, c),$$

the two payoff vectors achieve the same expected utility u .

By Lemma A.9, there exists a signal $\mathcal{S}^0 = \langle \{\alpha^1, \alpha^2\}, \mathbf{q}^0 \rangle$ such that

$$\sum_{i=1}^I \pi_i (q_i^0(\alpha^1) \alpha_i^1 + q_i^0(\alpha^2) \alpha_i^2) - u > 2c(\mathcal{S}^0, \pi). \quad (11)$$

We use this signal to construct an improved strategy in menu $\mathcal{A} \cup \{\alpha^k\}$ for either $k = 1$ or $k = 2$. Specifically, suppose that the agent first draws signal \mathcal{S}^0 . If its realization α^ℓ is available, $\ell = k$, the agent implements that action and otherwise proceeds with the optimal strategy for $(\mathcal{A}, \pi^\ell, c)$, where π^ℓ is the posterior belief after observing α^ℓ . The welfare of this strategy in menu $\mathcal{A} \cup \{\alpha^k\}$ is

$$V^k := \left[\sum_{i=1}^I \pi_i q_i^0(\alpha^k) \alpha_i^k \right] + (\mathbf{q}^0(\alpha^{-k}) \cdot \pi) W(\mathcal{A}, \pi^{-k}, c) - c(\mathcal{S}^0, \pi).$$

It is comprised of the agent's expected continuation utility after either of the two outcomes of the binary signal \mathcal{S}^0 , net its information costs.

The sum can be written as

$$\begin{aligned} V^1 + V^2 &= \sum_{i=1}^I \pi_i [q_i^0(\alpha^1) \alpha_i^1 + q_i^0(\alpha^2) \alpha_i^2] - 2c(\mathcal{S}^0, \pi) \\ &\quad + (\mathbf{q}^0(\alpha^1) \cdot \pi) W(\mathcal{A}, \pi^{\alpha^1}, c) + (\mathbf{q}^0(\alpha^2) \cdot \pi) W(\mathcal{A}, \pi^{\alpha^2}, c). \end{aligned}$$

The first line is strictly larger than $W(\mathcal{A}, \boldsymbol{\pi}, c)$ by Equation (11), and the second is weakly larger than that by prior-convexity of W (Lemma A.5). As a consequence, $V^k > W(\mathcal{A}, \boldsymbol{\pi}, c)$ for at least one k . Since the strategy is feasible, it also follows that $W(\mathcal{A} \cup \{\boldsymbol{\alpha}^k\}, \boldsymbol{\pi}, c) \geq V^k$. Because the addition of $\boldsymbol{\alpha}^k$ to the menu \mathcal{A} generates additional learning opportunities, it is not an ignorance equivalent. \square

Corollaries.

Proof of Corollary 1: Let $\boldsymbol{\alpha}$ denote the ignorance equivalent, and consider any learning strategy \mathcal{S} . Uniqueness of the ignorance equivalent implies that $\boldsymbol{\alpha}$ is dominated by any optimal signal according to Definition 2. So when this inequality does not hold, \mathcal{S} cannot be not optimal in either RI problem.

Conversely, unconditional implementation of the ignorance equivalent must, by Definition 1, achieve at least as much utility as following signal \mathcal{S} under both priors $\boldsymbol{\pi}$ and $\boldsymbol{\pi}'$. Dominance $\boldsymbol{\alpha} \succsim \mathcal{S}$ implies that the opposite inequality also holds at both $\boldsymbol{\pi}$ and $\boldsymbol{\pi}'$, and so \mathcal{S} achieves maximal welfare in both RI problems. \square

Lemma A.13. *The ignorance equivalent of RI problem $(\mathcal{A}, \boldsymbol{\pi}, c)$ can be computed from an optimal learning strategy \mathcal{S} as*

$$\boldsymbol{\alpha} = \mathbf{a}^{\mathcal{S}} - c(\mathcal{S}, \boldsymbol{\pi})\mathbf{1} + (\mathbf{1}\boldsymbol{\pi}^{\top} - \text{diag}(\mathbf{1}))\nabla_{\boldsymbol{\pi}} c(\mathcal{S}, \boldsymbol{\pi})$$

where $\mathbf{a}^{\mathcal{S}}$ captures the expected consumption utility that the agent achieves by following \mathcal{S} conditional on state i , $a_i^{\mathcal{S}} = \sum_{\mathbf{a} \in \mathcal{A}} q_i(\mathbf{a})a_i$, $\mathbf{1}$ is an I -dimensional vector of ones and $\text{diag}(\mathbf{1})$ is the I -dimensional identity matrix.

Proof. Unconditional implementation of the payoff vector $\boldsymbol{\alpha}$ achieves the same welfare as following signal \mathcal{S} , since

$$\boldsymbol{\pi} \cdot \boldsymbol{\alpha} = \boldsymbol{\pi} \cdot \mathbf{a}^{\mathcal{S}} - c(\mathcal{S}, \boldsymbol{\pi}) + \underbrace{(\boldsymbol{\pi}^{\top} - \boldsymbol{\pi}^{\top})\nabla_{\boldsymbol{\pi}} c(\mathcal{S}, \boldsymbol{\pi})}_{=0}.$$

Nevertheless, the payoff vector $\boldsymbol{\alpha}$ is dominated by signal \mathcal{S} , since for all $\boldsymbol{\rho} \in \Delta\mathcal{I}$,

$$\boldsymbol{\rho} \cdot \boldsymbol{\alpha} = \boldsymbol{\rho} \cdot \mathbf{a}^{\mathcal{S}} - c(\mathcal{S}, \boldsymbol{\pi}) + (\boldsymbol{\pi}^{\top} - \boldsymbol{\rho}^{\top})\nabla_{\boldsymbol{\pi}} c(\mathcal{S}, \boldsymbol{\pi}) \leq \boldsymbol{\rho} \cdot \mathbf{a}^{\mathcal{S}} - c(\mathcal{S}, \boldsymbol{\rho}),$$

where the inequality follows by prior-concavity (C2) of the cost function. Indeed,

since the inequality

$$(1-t)c(\mathcal{S}, \boldsymbol{\pi}) + tc(\mathcal{S}, \boldsymbol{\rho}) \leq c(\mathcal{S}, (1-t)\boldsymbol{\pi} + t\boldsymbol{\rho}) \quad \forall t \in [0, 1]$$

is binding at $t = 0$, the local derivative with respect to t is steeper on the right,

$$c(\mathcal{S}, \boldsymbol{\rho}) - c(\mathcal{S}, \boldsymbol{\pi}) \leq (\boldsymbol{\rho}^\top - \boldsymbol{\pi}^\top) \nabla_{\boldsymbol{\pi}} c(\mathcal{S}, \boldsymbol{\pi}). \quad \square$$

Proof of Corollary 2: Consider a sequence of beliefs $\{\boldsymbol{\pi}^n\}_{n=0}^\infty$ that converges to a prior $\boldsymbol{\pi}^0 \in \text{int}(\Delta\mathcal{I})$, and let $\boldsymbol{\alpha}^n$ denote the ignorance equivalent of RI problem $(\mathcal{A}, \boldsymbol{\pi}^n, c)$, and \mathcal{S}^n an optimal learning strategy. By Lemma A.7, the correspondence of optimal learning strategies is upper hemicontinuous with nonempty and compact values. In other words, there exists a convergent subsequence \mathcal{S}^{n_k} such that $\mathcal{S}^0 = \lim_{k \rightarrow \infty} \mathcal{S}^{n_k}$ is optimal in RI problem $(\mathcal{A}, \boldsymbol{\pi}^0, c)$. The \mathcal{S}^0 -dominated payoff vector that maximizes expected payoff under $\boldsymbol{\pi}$ is thus, by Theorem 1, equal to the ignorance equivalent $\boldsymbol{\alpha}^0$.

Moreover, by uniqueness of the ignorance equivalent, any convergent subsequence of signals generates the exact same limit vector $\boldsymbol{\alpha}^0$. It is a well-known result from real analysis that uniqueness of the limit implies that any bounded sequence,²¹ and hence $\{\boldsymbol{\alpha}^n\}_{n=0}^\infty$ itself, converges to $\boldsymbol{\alpha}^0$. \square

Proof of Corollary 3: Let \mathcal{S} and \mathcal{S}^+ denote optimal learning strategies for RI problems $(\mathcal{A}, \boldsymbol{\pi}, c)$ and $(\mathcal{A} \cup \{\boldsymbol{\alpha}\}, \boldsymbol{\rho}, c)$ respectively. By Theorem 1, $\boldsymbol{\alpha}$ is weakly dominated by \mathcal{S} under any belief. In particular, this implies that whenever \mathcal{S}^+ recommends implementation of $\boldsymbol{\alpha}$ at some posterior $\boldsymbol{\sigma}$, the agent achieves weakly higher welfare by relying on \mathcal{S} instead. The one-shot implementation \mathcal{S}^ρ of this strategy is admissible in $(\mathcal{A}, \boldsymbol{\rho}, c)$, yet achieves weakly higher welfare than \mathcal{S}^+ by (C4). This implies

$$W(\mathcal{A}, \boldsymbol{\rho}, c) \geq \boldsymbol{\rho} \cdot \boldsymbol{\alpha}^{\mathcal{S}^\rho} - c(\mathcal{S}^\rho, \boldsymbol{\rho}) \geq \boldsymbol{\rho} \cdot \boldsymbol{\alpha}^{\mathcal{S}^+} - c(\mathcal{S}^+, \boldsymbol{\rho}) = W(\mathcal{A} \cup \{\boldsymbol{\alpha}\}, \boldsymbol{\rho}, c),$$

with the opposite inequality binding by menu-monotonicity. In particular, \mathcal{S}^ρ is optimal in both RI problems $(\mathcal{A}, \boldsymbol{\rho}, c)$ and $(\mathcal{A} \cup \{\boldsymbol{\alpha}\}, \boldsymbol{\rho}, c)$. By Theorem 1, the ignorance

²¹By contradiction, suppose the sequence $\{\boldsymbol{\alpha}^n\}_{n=0}^\infty$ does not converge. By definition, this implies that there exists $\varepsilon > 0$ and a subsequence $\{\boldsymbol{\alpha}^{n_k}\}$ with $\|\boldsymbol{\alpha}^{n_k} - \boldsymbol{\alpha}^0\| > \varepsilon$ for all $k \in \mathbb{N}$. Still, the associated learning strategies have bounded conditionals, and thus admit a convergent subsequence. The Bolzano-Weierstrass Theorem asserts that this bounded subsequence admits a convergent sub-subsequence, but its limit payoff vector must be different from $\boldsymbol{\alpha}^0$.

equivalent for both problems is thus equal to the unique $\arg \max_{\mathbf{x} \succsim_{\mathcal{S}^\rho} \boldsymbol{\pi}} \boldsymbol{\pi} \cdot \mathbf{x}$. \square

A.3 Learning-Proof Menu

Proof of Theorem 2: Let \mathcal{A}' denote the smallest set that satisfies the two properties. Fix any prior $\boldsymbol{\pi} \in \text{int}(\Delta\mathcal{I})$ and let \mathcal{S} denote a learning strategy that is optimal in $(\mathcal{A}, \boldsymbol{\pi}, c)$. By property (a), \mathcal{A}' needs to contain the ignorance equivalent of RI problem $(\mathcal{A}', \boldsymbol{\pi}, c)$. By property (b), strategy \mathcal{S} is also optimal in $(\mathcal{A}', \boldsymbol{\pi}, c)$. The most attractive \mathcal{S} -dominated payoff vector thus identifies the ignorance equivalent under both RI problems $(\mathcal{A}', \boldsymbol{\pi}, c)$ and $(\mathcal{A}, \boldsymbol{\pi}, c)$ by Theorem 1. In particular, the two ignorance equivalents are equal to each other, and so \mathcal{A}' contains all the elements of $\bar{\mathcal{A}}$ according to Definition 3.

Conversely, fix any belief $\boldsymbol{\pi}$ and let \mathcal{S}' denote an optimal strategy in RI problem $(\mathcal{A} \cup \bar{\mathcal{A}}, \boldsymbol{\pi}, c)$. By Lemma A.1, it is without loss of generality to assume that \mathcal{S}' only relies on finitely many actions $A \subseteq \bar{\mathcal{A}}$ in addition to those in \mathcal{A} . By iterative application of Corollary 3, the ignorance equivalent of $(\mathcal{A} \cup A, \boldsymbol{\pi}, c)$ is the same as that of $(\mathcal{A}, \boldsymbol{\pi}, c)$, and hence the same is true of welfare $W(\mathcal{A} \cup \bar{\mathcal{A}}, \boldsymbol{\pi}, c) = W(\mathcal{A}, \boldsymbol{\pi}, c)$. By menu-monotonicity, the former is weakly larger than $W(\bar{\mathcal{A}}, \boldsymbol{\pi}, c)$. The agent can achieve this upper bound through unconditional implementation of $\boldsymbol{\alpha}^{(\mathcal{A}, \boldsymbol{\pi}, c)}$, making ignorance optimal in $(\bar{\mathcal{A}}, \boldsymbol{\pi}, c)$ and establishing property (a). For property (b), note that the equality of welfare across the two menus implies that any optimal strategy \mathcal{S} in $(\mathcal{A}, \boldsymbol{\pi}, c)$ remains optimal in $(\bar{\mathcal{A}}, \boldsymbol{\pi}, c)$ – provided that it is still feasible. And indeed, whenever \mathcal{S} recommends an action $\mathbf{a} \in \mathcal{A}$ at some posterior $\boldsymbol{\rho}$, Lemma A.8 guarantees that \mathbf{a} itself is the ignorance equivalent of RI problem $(\mathcal{A}, \boldsymbol{\rho}, c)$, and hence is included in $\bar{\mathcal{A}}$. \square

Proof of Corollary 4: We proceed in steps.

(a) \Leftrightarrow (b): Definition 3 states that $\mathbf{a} \in \bar{\mathcal{A}}$ if and only if it is the ignorance equivalent of $(\mathcal{A}, \boldsymbol{\pi}, c)$ for some prior $\boldsymbol{\pi} \in \Delta\mathcal{I}$. Since $\mathbf{a} \in \mathcal{A}$, Definition 1 collapses to just requiring that unconditional implementation of \mathbf{a} is optimal, $W(\mathcal{A}, \boldsymbol{\pi}, c) = W(\{\mathbf{a}\}, \boldsymbol{\pi}, c)$.

(b) \Leftrightarrow (c): Since unconditional implementation implies that \mathbf{a} is chosen with positive probability, (b) trivially implies (c). Lemma A.8 formalizes the converse claim.

(b) \Rightarrow (d): Proving the contrapositive claim, assume also that there exists $\boldsymbol{\rho} \in \Delta\mathcal{I}$

such that $\boldsymbol{\alpha}^{(\mathcal{A}, \rho, c)} > \mathbf{a}$. In particular, for any prior $\boldsymbol{\pi}$, we have $\boldsymbol{\pi} \cdot \mathbf{a} < \boldsymbol{\pi} \cdot \boldsymbol{\alpha}^{(\mathcal{A}, \rho, c)}$. By self-selection of the ignorance equivalent (Corollary 3), adding $\boldsymbol{\alpha}^{(\mathcal{A}, \rho, c)}$ to the menu is not welfare-enhancing under $\boldsymbol{\pi}$, implying in particular that its unconditional implementation can achieve at most utility $W(\mathcal{A}, \boldsymbol{\pi}, c)$. Taken together, at any prior $\boldsymbol{\pi}$,

$$\boldsymbol{\pi} \cdot \mathbf{a} < \boldsymbol{\pi} \cdot \boldsymbol{\alpha}^{(\mathcal{A}, \rho, c)} \leq W(\mathcal{A}, \boldsymbol{\pi}, c),$$

proving that unconditional implementation of \mathbf{a} is suboptimal.

(d) \Rightarrow (a): Proving the contrapositive claim, assume that \mathbf{a} is not part of the learning-proof menu $\bar{\mathcal{A}}$. Since we can write the learning-proof menu as the upper boundary of an intersection of half-spaces with positive orthogonality vectors by Equation (4), \mathbf{a} lies strictly below each individual half-space. As a consequence, the ray $\{\mathbf{a} + t\mathbf{1} \mid t \geq 0\}$ crosses the learning-proof menu at some point $\boldsymbol{\alpha} \in \bar{\mathcal{A}}$, and by Definition 3, this point represents the ignorance equivalent under some prior. \square

Menu Expansion. In this subsection, we rely on all cost properties (C1) to (C5). For any RI problem $(\mathcal{A}, \boldsymbol{\pi}, c)$, we define the set of $\boldsymbol{\pi}$ -dominated payoff vectors

$$\bar{\mathcal{A}}^\pi = \{\mathbf{x} \in \mathbb{R}^I \mid W(\mathcal{A} \cup \{\mathbf{x}\}, \boldsymbol{\pi}, c) \leq W(\mathcal{A}, \boldsymbol{\pi}, c)\} \quad (12)$$

as those that do not increase the welfare of an agent with prior $\boldsymbol{\pi}$ when added to the menu. We first establish Theorem 3, which ensures that $\bar{\mathcal{A}}^\pi$ can equivalently be stated by replacing all references to \mathcal{A} with the ignorance equivalent $\{\boldsymbol{\alpha}\}$.

Proof of Theorem 3: We start by proving part (b). Let \mathcal{S}^0 denote an optimal learning strategy for RI problem $(\{\boldsymbol{\alpha}, \mathbf{a}^+\}, \boldsymbol{\pi}, c)$ and \mathcal{S}^1 an optimal learning strategy for $(\mathcal{A}, \boldsymbol{\pi}, c)$. Consider a sequential strategy where the agent first draws \mathcal{S}^0 and follows its recommendation except for when it evaluates to $\boldsymbol{\alpha}$, when she draws and follows \mathcal{S}^1 . The sequential strategy is weakly more attractive because \mathcal{S}^1 dominates $\boldsymbol{\alpha}$, implying $W(\mathcal{A} \cup \{\mathbf{a}^+\}, \boldsymbol{\pi}, c) \geq W(\{\boldsymbol{\alpha}, \mathbf{a}^+\}, \boldsymbol{\pi}, c)$.

We now establish part (a). The backwards implication is a direct consequence of the argument we just made, since $W(\{\boldsymbol{\alpha}, \mathbf{a}^+\}, \boldsymbol{\pi}, c) > W(\{\boldsymbol{\alpha}\}, \boldsymbol{\pi}, c)$ implies that

$$W(\mathcal{A} \cup \{\mathbf{a}^+\}, \boldsymbol{\pi}, c) \geq W(\{\boldsymbol{\alpha}, \mathbf{a}^+\}, \boldsymbol{\pi}, c) > W(\{\boldsymbol{\alpha}\}, \boldsymbol{\pi}, c) = W(\mathcal{A}, \boldsymbol{\pi}, c),$$

where the first inequality restates claim (b) and the last equality follows from the

definition of the ignorance equivalent.

Next, we provide a direct proof of the forward implication,

$$W(\mathcal{A} \cup \{\mathbf{a}^+\}, \boldsymbol{\pi}, c) > W(\mathcal{A}, \boldsymbol{\pi}, c) \implies W(\{\boldsymbol{\alpha}, \mathbf{a}^+\}, \boldsymbol{\pi}, c) > W(\{\boldsymbol{\alpha}\}, \boldsymbol{\pi}, c).$$

To do so, let $\mathcal{S}^+ = \langle \mathcal{A} \cup \{\mathbf{a}^+\}, \mathbf{q}^+ \rangle$ denote an optimal learning strategy for RI problem $(\mathcal{A} \cup \{\mathbf{a}^+\}, \boldsymbol{\pi}, c)$, and $\mathcal{S} = \langle \mathcal{A}, \mathbf{q} \rangle$ one for $(\mathcal{A}, \boldsymbol{\pi}, c)$. Let $\Delta = W(\mathcal{A} \cup \{\mathbf{a}^+\}, \boldsymbol{\pi}, c) - W(\mathcal{A}, \boldsymbol{\pi}, c) > 0$ denote the difference in the welfare between the two strategies.

Consider an agent who relies on \mathcal{S}^+ with probability ε and on \mathcal{S} otherwise. Since the gains from \mathcal{S}^+ are realized only with probability ε , this strategy yields welfare $W(\mathcal{A}, \boldsymbol{\pi}, c) + \varepsilon\Delta$. Note also that mixing changes the marginal likelihood that \mathbf{a}^+ is implemented, but not its associated posterior $\boldsymbol{\pi}^+$.

We now suggest a Blackwell-equivalent implementation strategy that proceeds in two steps: First, the agent draws a binary signal $\mathcal{S}_\varepsilon^0 = \langle \{-, +\}, \mathbf{q}_\varepsilon^0 \rangle$ which pools all realizations other than \mathbf{a}^+ by returning $-$ with probability $\mathbf{q}_\varepsilon^0(-) = 1 - \varepsilon\mathbf{q}^+(\mathbf{a}^+)$. If $\mathcal{S}_\varepsilon^0$ evaluates to $+$, the agent implements \mathbf{a}^+ at the same posterior $\boldsymbol{\pi}^+$ as above. If $\mathcal{S}_\varepsilon^0$ evaluates to $-$, the agent updates her belief to $\boldsymbol{\pi}_\varepsilon^-$ and draws $\mathcal{S}_\varepsilon^1$, which conditions on ‘not \mathbf{a}^+ ’ by returning $\mathbf{a} \in \mathcal{A}$ with probability $\frac{\varepsilon q_i^+(\mathbf{a}) + (1-\varepsilon)q_i(\mathbf{a})}{1-\varepsilon q_i^+(\mathbf{a})}$ in state i . Since the agent is indifferent across all Blackwell-equivalent sequential information strategies under (C4) and (C5), this strategy too yields welfare $W(\mathcal{A}, \boldsymbol{\pi}, c) + \varepsilon\Delta$. Welfare can only increase if we replace signal $\mathcal{S}_\varepsilon^1$ by the ignorance equivalent $\boldsymbol{\alpha}^\varepsilon$ of the corresponding RI problem $(\mathcal{A}, \boldsymbol{\pi}_\varepsilon^-, c)$, hence

$$(\boldsymbol{\pi} \cdot \mathbf{q}_\varepsilon^0(-))(\boldsymbol{\pi}_\varepsilon^- \cdot \boldsymbol{\alpha}^\varepsilon) + (\boldsymbol{\pi} \cdot \mathbf{q}_\varepsilon^0(+))(\boldsymbol{\pi}^+ \cdot \mathbf{a}^+) - c(\mathcal{S}_\varepsilon^0, \boldsymbol{\pi}) \geq W(\mathcal{A}, \boldsymbol{\pi}, c) + \varepsilon\Delta. \quad (13)$$

By the law of total probability, the average posterior is equal to the prior, allowing us to express the posterior $\boldsymbol{\pi}_\varepsilon^-$ as a function of the prior $\boldsymbol{\pi}$ and the posterior $\boldsymbol{\pi}^+$,

$$(\boldsymbol{\pi} \cdot \mathbf{q}_\varepsilon^0(-))\boldsymbol{\pi}_\varepsilon^- = \boldsymbol{\pi} - (\boldsymbol{\pi} \cdot \mathbf{q}_\varepsilon^0(+))\boldsymbol{\pi}^+ = \boldsymbol{\pi} - \varepsilon(\boldsymbol{\pi} \cdot \mathbf{q}^+(\mathbf{a}^+))\boldsymbol{\pi}^+,$$

with both $\boldsymbol{\pi}$ and $\boldsymbol{\pi}^+$ independent of ε .

We now show that replacing $\boldsymbol{\alpha}^\varepsilon$ with $\boldsymbol{\alpha}$ still achieves a positive welfare gain for small enough ε . First, prior-continuity of the ignorance equivalent (Corollary 2)

implies that there exists $\varepsilon > 0$ such that

$$(\boldsymbol{\pi} \cdot \mathbf{q}^+(\mathbf{a}^+))\boldsymbol{\pi}^+ \cdot (\boldsymbol{\alpha} - \boldsymbol{\alpha}^\varepsilon) < \frac{\Delta}{2}.$$

Second, the self-selection property of the ignorance equivalent ([Corollary 3](#)) further implies that $\boldsymbol{\pi} \cdot (\boldsymbol{\alpha} - \boldsymbol{\alpha}^\varepsilon) \geq 0$, and so the welfare loss of replacing $\boldsymbol{\alpha}^\varepsilon$ with $\boldsymbol{\alpha}$ is bounded below by

$$(\boldsymbol{\pi} \cdot \mathbf{q}_\varepsilon^0(-))\boldsymbol{\pi}_\varepsilon^- \cdot (\boldsymbol{\alpha} - \boldsymbol{\alpha}^\varepsilon) = \boldsymbol{\pi} \cdot (\boldsymbol{\alpha} - \boldsymbol{\alpha}^\varepsilon) - \varepsilon(\boldsymbol{\pi} \cdot \mathbf{q}^+(\mathbf{a}^+))\boldsymbol{\pi}^+ \cdot (\boldsymbol{\alpha} - \boldsymbol{\alpha}^\varepsilon) > -\varepsilon \frac{\Delta}{2}. \quad (14)$$

Taken together, [Equations \(13\)](#) and [\(14\)](#) imply that the agent can achieve welfare

$$(\boldsymbol{\pi} \cdot \mathbf{q}_\varepsilon^0(-))\boldsymbol{\pi}_\varepsilon^- \cdot \boldsymbol{\alpha} + (\boldsymbol{\pi} \cdot \mathbf{q}_\varepsilon^0(+))(\boldsymbol{\pi}^+ \cdot \mathbf{a}^+) - c(\mathcal{S}_\varepsilon^0, \boldsymbol{\pi}) \geq W(\mathcal{A}, \boldsymbol{\pi}, c) + \varepsilon \frac{\Delta}{2}$$

by drawing $\mathcal{S}_\varepsilon^0$ and implementing $\boldsymbol{\alpha}$ upon realization – and \mathbf{a}^+ otherwise. Since this strategy is feasible in RI problem $(\{\boldsymbol{\alpha}, \mathbf{a}^+\}, \boldsymbol{\pi}, c)$, it implies in particular that $W(\{\boldsymbol{\alpha}, \mathbf{a}^+\}, \boldsymbol{\pi}, c) > W(\{\boldsymbol{\alpha}\}, \boldsymbol{\pi}, c)$. \square

Inductive application of this result yields a binary characterization of situations where ignorance is optimal.

Proof of [Corollary 5](#): Suppose first that unconditional implementation of $\mathbf{a} \in \mathcal{A}$ is optimal in RI problem $(\mathcal{A}, \boldsymbol{\pi}, c)$. By optimality, $W(\{\mathbf{a}\}, \boldsymbol{\pi}, c) = W(\mathcal{A}, \boldsymbol{\pi}, c)$, and by menu-monotonicity, the latter is weakly larger than $W(\{\mathbf{a}, \mathbf{a}'\}, \boldsymbol{\pi}, c)$ for each $\mathbf{a}' \in \mathcal{A}$.

Conversely, suppose condition $W(\{\mathbf{a}, \mathbf{a}'\}, \boldsymbol{\pi}, c) \geq W(\{\mathbf{a}\}, \boldsymbol{\pi}, c)$ holds for each $\mathbf{a}' \in \mathcal{A}$. We prove by induction that unconditional implementation of \mathbf{a} is optimal in RI problem $(A, \boldsymbol{\pi}, c)$, starting with the trivial case $A = \{\mathbf{a}\}$ and adding actions one-by-one. For the inductive step, assume $W(\{\mathbf{a}\}, \boldsymbol{\pi}, c) = W(A, \boldsymbol{\pi}, c)$ for some subset $A \subseteq \mathcal{A}$, and consider what happens when $\mathbf{a}' \in \mathcal{A} \setminus A$ is added to the menu. Since $\mathbf{a} \in A$, the assumption satisfies all the conditions of [Definition 1](#), and \mathbf{a} denotes the ignorance equivalent of $(A, \boldsymbol{\pi}, c)$. By [Theorem 3\(a\)](#), the condition $W(\{\mathbf{a}, \mathbf{a}'\}, \boldsymbol{\pi}, c) \leq W(\{\mathbf{a}\}, \boldsymbol{\pi}, c)$ then implies that $W(A \cup \{\mathbf{a}'\}, \boldsymbol{\pi}, c) \leq W(A, \boldsymbol{\pi}, c)$. Menu-monotonicity implies the opposite inequality and thus

$$W(A \cup \{\mathbf{a}'\}, \boldsymbol{\pi}, c) = W(A, \boldsymbol{\pi}, c) = W(\{\mathbf{a}\}, \boldsymbol{\pi}, c). \quad \square$$

We now turn back to our definition of $\boldsymbol{\pi}$ -dominated payoff vectors in Equation (12) and rewrite it in a way that fits Lemma A.11. To do so, we let $\mathcal{S}_\pi^{(k,\boldsymbol{\sigma})}$ denote the binary signal that, starting from prior $\boldsymbol{\pi}$, reaches posterior $\boldsymbol{\sigma}$ with marginal likelihood $k \in (0, 1)$. By the law of total probability, the same signal reaches posterior $\frac{1}{1-k}\boldsymbol{\pi} - \frac{k}{1-k}\boldsymbol{\sigma}$ with marginal likelihood $1 - k$. For $k > 0$ small enough, all probabilities are non-degenerate. We also define function $\Phi_\pi : \Delta\mathcal{I} \rightarrow \mathbb{R}$ as

$$\begin{aligned} \Phi_\pi(\boldsymbol{\sigma}) &:= \boldsymbol{\sigma} \cdot \boldsymbol{\alpha}^{(\mathcal{A}, \boldsymbol{\pi}, c)} + \lim_{k \downarrow 0} \frac{1}{k} c(\mathcal{S}_\pi^{(k,\boldsymbol{\sigma})}, \boldsymbol{\pi}) \\ &= \boldsymbol{\sigma} \cdot \boldsymbol{\alpha}^{(\mathcal{A}, \boldsymbol{\pi}, c)} + \phi(\boldsymbol{\sigma}) + \frac{\partial}{\partial k} \left[(1-k)\phi \left(\frac{1}{1-k}\boldsymbol{\pi} - \frac{k}{1-k}\boldsymbol{\sigma} \right) \right]_{k=0} \\ &= \boldsymbol{\sigma} \cdot \boldsymbol{\alpha}^{(\mathcal{A}, \boldsymbol{\pi}, c)} + \phi(\boldsymbol{\sigma}) - \phi(\boldsymbol{\pi}) + \nabla \phi(\boldsymbol{\pi}) \cdot (\boldsymbol{\pi} - \boldsymbol{\sigma}), \end{aligned} \quad (15)$$

and use it to restate $\bar{\mathcal{A}}^\pi$.

Lemma A.14. *The function Φ_π given in Equation (15) is convex and finite-valued, and $\bar{\mathcal{A}}^\pi = \{\boldsymbol{x} \in \mathbb{R}^I \mid \boldsymbol{\sigma} \cdot \boldsymbol{x} \leq \Phi_\pi(\boldsymbol{\sigma}) \ \forall \boldsymbol{\sigma} \in \Delta\mathcal{I}\}$.*

Proof. By Theorem 3(a), $\bar{\mathcal{A}}^\pi$ can be stated with reference to the ignorance equivalent $\boldsymbol{\alpha} = \boldsymbol{\alpha}^{(\mathcal{A}, \boldsymbol{\pi}, c)}$ alone, $\bar{\mathcal{A}}^\pi = \{\boldsymbol{x} \in \mathbb{R}^I \mid W(\{\boldsymbol{\alpha}, \boldsymbol{x}\}, \boldsymbol{\pi}, c) \leq W(\{\boldsymbol{\alpha}\}, \boldsymbol{\pi}, c)\}$. And since the menu $\{\boldsymbol{\alpha}, \boldsymbol{x}\}$ is binary, any feasible strategy can be described as drawing $\mathcal{S}_\pi^{(k,\boldsymbol{\sigma})}$ for a small enough k , and implementing \boldsymbol{x} at posterior $\boldsymbol{\rho}$ and $\boldsymbol{\alpha}$ otherwise. Relative to unconditional implementation of $\boldsymbol{\alpha}$, relying on this costly signal improves welfare by

$$k\boldsymbol{\sigma} \cdot \boldsymbol{x} + (\boldsymbol{\pi} - k\boldsymbol{\sigma}) \cdot \boldsymbol{\alpha} - c(\mathcal{S}_\pi^{(k,\boldsymbol{\sigma})}, \boldsymbol{\pi}) - \boldsymbol{\pi} \cdot \boldsymbol{\alpha} = k\boldsymbol{\sigma} \cdot (\boldsymbol{x} - \boldsymbol{\alpha}) - c(\mathcal{S}_\pi^{(k,\boldsymbol{\sigma})}, \boldsymbol{\pi}),$$

which has the same sign as $\boldsymbol{\sigma} \cdot (\boldsymbol{x} - \boldsymbol{\alpha}) - \frac{1}{k}c(\mathcal{S}_\pi^{(k,\boldsymbol{\sigma})}, \boldsymbol{\pi})$. As k converges to zero from above, $k \downarrow 0$, the weighted cost term shrinks. This is because for any $t \in [0, 1]$, $\mathcal{S}_\pi^{(tk,\boldsymbol{\sigma})}$ can be implemented in two steps by an uninformative coin flip that triggers a draw of $\mathcal{S}_\pi^{(k,\boldsymbol{\sigma})}$ with probability t , and otherwise recommends action $\boldsymbol{\alpha}$. By Blackwell monotonicity (C2) and sequential optimality (C4), this implies the monotonic ranking $\frac{1}{tk}c(\mathcal{S}_\pi^{(tk,\boldsymbol{\sigma})}, \boldsymbol{\pi}) \leq \frac{1}{k}c(\mathcal{S}_\pi^{(k,\boldsymbol{\sigma})}, \boldsymbol{\pi})$. In other words, the signal is welfare-enhancing for any $k > 0$ if and only if

$$\boldsymbol{\sigma} \cdot \boldsymbol{x} > \inf_{k>0} \left\{ \boldsymbol{\sigma} \cdot \boldsymbol{\alpha} + \frac{1}{k}c(\mathcal{S}_\pi^{(k,\boldsymbol{\sigma})}, \boldsymbol{\pi}) \right\} = \Phi_\pi(\boldsymbol{\sigma}) \quad \forall \boldsymbol{\sigma} \in \Delta\mathcal{I},$$

warranting the suggested formula for $\bar{\mathcal{A}}^\pi$. And since costs are finite for any small enough $k > 0$, the infimum $\Phi_\pi(\boldsymbol{\sigma})$ is finite-valued.

It remains to show that the function $\Phi_\pi : \Delta\mathcal{I} \rightarrow \mathbb{R}$ is convex. Consider a two-stage process that first flips an uninformative coin that triggers a draw of $\mathcal{S}_\pi^{(k,\boldsymbol{\sigma})}$ with probability t , and otherwise triggers a draw of $\mathcal{S}_\pi^{(k,\boldsymbol{\sigma}')}$. Overall, this process recommends action \mathbf{x} at posterior $t\boldsymbol{\sigma} + (1-t)\boldsymbol{\sigma}'$ with marginal likelihood k . Again, the one-shot implementation is weakly cheaper by (C2) and (C4), $c(\mathcal{S}_\pi^{(k,t\boldsymbol{\sigma}+(1-t)\boldsymbol{\sigma}')} , \boldsymbol{\pi}) \leq tc(\mathcal{S}_\pi^{(k,\boldsymbol{\sigma})} , \boldsymbol{\pi}) + (1-t)c(\mathcal{S}_\pi^{(k,\boldsymbol{\sigma}')} , \boldsymbol{\pi})$. Convexity follows since

$$\begin{aligned} \Phi_\pi(t\boldsymbol{\sigma} + (1-t)\boldsymbol{\sigma}') &= (t\boldsymbol{\sigma} + (1-t)\boldsymbol{\sigma}') \cdot \boldsymbol{\alpha} + \lim_{k \downarrow 0} \frac{1}{k} c(\mathcal{S}_\pi^{(k,t\boldsymbol{\sigma}+(1-t)\boldsymbol{\sigma}')} , \boldsymbol{\pi}) \\ &\leq t \left[\boldsymbol{\sigma} \cdot \boldsymbol{\alpha} + \lim_{k \downarrow 0} \frac{1}{k} c(\mathcal{S}_\pi^{(k,\boldsymbol{\sigma})} , \boldsymbol{\pi}) \right] + (1-t) \left[\boldsymbol{\sigma}' \cdot \boldsymbol{\alpha} + \lim_{k \downarrow 0} \frac{1}{k} c(\mathcal{S}_\pi^{(k,\boldsymbol{\sigma}')} , \boldsymbol{\pi}) \right] \\ &= t\Phi_\pi(\boldsymbol{\sigma}) + (1-t)\Phi_\pi(\boldsymbol{\sigma}'). \quad \square \end{aligned}$$

We now show that under any other prior $\boldsymbol{\rho}$, the set $\bar{\mathcal{A}}^\pi$ contains a point that is implemented unconditionally when present.

Lemma A.15. *For any prior $\boldsymbol{\rho} \in \text{int}(\Delta\mathcal{I})$, unconditional implementation of the payoff vector $\mathbf{a}^\rho = \arg \max \{ \boldsymbol{\rho} \cdot \mathbf{x} \mid \mathbf{x} \in \bar{\mathcal{A}}^\pi \}$ is optimal in any RI problem $(\mathcal{A}', \boldsymbol{\rho}, c)$ where the finite menu $\mathcal{A}' \subseteq \mathcal{A} \cup \bar{\mathcal{A}}^\pi$ contains \mathbf{a}^ρ .*

Proof. By Lemmas A.11 and A.14, the point \mathbf{a}^ρ satisfies $\boldsymbol{\rho} \cdot \mathbf{a}^\rho = \Phi_\pi(\boldsymbol{\rho})$. By contradiction, assume that there exists a learning strategy $\mathcal{S} = \langle \mathcal{A}', \mathbf{q} \rangle$ that achieves higher welfare than unconditional implementation of \mathbf{a}^ρ under prior $\boldsymbol{\rho}$,

$$\sum_{i \in \mathcal{I}} \rho_i \sum_{\mathbf{a} \in \mathcal{A}'} q_i(\mathbf{a}) a_i - c(\mathcal{S}, \boldsymbol{\rho}) > \boldsymbol{\rho} \cdot \mathbf{a}^\rho = \Phi_\pi(\boldsymbol{\rho}).$$

By definition of Φ_π , this implies that there exists $k > 0$ small enough such that $\sum_{i \in \mathcal{I}} \rho_i \sum_{\mathbf{a} \in \mathcal{A}'} q_i(\mathbf{a}) a_i - c(\mathcal{S}, \boldsymbol{\rho}) > \boldsymbol{\rho} \cdot \boldsymbol{\alpha} + \frac{1}{k} c(\mathcal{S}_\pi^{(k,\boldsymbol{\rho})} , \boldsymbol{\pi})$. Rearranging terms, we obtain

$$k \sum_{i \in \mathcal{I}} \rho_i \sum_{\mathbf{a} \in \mathcal{A}'} q_i(\mathbf{a}) a_i + (\boldsymbol{\pi} - k\boldsymbol{\rho}) \cdot \boldsymbol{\alpha} - c(\mathcal{S}_\pi^{(k,\boldsymbol{\rho})} , \boldsymbol{\pi}) - kc(\mathcal{S}, \boldsymbol{\rho}) > \boldsymbol{\pi} \cdot \boldsymbol{\alpha}.$$

This strict inequality implies that unconditional implementation of $\boldsymbol{\alpha}$ is not optimal in RI problem $(\mathcal{A}' \cup \{\boldsymbol{\alpha}\}, \boldsymbol{\pi}, c)$; it is strictly dominated by a sequential strategy where $\mathcal{S}_\pi^{(k,\boldsymbol{\rho})}$ is drawn first, and upon realization \mathbf{x} , the agent draws and follows signal

\mathcal{S} , otherwise she implements α . This contradicts the binary characterization for optimality of ignorance (Corollary 5) since by definition of $\bar{\mathcal{A}}^\pi$, none of the actions $\mathbf{a} \in \mathcal{A}' \subseteq \mathcal{A} \cup \bar{\mathcal{A}}^\pi$ yields a welfare improvement by itself. \square

Building on this, we prove Theorem 4, which states that any anchor can be ‘activated’ by adding the right action to the menu.

Proof of Theorem 4: Suppose first that \mathbf{a} is an anchor action. Since anchor \mathbf{a} is part of the learning-proof menu, there exists a prior $\rho \in \text{int}(\Delta\mathcal{I})$ such that \mathbf{a} is implemented unconditionally in (\mathcal{A}, ρ, c) , and thus denotes the ignorance equivalent. Fix any $k \in (0, 1)$ small enough such that the belief $\rho^+ = \frac{1}{1-k}\pi - \frac{k}{1-k}\rho$ has full support. By Lemmas A.11 and A.14, there exist payoff vectors $\alpha, \mathbf{a}^+ \in \bar{\mathcal{A}}^\rho$ such that $\pi \cdot \alpha = \Phi_\rho(\pi)$ and $\rho^+ \cdot \mathbf{a}^+ = \Phi_\rho(\rho^+)$. Moreover, by Lemma A.15, unconditional implementation of α and \mathbf{a}^+ is optimal in RI problems $(\mathcal{A} \cup \{\mathbf{a}^+, \alpha\}, \pi, c)$ and $(\mathcal{A} \cup \{\mathbf{a}^+, \alpha\}, \rho^+, c)$, respectively. In particular,

$$W(\mathcal{A} \cup \{\mathbf{a}^+\}, \pi, c) \leq W(\mathcal{A} \cup \{\mathbf{a}^+, \alpha\}, \pi, c) = \pi \cdot \alpha.$$

We now show that the agent can achieve this upper bound for welfare by following a learning strategy $\mathcal{S}_\pi^{(t, \rho)}$ that implements action \mathbf{a} at posterior ρ with marginal likelihood t , and otherwise implements action \mathbf{a}^+ at posterior ρ^+ .

To do so, note that for any $k > 0$ small enough, the following two sequential learning strategies are Blackwell equivalent for an agent with prior ρ :

- Draw $\mathcal{S}_\rho^{(k, \pi)}$ and then, conditional on reaching posterior π , draw $\mathcal{S}_\pi^{(t, \rho)}$.
- Flip an uninformative coin that triggers a draw of $\mathcal{S}_\rho^{(k(1-t)/(1-tk), \rho^+)}$ with probability $1 - tk$ and reveals no information otherwise.

Since sequential implementation yields neither gains nor losses by (C4) and (C5), both strategies have the same expected cost,

$$c(\mathcal{S}_\rho^{(k, \pi)}, \rho) + kc(\mathcal{S}_\pi^{(t, \rho)}, \pi) = (1 - tk)c(\mathcal{S}_\rho^{(k(1-t)/(1-tk), \rho^+)}, \rho).$$

Rearranging terms and taking the limit $k \downarrow 0$, we find

$$\begin{aligned}
c(\mathcal{S}_\pi^{(t,\rho)}, \pi) &= (1-t) \lim_{k \downarrow 0} \frac{1-tk}{k(1-t)} c(\mathcal{S}_\rho^{(k(1-t)/(1-tk), \rho^+)}, \rho) - \lim_{k \downarrow 0} \frac{1}{k} c(\mathcal{S}_\rho^{(k, \pi)}, \rho) \\
&\stackrel{(15)}{=} (1-t)[\Phi_\rho(\rho^+) - \rho^+ \cdot \mathbf{a}] - [\Phi_\rho(\pi) - \pi \cdot \mathbf{a}] \\
&= (1-t)\rho^+ \cdot (\mathbf{a}^+ - \mathbf{a}) - \pi \cdot (\boldsymbol{\alpha} - \mathbf{a}).
\end{aligned}$$

Subtracting this cost from the consumption utility $t\rho \cdot \mathbf{a} + (1-t)\rho^+ \cdot \mathbf{a}^+$, we obtain the welfare under signal $\mathcal{S}_\pi^{(t,\rho)}$,

$$\begin{aligned}
&t\rho \cdot \mathbf{a} + (1-t)\rho^+ \cdot \mathbf{a}^+ - [(1-t)\rho^+ \cdot (\mathbf{a}^+ - \mathbf{a}) - \pi \cdot (\boldsymbol{\alpha} - \mathbf{a})] \\
&= \underbrace{[t\rho + (1-t)\rho^+] \cdot \mathbf{a}}_{=\pi} + \pi \cdot (\boldsymbol{\alpha} - \mathbf{a}) = \pi \cdot \boldsymbol{\alpha}.
\end{aligned}$$

As a result, it is optimal for the agent to implement \mathbf{a} with positive probability $t > 0$ in RI problem $(\mathcal{A} \cup \{\mathbf{a}^+\}, \pi, c)$.

Conversely, suppose $\mathbf{a} \in \mathcal{A} \setminus \bar{\mathcal{A}}$ is not an anchor. By [Corollary 4\(d\)](#), there exists a prior ρ such that $\mathbf{a} < \boldsymbol{\alpha}^{(\mathcal{A}, \rho, c)}$. By construction of the ignorance equivalent in [Theorem 1](#), \mathbf{a} is thus strictly dominated by any optimal learning strategy \mathcal{S} of RI problem (\mathcal{A}, ρ, c) . In particular, as long as the menu contains \mathcal{A} , the agent is always strictly better off following strategy \mathcal{S} rather than implementing \mathbf{a} , no matter her belief. By sequential optimality [\(C4\)](#), implementing \mathbf{a} is also strictly suboptimal among one-shot strategies. \square

Proof of [Corollary 6](#): The agent's willingness to pay is equal to the expected change in welfare, $\sum_{s \in S} (\pi \cdot \mathbf{q}(s)) W(\mathcal{A}, \pi^s, c) - W(\mathcal{A}, \pi, c)$. The claim then follows because [Theorem 2](#) implies that for all priors $\rho \in \text{int}(\Delta \mathcal{I})$,

$$W(\mathcal{A}, \rho, c) = \rho \cdot \boldsymbol{\alpha}^{(\mathcal{A}, \rho, c)} \stackrel{2(b)}{=} \rho \cdot \boldsymbol{\alpha}^{(\bar{\mathcal{A}}, \rho, c)} = W(\bar{\mathcal{A}}, \rho, c) \stackrel{2(a)}{=} \max_{\mathbf{a} \in \bar{\mathcal{A}}} \rho \cdot \mathbf{a}.$$

The result generalizes to the boundary of the simplex by continuity of the welfare function [Lemma A.7](#). \square

Proof of [Corollary 7](#): Suppose the principal offers the menu $\mathcal{A}^0 \cup \{\mathbf{a}^s \mid s \in S\}$ with $\mathbf{a}^s = \arg \max \{\pi^s \cdot \mathbf{a} \mid \mathbf{a} \in \bar{\mathcal{A}}^\pi\}$. Since none of the added options are welfare-enhancing, the agent's welfare is bounded above by $\pi \cdot \boldsymbol{\alpha}$, and we now show that the

agent achieves this upper bound by drawing \mathcal{S} and selecting action \mathbf{a}^s after realization s . Doing so yields consumption utility $\boldsymbol{\pi}^s \cdot \mathbf{a}^s$ with marginal likelihood $q(s) = \boldsymbol{\pi} \cdot \mathbf{q}(s)$ for a cost of $c(\mathcal{S}, \boldsymbol{\pi})$. By [Lemmas A.11](#) and [A.14](#), we can write

$$\begin{aligned} \sum_{s \in S} q(s)(\boldsymbol{\pi}^s \cdot \mathbf{a}^s) - c(\mathcal{S}, \boldsymbol{\pi}) &= \sum_{s \in S} q(s)\Phi_{\boldsymbol{\pi}}(\boldsymbol{\pi}^s) - c(\mathcal{S}, \boldsymbol{\pi}) \\ &= \sum_{s \in S} q(s) \left(\boldsymbol{\pi}^s \cdot \boldsymbol{\alpha} + \lim_{k \downarrow 0} \frac{1}{k} c(\mathcal{S}_{\boldsymbol{\pi}}^{(k, \boldsymbol{\pi}^s)}, \boldsymbol{\pi}) \right) - c(\mathcal{S}, \boldsymbol{\pi}) \\ &= \boldsymbol{\pi} \cdot \boldsymbol{\alpha} + \lim_{k \downarrow 0} \frac{1}{k} \left(\sum_{s \in S} q(s) c(\mathcal{S}_{\boldsymbol{\pi}}^{(k, \boldsymbol{\pi}^s)}, \boldsymbol{\pi}) \right) - c(\mathcal{S}, \boldsymbol{\pi}) \end{aligned} \tag{16}$$

The term in parentheses captures the cost of a sequential strategy where an assistant first draws an uninformative signal that returns each $s \in S$ with unconditional probability $q(s)$, and then flips an informative coin that yields heads at posterior $\boldsymbol{\pi}^s$ with probability k . Suppose the agent announces s if the coin returns heads and says nothing otherwise. The assistant's announcement is then Blackwell equivalent to a strategy where the agent draws \mathcal{S} with probability k and maintains prior $\boldsymbol{\pi}$ otherwise. Since the agent is indifferent across all sequential strategies [\(C4\)](#) and [\(C5\)](#) and costs are Blackwell-monotone [\(C2\)](#), the term $\sum_{s \in S} q(s) c(\mathcal{S}_{\boldsymbol{\pi}}^{(k, \boldsymbol{\pi}^s)}, \boldsymbol{\pi})$ is equal to $kc(\mathcal{S}, \boldsymbol{\pi})$. From [Equation \(16\)](#), it follows that \mathcal{S} achieves $\boldsymbol{\pi} \cdot \boldsymbol{\alpha}$ and is thus indeed optimal for the agent. And since the agent receives the payoff $\boldsymbol{\pi} \cdot \boldsymbol{\alpha} = W(\mathcal{A}^0, \boldsymbol{\pi}, c)$ that he is guaranteed in any offer $A \supseteq \mathcal{A}^0$, the principal is making minimal transfers. \square

References

- Roc Armenter, Michèle Müller-Ippen, and Zachary Stangebye. Geometric methods for finite rational inattention. Available at <https://www.philadelphiafed.org/-/media/frbp/assets/working-papers/2021/wp21-30.pdf>, 2021.
- Dirk Bergemann and Juuso Välimäki. Information acquisition and efficient mechanism design. *Econometrica*, 70(3):1007–1033, 2002. doi: <https://doi.org/10.1111/1468-0262.00317>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/1468-0262.00317>.
- David Blackwell. Equivalent comparisons of experiments. *The annals of mathematical statistics*, pages 265–272, 1953.

- Alexander W. Bloedel and Weijie Zhong. The cost of optimally acquired information. *Unpublished working paper*, 2020. Available at <https://drive.google.com/file/d/1HC5WxR6p5sB-6FvzXwf73rrFzx39UXYw/view?usp=sharing>.
- Jonathan M. Borwein and Qiji J. Zhu. *Techniques of Variational Analysis*. CMS Books in Mathematics. Springer New York, New York, NY, 2005. ISBN 0387242988.
- Andrew Caplin, Mark Dean, and John Leahy. Rational inattention, optimal consideration sets, and stochastic choice. *Review of Economic Studies*, 86(3):1061–1094, 07 2018. ISSN 0034-6527. doi: 10.1093/restud/rdy037. URL <https://doi.org/10.1093/restud/rdy037>.
- Andrew Caplin, Mark Dean, and John Leahy. Rationally inattentive behavior: Characterizing and generalizing shannon entropy. *Journal of Political Economy*, 130(6):1676–1715, 2022. doi: 10.1086/719276. URL <https://doi.org/10.1086/719276>.
- D.L. Cohn. *Measure Theory*. Birkhäuser Boston, 1997. ISBN 9780817630034. URL <https://books.google.com/books?id=vRxV2FwJvoAC>.
- Jacques Crémer and Fahad Khalil. Gathering information before signing a contract. *The American Economic Review*, 82(3):566–578, 1992. ISSN 00028282. URL <http://www.jstor.org/stable/2117322>.
- Tommaso Denti, Massimo Marinacci, and Aldo Rustichini. Experimental cost of information. *American Economic Review*, 112(9):3106–23, September 2022. doi: 10.1257/aer.20210879. URL <https://www.aeaweb.org/articles?id=10.1257/aer.20210879>.
- Matthew Gentzkow and Emir Kamenica. Costly persuasion. *American Economic Review*, 104(5):457–462, 2014.
- Benjamin M. Hébert and Michael Woodford. Rational inattention when decisions take time. Technical report, National Bureau of Economic Research, 2019.
- Benjamin M. Hébert and Michael Woodford. Neighborhood-based information costs. Technical report, National Bureau of Economic Research, 2020.
- Bengt Holmström and Paul Milgrom. Regulating trade among agents. *Journal of Institutional and Theoretical Economics (JITE) / Zeitschrift für die gesamte*

- Staatswissenschaft*, 146(1):85–105, 1990. ISSN 09324569. URL <http://www.jstor.org/stable/40751306>.
- Hideshi Itoh. Incentives to help in multi-agent situations. *Econometrica*, 59(3): 611–636, 1991. ISSN 00129682, 14680262. URL <http://www.jstor.org/stable/2938221>.
- Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. *American Economic Review*, 101(6):2590–2615, 2011.
- David M Kreps. A representation theorem for "preference for flexibility". *Econometrica: Journal of the Econometric Society*, pages 565–577, 1979.
- Bartosz Maćkowiak, Filip Matějka, and Mirko Wiederholt. Survey: Rational inattention, a disciplined behavioral model. Working paper, Center for Economic and Policy Research, October 2018.
- Filip Matějka and Alisdair McKay. Rational inattention to discrete choices: A new foundation for the multinomial logit model. *American Economic Review*, 105(1): 272–98, January 2015. doi: 10.1257/aer.20130047. URL <http://www.aeaweb.org/articles?id=10.1257/aer.20130047>.
- Jeffrey Mensch. Cardinal representations of information. *Unpublished working paper*, 2018. Available at <https://ssrn.com/abstract=3148954>.
- Stephen Morris and Philipp Strack. The Wald problem and the relation of sequential sampling and ex-ante information costs. *Unpublished working paper*, 2019. Available at <https://ssrn.com/abstract=2991567>.
- Luciano Pomatto, Philipp Strack, and Omer Tamuz. The cost of information. *arXiv preprint arXiv:1812.04211*, 2018.
- Christopher A. Sims. Implications of rational inattention. *Journal of Monetary Economics*, 50(3):665–690, 2003.
- Elias Tsakas. Robust scoring rules. *Theoretical Economics*, 15(3):955–987, 2020. URL <https://econtheory.org/ojs/index.php/te/article/view/20200955/0>.

Gerhard Winkler. Extreme points of moment sets. *Mathematics of Operations Research*, 13(4):581–587, 1988. ISSN 0364765X, 15265471. URL <http://www.jstor.org/stable/3689944>.

Nathan Yoder. Designing incentives for heterogeneous researchers. *Journal of Political Economy*, 130(8):2018–2054, 2022. doi: 10.1086/720072. URL <https://doi.org/10.1086/720072>.

Weijie Zhong. Optimal dynamic information acquisition. *Unpublished working paper*, 2019. Available at https://wjzhong.com/workingpapers/info_acquisition/Dynamic_info_acquisition_main.pdf.

B Online Appendix

B.1 Additional Explanations

B.1.1 Duality

Proof of Lemma A.10: Let V denote the Banach space of all finite signed measures over $\Delta\mathcal{I}$, and $V_+ = \{v \in V \mid v \geq 0\}$ the subset of nonnegative measures. Along with the total variation norm $\|\cdot\|_V$, this forms a Banach space [Cohn, 1997, p.219]. We endow $\mathbb{R}^{\mathcal{I}}$ with the ℓ_1 norm $\|\cdot\|_1$ for convenience, and let $M : V \rightarrow \mathbb{R}^{\mathcal{I}}$ denote the linear map that assigns to each measure ν the Lebesgue integral $M\nu = \int_{\Delta\mathcal{I}} \rho d\nu$. The map M is bounded since $\int_{\Delta\mathcal{I}} \rho_i d\nu \in [-\nu^-(\Delta\mathcal{I}), \nu^+(\Delta\mathcal{I})]$ for each i and hence $\|M\nu\|_1 \leq |\mathcal{I}|(\nu^+(\Delta\mathcal{I}) + \nu^-(\Delta\mathcal{I})) = |\mathcal{I}|\|\nu\|_V$, where ν^+ and $-\nu^-$ denote the upper and lower variation respectively. For any set B , we further define the indicator function $\chi_B(b)$ that is 0 whenever $b \in B$ and ∞ otherwise. This function is lower semi-continuous and convex whenever B is closed and convex.

We now define two functions $f : V \rightarrow \mathbb{R} \cup \{+\infty\}$ and $g : \mathbb{R}^{\mathcal{I}} \times \mathbb{R} \rightarrow \mathbb{R} \cup \{+\infty\}$ as

$$f(\nu) = \int_{\Delta\mathcal{I}} \Phi(\rho) d\nu + \chi_{V_+}(\nu) \quad \text{and} \quad g(\mathbf{x}) = \chi_{\{\boldsymbol{\pi}\}}(\mathbf{x}).$$

Since convexity and lower semi-continuity are preserved under addition, f and g are both convex and lower-semicontinuous. Moreover, in the notation of Borwein and Zhu [2005, Theorem 4.4.3],

$$\mathbf{0} \in \text{core}(\text{dom}g - M\text{dom}f) = \text{core}(\{\boldsymbol{\pi}\} - [0, \infty)^{\mathcal{I}+1}),$$

and therefore strong duality holds:

$$\inf_{\nu \in V} f(\nu) + g(M\nu) = \sup_{\langle \mathbf{x}^*, \cdot \rangle \in (\mathbb{R}^{\mathcal{I}})^*} -f^*(M^*\langle \mathbf{x}^*, \cdot \rangle) - g^*(-\langle \mathbf{x}^*, \cdot \rangle), \quad (17)$$

where $(\mathbb{R}^{\mathcal{I}})^*$ denotes the space of all linear forms over $\mathbb{R}^{\mathcal{I}}$, M^* is the adjoint operator

and the convex conjugates assign to each linear form a supremum

$$\begin{aligned}
f^*(\langle \nu^*, \cdot \rangle) &= \sup_{v \in V} \langle \nu^*, v \rangle - f(v) = \sup_{v \in V_+} \langle \nu^*, v \rangle - \int_{\Delta \mathcal{I}} \Phi(\boldsymbol{\rho}) dv \\
&= \begin{cases} 0 & \text{if } \langle \nu^*, v \rangle \leq \int_{\Delta \mathcal{I}} \Phi(\boldsymbol{\rho}) dv \quad \forall v \in V_+ \\ \infty & \text{otherwise,} \end{cases} \\
g^*(\langle \mathbf{x}^*, \cdot \rangle) &= \sup_{\mathbf{x} \in \mathbb{R}^{\mathcal{I}}} \langle \mathbf{x}^*, \mathbf{x} \rangle - g(\mathbf{x}) = \langle \mathbf{x}^*, \boldsymbol{\pi} \rangle.
\end{aligned}$$

Using these expressions, we can equivalently write the duality (17) as

$$\begin{cases} \inf_{\nu \in V_+} & \int_{\Delta \mathcal{I}} \Phi(\boldsymbol{\rho}) dv \\ \text{s.t.} & \int_{\Delta \mathcal{I}} \boldsymbol{\rho} dv = \boldsymbol{\pi} \end{cases} = \begin{cases} \sup_{\langle \mathbf{x}^*, \cdot \rangle \in (\mathbb{R}^{\mathcal{I}})^*} & \langle \mathbf{x}^*, \boldsymbol{\pi} \rangle \\ \text{s.t.} & \langle M^* \langle \mathbf{x}^*, \cdot \rangle, v \rangle \leq \int_{\Delta \mathcal{I}} \Phi(\boldsymbol{\rho}) dv \quad \forall v \in V_+. \end{cases}$$

The constraint on the left side implies in particular that

$$\nu(\Delta \mathcal{I}) = \int_{\Delta \mathcal{I}} 1 d\nu = \mathbf{1} \cdot \int_{\Delta \mathcal{I}} \boldsymbol{\rho} dv = \mathbf{1} \cdot \boldsymbol{\pi} = 1,$$

and hence it is without loss of generality to restrict attention to probability measures $\nu \in \Delta \Delta \mathcal{I}$, making this infimum equal to p^* from the claim.

As for the right side, the definition of the adjoint states that $\langle M^* \langle \mathbf{x}^*, \cdot \rangle, v \rangle = \langle \mathbf{x}^*, Mv \rangle$, and since any linear form over $\mathbb{R}^{\mathcal{I}}$ can be expressed as a scalar product with a vector in $\mathbb{R}^{\mathcal{I}}$, we can also write the duality (17) as

$$p^* = \begin{cases} \sup_{\mathbf{a} \in \mathbb{R}^{\mathcal{I}}} & \mathbf{a} \cdot \boldsymbol{\pi} \\ \text{s.t.} & \mathbf{a} \cdot Mv = \mathbf{a} \cdot \int_{\Delta \mathcal{I}} \boldsymbol{\rho} dv \leq \int_{\Delta \mathcal{I}} \Phi(\boldsymbol{\rho}) dv \quad \forall v \in V_+. \end{cases}$$

We rewrite the constraint as

$$0 \leq \int_{\Delta \mathcal{I}} (\Phi(\boldsymbol{\rho}) - \mathbf{a} \cdot \boldsymbol{\rho}) dv \quad \forall v \in V_+,$$

and make a few simplifying observations: First, the inequality is unaffected by positive scaling of v , so it is without loss of generality to restrict attention to probability measures $v \in \Delta \Delta \mathcal{I}$. Moreover, the right side is most binding when v is a degenerate

distribution, so it is also without loss of generality to simply write (17) as

$$p^* = \begin{cases} \sup_{\mathbf{a} \in \mathbb{R}^I} & \mathbf{a} \cdot \boldsymbol{\pi} \\ \text{s.t.} & 0 \leq \Phi(\boldsymbol{\rho}) - \mathbf{a} \cdot \boldsymbol{\rho} \quad \forall \boldsymbol{\rho} \in \Delta\mathcal{I}, \end{cases}$$

which is now equal to d^* from the claim.

Moreover, the Fenchel duality theorem also asserts that whenever $p^* = d^*$ is finite, then the supremum is attained. \square

B.1.2 Posterior-separable Costs

Although our proof of [Theorem 1](#) relies on prior-concavity (C2) and sequential learning (C4), a unique ignorance equivalent exists for a larger class of cost functions.²² Indeed, consider any posterior-separable cost of the form

$$c(\langle S, \mathbf{q} \rangle, \boldsymbol{\pi}) = \sum_{s \in \mathcal{S}} (\boldsymbol{\pi} \cdot \mathbf{q}(s)) \phi^\pi(\boldsymbol{\pi}^s) - \phi^\pi(\boldsymbol{\pi}), \quad (\text{PS})$$

which differs from [Equation \(UPS+\)](#) in that the convex and differentiable potential ϕ^π can depend on the prior, and does not need to induce prior-concavity.

Let \mathcal{S} denote a learning strategy that is optimal in RI problem $(\mathcal{A}, \boldsymbol{\pi}, c)$. By [Caplin et al. \[2022, Lemma 1\]](#), there exists a vector $\boldsymbol{\theta} \in \mathbb{R}^I$ such that, for any \mathbf{a} that is chosen with positive probability and for any optimal posterior $\boldsymbol{\pi}^s$,

$$\boldsymbol{\rho} \cdot \mathbf{a}' - \phi^\pi(\boldsymbol{\rho}) - \boldsymbol{\rho} \cdot \boldsymbol{\theta} \leq \boldsymbol{\pi}^s \cdot \mathbf{a} - \phi^\pi(\boldsymbol{\pi}^s) - \boldsymbol{\pi}^s \cdot \boldsymbol{\theta} \quad \forall \boldsymbol{\rho} \in \Delta\mathcal{I}, \forall \mathbf{a}' \in \mathcal{A}. \quad (18)$$

Moreover, the existence of such a $\boldsymbol{\theta}$ guarantees optimality of the signal \mathcal{S} .

By symmetry, the value on right side of the inequality needs to be the same, independently of the choice of the optimal action \mathbf{a} or the optimal posterior $\boldsymbol{\pi}^s$. Moreover, since the inequality also holds for any $\boldsymbol{\theta} + k\mathbf{1}$ with $k \in \mathbb{R}$, it is without loss of generality to normalize the right side to zero. Under this normalization, we can also express welfare succinctly as

$$W(\mathcal{A}, \boldsymbol{\pi}, c) = \sum_{\mathbf{a} \in \mathcal{A}} (\boldsymbol{\pi} \cdot \mathbf{q}(\mathbf{a})) (\boldsymbol{\pi}^{\mathbf{a}} \cdot \mathbf{a} - \phi^\pi(\boldsymbol{\pi}^{\mathbf{a}})) + \phi^\pi(\boldsymbol{\pi}) = \boldsymbol{\pi} \cdot \boldsymbol{\theta} + \phi^\pi(\boldsymbol{\pi}).$$

²²We thank an anonymous referee for pointing this out to us.

The ignorance equivalent can be obtained from this optimal Lagrangian vector as

$$\boldsymbol{\alpha} := \nabla \phi^\pi(\boldsymbol{\pi}) + \boldsymbol{\theta} + (\phi^\pi(\boldsymbol{\pi}) - \boldsymbol{\pi} \cdot \nabla \phi^\pi(\boldsymbol{\pi}))\mathbf{1}. \quad (19)$$

Indeed, this payoff vector $\boldsymbol{\alpha}$ achieves the same welfare as the optimal signal since

$$\boldsymbol{\pi} \cdot \boldsymbol{\alpha} = \boldsymbol{\pi} \cdot \boldsymbol{\theta} + \phi^\pi(\boldsymbol{\pi}) = W(\mathcal{A}, \boldsymbol{\pi}, c). \quad (20)$$

Moreover, the addition of $\boldsymbol{\alpha}$ to the menu does not introduce any new learning opportunities since

$$\boldsymbol{\rho} \cdot \boldsymbol{\alpha} - \phi^\pi(\boldsymbol{\rho}) - \boldsymbol{\rho} \cdot \boldsymbol{\theta} = \phi^\pi(\boldsymbol{\pi}) - \phi^\pi(\boldsymbol{\rho}) - (\boldsymbol{\pi} - \boldsymbol{\rho}) \cdot \nabla \phi^\pi(\boldsymbol{\pi}) \leq 0 \quad (21)$$

by concavity of ϕ . In other words, Equation (18) holds with the same $\boldsymbol{\theta}$ for the degenerate strategy that puts full weight on $\boldsymbol{\alpha}$, making this optimal even in the large menu $\mathcal{A} \cup \{\boldsymbol{\alpha}\}$. Moreover, both constraints (20) and (21) have to be satisfied, making the payoff vector in (19) the unique candidate.

B.2 Examples

B.2.1 Non-Existence of an Ignorance Equivalence

In this section, we show by example that an ignorance equivalent may fail to exist when the cost function does not satisfy (C4).

In particular, we consider prior-independent costs

$$c(\langle S, \mathbf{q} \rangle, \boldsymbol{\pi}) = \left[\frac{1}{I} \sum_{i \in \mathcal{I}} H(q_i) - H \left(\frac{1}{I} \sum_{i \in \mathcal{I}} q_i \right) \right]^2,$$

where $H(p) = \sum_{s \in \mathcal{S}} p(s) \ln(p(s))$ denotes the Shannon entropy of probability mass function p . This cost function measures the square of the mutual information between the conditional probabilities q_i and the mean probability $\frac{1}{I} \sum_{i \in \mathcal{I}} q_i$. It is well known that mutual information is continuous and strictly increasing in the Blackwell order. By taking the unconditional mean rather than (as is standard) weighting each term by π_i , the cost function becomes prior-independent and therefore trivially prior-concave. In addition, let $\mathbf{a}^1 = (1, 0)$, $\mathbf{a}^2 = (0, 1)$, $\mathcal{A} = \{\mathbf{a}^1, \mathbf{a}^2\}$ and $\boldsymbol{\pi} = (\frac{1}{2}, \frac{1}{2})$. The exact

choice of the menu and the prior matter only in so far as we want to guarantee that some learning is optimal in RI problem $(\mathcal{A}, \boldsymbol{\pi}, c)$. We claim that this RI problem does not admit an ignorance equivalent.

Indeed, let $\mathcal{S} = \langle \mathcal{A}, \mathbf{q} \rangle$ denote an optimal learning strategy, with expected state-wise consumption of $\mathbf{a}^{\mathcal{S}}$. For convenience, we determine \mathcal{S} using Mathematica (see code at the end of this section) and verify that it involves a strictly positive amount of learning, $c(\mathcal{S}, \boldsymbol{\pi}) > 0$. By definition, the expected utility of any candidate ignorance equivalent $\boldsymbol{\alpha}$ has to be equal to the net utility as \mathcal{S} ,

$$\boldsymbol{\pi} \cdot \boldsymbol{\alpha} = W(\{\boldsymbol{\alpha}\}, \boldsymbol{\pi}, c) = W(\mathcal{A}, \boldsymbol{\pi}, c) = \boldsymbol{\pi} \cdot \mathbf{a}^{\mathcal{S}} - c(\mathcal{S}, \boldsymbol{\pi}).$$

At the same time, an agent facing RI problem $(\mathcal{A} \cup \{\boldsymbol{\alpha}\}, \boldsymbol{\pi}, c)$ could follow signal $\tilde{\mathcal{S}} = \langle \mathcal{A} \cup \{\boldsymbol{\alpha}\}, \tilde{\mathbf{q}} \rangle$ with $q_i(\boldsymbol{\alpha}) \equiv t \in (0, 1)$ and $\mathbf{q}(\mathbf{a}^k) \equiv (1 - t)\mathbf{q}(\mathbf{a}^k)$. This strategy is Blackwell equivalent to a two-stage strategy where the agent first tosses an uninformative coin. The coin returns heads with probability t , prompting the agent to implement $\boldsymbol{\alpha}$, and tails with probability $1 - t$, prompting her to follow \mathcal{S} . Beliefs are unaffected by the coin toss, and so the agent's expected consumption utility is equal to $t\boldsymbol{\pi} \cdot \boldsymbol{\alpha} + (1 - t)\boldsymbol{\pi} \cdot \mathbf{a}^{\mathcal{S}}$. However, the cost of strategy $\tilde{\mathcal{S}}$ is strictly below that of its two-stage implementation, with $c(\tilde{\mathcal{S}}, \boldsymbol{\pi}) = (1 - t)^2 c(\mathcal{S}, \boldsymbol{\pi})$. Consequently, the net utility of $\tilde{\mathcal{S}}$ exceeds $W(\mathcal{A}, \boldsymbol{\pi}, c)$, implying that adding $\boldsymbol{\alpha}$ to the menu is strictly welfare enhancing,

$$\begin{aligned} W(\mathcal{A} \cup \{\boldsymbol{\alpha}\}, \boldsymbol{\pi}, c) &\geq t\boldsymbol{\pi} \cdot \boldsymbol{\alpha} + (1 - t)\boldsymbol{\pi} \cdot \mathbf{a}^{\mathcal{S}} - (1 - t)^2 c(\mathcal{S}, \boldsymbol{\pi}) \\ &= tW(\{\boldsymbol{\alpha}\}, \boldsymbol{\pi}, c) + (1 - t)W(\mathcal{A}, \boldsymbol{\pi}, c) + t(1 - t)c(\mathcal{S}, \boldsymbol{\pi}) \\ &> W(\mathcal{A}, \boldsymbol{\pi}, c). \end{aligned}$$

The incompatibility between the two properties of Definition 1 is linked to the fact that c violates sequential-learning proofness (C4). This means that even though \mathcal{S} dominates $\boldsymbol{\alpha}$, the agent is not able to realize the full welfare potential of $(\mathcal{A} \cup \{\boldsymbol{\alpha}\}, \boldsymbol{\pi}, c)$ without relying on the ‘moderate’ payoffs $\boldsymbol{\alpha}$.²³

²³To avoid confusion: The issue is not that $\tilde{\mathcal{S}}$ is strictly cheaper than the two-step implementation described above. (C4) allows for that. The issue is that \mathcal{S} is strictly more expensive than a two-stage implementation where the agent first follows $\tilde{\mathcal{S}}$ and replaces action $\boldsymbol{\alpha}$ with another instance of \mathcal{S} . The former costs $c(\mathcal{S}, \boldsymbol{\pi})$ while the latter merely costs $(1 - t)^2 c(\mathcal{S}, \boldsymbol{\pi}) + tc(\mathcal{S}, \boldsymbol{\pi})$.

Code: Non-existence under violation of (C4).

```
1 Clear["Global`*"];
```

We consider a setup with two states and two actions:

```
2 a1={1,0};
```

```
3 a2={0,1};
```

The prior is equal to

```
4 pi1=0.5;
```

```
5 pi={pi1,1-pi1};
```

A learning strategy is written as a matrix where entry (i,a) denotes the likelihood of implementing action a conditional on state i. Costs are given by the prior-independent function

```
6 p={1/2,1/2};
```

```
7 c[q_?(MatrixQ[#,NumericQ]&),pi_]:= Sum[ If[ q[[i,a]]>0, p[[i]]q[[i,a]]
  Log[q[[i,a]] / p.q[[;,a]] , 0] , {i,1,Dimensions[q][[1]]} ,
  {a,1,Dimensions[q][[2]]} ]^2
```

This function denotes the squared reduction in Shannon entropy, but relative to the fixed state distribution p rather than the prior pi. This has the advantage of ensuring that the function trivially satisfies (C1) and (C2).

We now compute the optimal signal:

```
8 q={{q1a1,1-q1a1},{1-q2a2,q2a2}};
```

```
9 W[q_?(MatrixQ[#,NumericQ]&),pi_] := pi.Total[{a1,a2}*Transpose[q]]-c[q,pi]
```

```
10 NMaximize[{W[q,pi],0<=q1a1<=1,0<=q2a2<=1},{q1a1,q2a2]}
```

```
>> {0.777046, {q1a1 -> 0.855844, q2a2 -> 0.855844}}
```

```
11 qopt={{q1a1,1-q1a1},{1-q2a2,q2a2}}/.%[[2]];
```

This involves a strictly positive amount of learning,

```
12 c[qopt,c]>0
```

```
>> True
```

B.2.2 Failure of Self-Selection

To illustrate the role that (C2) plays in our results, we here show that self-selection (Corollary 3) fails for a concrete example with a cost function that satisfies all properties save for prior-concavity. Self-selection is crucial for our subsequent results, as the construction of the learning-proof menu rests upon the absence of welfare gains.

Specifically, we consider a uniformly posterior-separable cost function with the convex and differentiable potential $\phi(\boldsymbol{\pi}) = -\sum_{i \in \mathcal{I}} \pi_i(1 - \pi_i)$. We study a three-state problem with two actions $\mathcal{A} = \{(1, 0, 0), (0, 1, 0)\}$ and a uniform prior $\boldsymbol{\pi} = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$. To streamline exposition, we verify all claims algebraically with Mathematica and provide the code at the end of this section.

Using the tools from Online Appendix B.1.2, we locate the unique ignorance equivalent equivalent for this posterior separable function at $\boldsymbol{\alpha} = (\frac{31}{24}, \frac{31}{24}, \frac{19}{24})$. We then show that adding this action $\boldsymbol{\alpha}$ to the menu is strictly welfare-enhancing for an agent who puts a higher likelihood on state 3, with prior $\boldsymbol{\rho} = (\frac{1}{10}, \frac{1}{10}, \frac{8}{10})$. In particular,

$$W(\mathcal{A}, \boldsymbol{\rho}, c) = 0.18 < \boldsymbol{\rho} \cdot \boldsymbol{\alpha} = 0.225 \leq W(\mathcal{A} \cup \{\boldsymbol{\alpha}\}, \boldsymbol{\rho}, c),$$

thus contradicting Corollary 3. Finally, we show that that this cost function violates (only) prior-concavity (C2). Intuitively, prior-concavity allows us to connect global welfare gains to local welfare gains in the ‘necessity of dominance’ argument (Section 3). Without it, there may be strict welfare gains at prior $\boldsymbol{\rho}$ even though there are none for local priors around $\boldsymbol{\pi}$.

Code: Violation of Corollary 3 without (C2).

```
1 Clear["Global`*"];
```

There are three states and two actions with payoffs

```
2 a1={1,0,0};
```

```
3 a2={0,1,0};
```

The UPS potential is given by

```
4 phi[{r1_,r2_,r3_}]:=Sum[-r (1-r),{r,{r1,r2,r3}}]
```

For notational convenience, we explicitly define its gradient as

```
5 Dphi[{p1_,p2_,p3_}]=D[phi[{p1,p2,p3}],{{p1,p2,p3},1}];
```

and verify that the potential is convex by checking for positive Eigenvalues of the

Hessian,

6 Eigenvalues[D[$\phi[\{p_1, p_2, p_3\}], \{p_1, p_2, p_3\}, 2]$]
» {2,2,2}

We claim that the ignorance equivalent at

7 $\pi = \{1/3, 1/3, 1/3\}$;

is equal to

8 $\alpha = \{5/8, 5/8, 1/8\}$;

and hence using formula (19),

9 $\theta = \alpha - D\phi[\pi] - \{1, 1, 1\}(\phi[\pi] - \pi \cdot D\phi[\pi])$
» {31/24, 31/24, 19/24}

This satisfies the constraints (18) since for any beliefs r and any available action a , the expression $r \cdot a - \phi[r] - r \cdot \theta$ is maximized at 0.

10 $r = \{r_1, r_2, r_3\}$;

11 Maximize[$\{r \cdot a_1 - \phi[r] - r \cdot \theta, \{0 \leq r_1, 0 \leq r_2, 0 \leq r_3, r_1 + r_2 + r_3 == 1\}\}, \{r_1, r_2, r_3\}$]

12 Maximize[$\{r \cdot a_2 - \phi[r] - r \cdot \theta, \{0 \leq r_1, 0 \leq r_2, 0 \leq r_3, r_1 + r_2 + r_3 == 1\}\}, \{r_1, r_2, r_3\}$]

» {0, { $r_1 \rightarrow 7/12, r_2 \rightarrow 1/12, r_3 \rightarrow 1/3$ }}

» {0, { $r_1 \rightarrow 1/12, r_2 \rightarrow 7/12, r_3 \rightarrow 1/3$ }}

For this reason, we know that we have found the correct ignorance equivalent.

We now show that adding this α to the menu is strictly welfare enhancing for an agent with prior

13 $\rho = \{1/10, 1/10, 8/10\}$;

Sticking to the posterior-based approach, we know we can find $W(\{a_1, a_2\}, \rho, c)$ by maximizing over Bayes-plausible posteriors. In particular, we maximize $q(r \cdot a_1 - \phi[r]) + (1-q)(s \cdot a_2 - \phi[s])$, which is the convex combination over the agent's net utility at r and s , weighed according to q . The constraints ensure that both posteriors are valid, and that the average posterior equals the prior ρ . Doing so determines the maximal net utility of the agent, and we obtain actual welfare by adding $\phi[\rho]$ back in.

14 $r = \{r_1, r_2, r_3\}$;

15 $s = \{s_1, s_2, s_3\}$;

16 Maximize[$\{q(r \cdot a_1 - \phi[r]) + (1-q)(s \cdot a_2 - \phi[s]), \{0 \leq r_1, 0 \leq r_2, 0 \leq r_3, r_1 + r_2 + r_3 == 1, 0 \leq s_1, 0 \leq s_2, 0 \leq s_3, s_1 + s_2 + s_3 == 1, 0 \leq q \leq 1, q r + (1-q)s ==$

```

    rho}}, {r1, r2, r3, q, s1, s2, s3}]
17 W=%[[1]]+phi[rho] (* In mathematica, % refers to the last computed output. *)
>> {13/25, {r1->1/5, r2->0, r3->4/5, q->1/2, s1->0, s2->1/5, s3->4/5}}
>> 9/50

```

In particular, the agent's welfare without access to α is equal to

```

18 W //N
>> 0.18

```

However, if α is available to the agent, she can in particular choose to implement it unconditionally, which yields a lower bound for $W(\{a1, a2, \alpha\}, \rho, c)$,

```

19 rho.alpha//N
>> 0.225

```

This shows that adding α affects the welfare of an agent who differs only in prior belief, contrary to the result we obtain in [Corollary 3](#).

Lastly, we verify that prior-concavity is the only property not satisfied by the UPS cost function considered here. Without loss of generality, we denote a signal by a list of state-dependent probabilities, where $qs=\{qs1, qs2, qs3\}$ denotes the likelihood for realization $s \in \{x, y, z\}$ conditional on each of the states $i \in \{1, 2, 3\}$,

```

20 c[q_, pi_] := Sum[(pi.qs) phi[(pi*qs)/pi.qs], {qs, q}] - phi[pi]

```

Continuity (C1) holds by continuity of the expectation and that of ϕ . Blackwell monotonicity holds because ϕ is convex (as verified above). [Bloedel and Zhong \[2020\]](#) show that indifference to sequential information acquisition (C4) and (C5) hold for any UPS functions, including this one. The differentiability of the potential further implies that ties are broken through learning (C3), as detailed in [Lemma A.2](#).

However, prior-concavity does not hold. We illustrate this by focusing on the optimal signal under prior π_i ,

```

21 q={r1, r2, r3}/(2pi) /.Maximize[{r.a1-phi[r]-r.theta,
    {0<=r1, 0<=r2, 0<=r3, r1+r2+r3==1}}, {r1, r2, r3}][[2]], {r1, r2, r3}/(2pi)}
    /.Maximize[{r.a2-phi[r]-r.theta, {0<=r1, 0<=r2, 0<=r3, r1+r2+r3==1}},
    {r1, r2, r3}][[2]]
>> {{7/8, 1/8, 1/2}, {1/8, 7/8, 1/2}}

```

This is optimal because it achieves the same welfare at π_i as unconditional implementation of α ,

22 $(\pi * q[[1, ; ;]]) . a1 + (\pi * q[[2, ; ;]]) . a2 - c[q, \pi] == \pi . \alpha$

» *True*


However, α is no longer dominated by this signal q because

23 $(\rho * q[[1, ; ;]]) . a1 + (\rho * q[[2, ; ;]]) . a2 - c[q, \rho] < \rho . \alpha$

» *True*

And indeed, we can verify that prior-concavity fails here,

24 `Plot[c[q, t pi + (1-t)rho], {t, 0, 1}]`



»

This breaks the step in the ‘dominance of necessity’ argument, where we go from an arbitrary prior where α is welfare-enhancing to a local perturbation with the same property.

B.2.3 Exogenous Information and Endogenous Learning

Figure 4 illustrates the interaction between exogenous and endogenous information for an RI agent of prior π with access to three actions and a learning technology that results in ignorance equivalent α and learning-proof menu $\bar{\mathcal{A}}$. Left to her own devices, the agent employs learning strategy $\mathcal{S}^{(q,r)}$ that yields posteriors π^q and π^r . If she instead obtains this signal from the principal, her welfare increases from $\pi \cdot \alpha$ to $\pi \cdot x^{(q,r)}$, and she would be willing to pay at most her endogenous information acquisition cost $c(\mathcal{S}^{(q,r)}, \pi)$. By sequential optimality (C4), the same is true for any Blackwell less informative signal: The potential welfare gain is exactly equal to the agent’s own learning cost as the signal purchase *replaces* endogenous learning. Suppose the





Figure 4: Willingness to pay for exogenous information.


principal offers a signal that yields more conclusive evidence in favor of state 1, yielding posteriors π^q and π^s . The agent still relies on the same two actions but implements the high-payoff action more often, so that welfare increases further to $\pi \cdot \mathbf{x}^{(q,s)}$. This signal is therefore more valuable to her than the first one, albeit only if she can acquire it at a cost strictly below her endogenous information acquisition cost. If the possible evidence in favor of state 1 grows even further, to posteriors π^q and π^t , the agent may eventually implement different actions, which further raises the welfare gains. Still, her willingness to pay remains strictly below her endogenous information acquisition cost. It is also possible to design signals that have no value to her. Here, this is the case if we add a third state that occurs with likelihood $\pi_3 = \frac{1}{2}$ in all previously discussed beliefs, and where all actions achieve a zero payoff. If, in addition, costs are linear in regards to that state, $c(\mathcal{S}, t\pi + (1-t)\pi') = tc(\mathcal{S}, \pi) + (1-t)c(\mathcal{S}, \pi')$ whenever $\pi_1/\pi_2 = \pi'_1/\pi'_2$, then the agent derives no benefit from learning about the relative likelihood of state 3: It does not help her identify the most profitable action, nor does it make her learning any easier.

B.2.4 Constrained Optimum

In [Figure 5\(a\)](#), we illustrate how to determine the optimal offer and acceptance strategies for one particular $t \notin T$. This state-independent transfer is above the threshold



(a) Construction



(b) Expected payoff for Agent 1

Figure 5: Partial trade when states are unverifiable

for unconditional acceptance \underline{t} but below the threshold for unconditional rejection \bar{t} .²⁴ As discussed above, Agent K can either accept the trade and then implement one of the actions $\mathbf{a} \in \mathcal{A}^K$ for a payoff of $\mathbf{a} - \mathbf{t}$, or he can settle for his outside option $\mathbf{0}$. Shifting all payoffs by \mathbf{t} , we can thus identify his optimal acceptance strategy from the learning-proof menu over $\mathcal{A}^K \cup \{\mathbf{t}\}$ under cost c^K (indicated in orange). Agent K rejects the offer at a posterior $\boldsymbol{\rho}^R$ orthogonal to the learning-proof menu at \mathbf{t} , and accepts the offer at a posterior $\boldsymbol{\rho}^A$ orthogonal at the other anchor action.

Using backwards induction, Agent 1 realizes that offering terms \mathbf{t} will yield expected payoff $\mathbf{t} \cdot \boldsymbol{\rho}^A$ conditional on trade happening, and $W(\mathcal{A}^1, \boldsymbol{\rho}^R, c^1)$ otherwise. By drawing equi-profit lines through \mathbf{t} and tangent to his learning-proof menu, he determines the ‘full-trade equivalent’ transfers \mathbf{t}^W . As far as Agent 1 is concerned, offering \mathbf{t} is equivalent to certain trade at \mathbf{t}^W — and this is true for all beliefs between $\boldsymbol{\rho}^A$ and $\boldsymbol{\rho}^R$. However, \mathbf{t}^W is not attractive enough to warrant an unconditional offer, as can be seen by the fact that it lies below one of the lower bounds from Equation (7). The learning-proof menu over $\mathcal{A}^1 \cup \{\mathbf{t}^W\}$ (drawn in dark red) reveals the equilibrium offer strategy: Agent 1 draws a noisy signal and either offers the trade \mathbf{t} at posterior $\boldsymbol{\rho}^O$ or implements the action that favors state one at posterior $\boldsymbol{\rho}^\emptyset$. If the trade is offered, Agent K draws a more accurate signal and either rejects at posterior $\boldsymbol{\rho}^R$ or accepts at posterior $\boldsymbol{\rho}^A$, in which case he goes on to implement action \circ .

Figure 5(b) plots the resulting payoff to Agent 1 across a range of state-independent transfers. For low prices $t < \underline{t}$, trade happens whenever Agent 1 wants it to, and thus his payoff increases monotonically with the price t . For high prices $t > \bar{t}$, trade never happens and Agent 1 receives his autarky payoff $W(\mathcal{A}^1, \boldsymbol{\pi}, c^1)$. Agent 1’s optimal price includes partial trade and corresponds to the equilibrium constructed in Figure 5(a).

B.2.5 Comparative Advantage and Repeated Trade

Notation. For each belief $\boldsymbol{\pi} \in \Delta\mathcal{I}$, we let \mathcal{S}^π denote the first signal acquired by the planner with prior $\boldsymbol{\pi}$ in an optimal learning sequence that minimizes the expected number of draws. We denote the associated distribution over posterior beliefs as $\nu^\pi \in \Delta\Delta\mathcal{I}$. By optimality,

$$W(\mathcal{A}^P, \boldsymbol{\pi}, c^P) = \int_{\boldsymbol{\rho} \in \Delta\mathcal{I}} W(\mathcal{A}^P, \boldsymbol{\rho}, c^P) d\nu^\pi - \min_{k \in \{1, \dots, K\}} c^k(\mathcal{S}^\pi, \boldsymbol{\pi}), \quad (22)$$

²⁴These thresholds can be computed analogously to Equation (8) and Equation (7), respectively, using Agent K ’s menu and cost.

and since the number of draws is minimized across all optimal learning sequences, a degenerate distribution indicates that the principal implements the best action unconditionally at this belief,

$$\nu^\pi = \delta_\pi \quad \implies \quad W(\mathcal{A}^P, \pi, c^P) = \max_{\mathbf{a} \in \mathcal{A}^P} \pi \cdot \mathbf{a}. \quad (23)$$

We also let \mathcal{K}^π denote the set of agents who can emulate the next step of the planner's strategy. Formally, if $\nu^\pi = \delta_\pi$, then agents in $\mathcal{K}^\pi = \arg \max_k \max_{\mathbf{a} \in \mathcal{A}^k} \pi \cdot \mathbf{a}$ are those with access to the optimal action; otherwise agents in $\mathcal{K}^\pi = \arg \min_k c^k(\mathcal{S}^\pi, \pi)$ are those who can implement \mathcal{S}^π at minimal cost. The minimality of draws also implies that $\mathcal{K}^\pi \cap \mathcal{K}^\rho = \emptyset$ for any $\rho \in \text{support}(\nu^\pi)$, for otherwise the principal could collapse signals \mathcal{S}^π and \mathcal{S}^ρ into a single draw.

Preliminaries. In order to guarantee that the learning-proof menu $\bar{\mathcal{A}}^P$ is well-defined, we need to verify that c^P satisfies properties (C1) to (C4). Continuity, Blackwell monotonicity and prior-concavity are maintained under minimization. Theorem 1 and Lemma 9 in [Bloedel and Zhong \[2020\]](#) ensure that these properties are also preserved under sequential optimization and that sequential learning cannot further improve upon c^P . And since in particular $c^P \leq c^1$, the principal also breaks all payoff-relevant ties through learning. Therefore, $\bar{\mathcal{A}}^P$ is well-defined and [Theorem 2](#) holds.

Game structure. We more formally describe the trading protocol as an infinite-period dynamic game. Each period consists of a simultaneous move game where all but one player can acquire any signal of their choice. (Drawing an uninformative signal models inaction.) One player can either learn or act: If there is no outstanding offer, the owner of the opportunity κ can extend a TIOLI offer to some other agent ℓ at terms $\mathbf{t} \in \mathbb{R}^I$, or implement an action $\mathbf{a} \in \mathcal{A}^\kappa$. If there is an outstanding offer at the beginning of the period, agent ℓ can accept or reject the offer. A rejection deactivates the offer, and the game proceeds to a period without outstanding offers. The game starts without any outstanding offer and with agent 1 as the owner. The game ends when an action is implemented, at which point the state i becomes public knowledge. Learning costs are incurred immediately and privately by the agent who draws the signal, but all agents are able to observe the resulting information, leading

to a common belief update at the end of each period. Trades represent transfer commitments that are due at the end of the game: Each agent pays t_i for any offers \mathbf{t} that he has accepted and receives t'_i from anyone who has accepted one of his offers \mathbf{t}' . In addition, the last owner reaps the payoffs from the implemented action, a_i .

Candidate equilibrium. We define a strategy profile by indexing each history by the agents' common belief $\boldsymbol{\pi}$ and the current owner κ , as well as potentially an outstanding offer from κ to some other agent ℓ with terms \mathbf{t} . Our candidate equilibrium is as follows:

- (i) At a history $(\boldsymbol{\pi}, \kappa)$ without an outstanding offer, but with $\kappa \in \mathcal{K}^\pi$ and $\nu^\pi = \delta_\pi$, owner κ implements the optimal action $\arg \max_{\mathbf{a} \in \mathcal{A}^\kappa} \boldsymbol{\pi} \cdot \mathbf{a}$. No other agent does any learning.
- (ii) At a history $(\boldsymbol{\pi}, \kappa)$ without an outstanding offer, but with $\kappa \in \mathcal{K}^\pi$ and $\nu^\pi \neq \delta_\pi$, owner κ draws signal \mathcal{S}^π . No other agent does any learning.
- (iii) At a history with an outstanding offer from owner κ to agent ℓ at terms \mathbf{t} under belief $\boldsymbol{\pi}$, agent ℓ analyzes RI problem $(\bar{\mathcal{A}}^P \cup \{\mathbf{t}\}, \boldsymbol{\pi}, c^\ell)$. He accepts the offer whenever there exists an optimal strategy that does not rely on \mathbf{t} . In particular, this includes all instances where the terms are at or below $\bar{\mathcal{A}}^P$. Agent ℓ rejects the offer if unconditional implementation of \mathbf{t} is optimal. (If indifferent, he defaults to acceptance.) If neither is optimal, he draws an optimal signal that maximizes $\boldsymbol{\pi} \cdot \mathbf{q}(\mathbf{t})$. No other agent does any learning.
- (iv) At a history $(\boldsymbol{\pi}, \kappa)$ without an outstanding offer and with $\kappa \notin \mathcal{K}^\pi$, owner κ extends a take-it-or-leave-it offer with terms $\arg \max_{\mathbf{t} \in \bar{\mathcal{A}}^P} \boldsymbol{\pi} \cdot \mathbf{t}$ to one of the agents in \mathcal{K}^π . No other agent does any learning.

Subgame payoffs. We claim that the expected continuation payoff after any history $(\boldsymbol{\pi}, \kappa)$ without an outstanding offer is equal to $W(\mathcal{A}^P, \boldsymbol{\pi}, c^P)$ for owner κ and zero for everyone else. (If there is an outstanding offer, the owner's payoff is bounded above by $W(\mathcal{A}^P, \boldsymbol{\pi}, c^P)$ but is possibly smaller.) As a consequence, everyone values opportunity ownership as granting access to the learning-proof menu $\bar{\mathcal{A}}^P$. We establish this claim by backwards induction:

- (i) At a terminal history $(\boldsymbol{\pi}, \kappa)$ without an outstanding offer, but with $\kappa \in \mathcal{K}^\pi$ and $\nu^\pi = \delta_\pi$, Equation (23) guarantees that owner κ earns an expected payoff of $W(\mathcal{A}^P, \boldsymbol{\pi}, c^P)$ by implementing his best action. Everyone else is inactive and earns a payoff of zero.
- (ii) At a history $(\boldsymbol{\pi}, \kappa)$ without an outstanding offer, but with $\kappa \in \mathcal{K}^\pi$ and $\nu^\pi \neq \delta_\pi$, owner κ incurs learning costs of $c^\kappa(\mathcal{S}^\pi, \boldsymbol{\pi})$. The game then proceeds with histories $(\boldsymbol{\rho}, \kappa)$, with $\boldsymbol{\rho}$ distributed according to ν^π . By backwards induction and Equation (22), the continuation payoff for owner κ is thus equal to

$$\int_{\boldsymbol{\rho} \in \Delta \mathcal{I}} W(\mathcal{A}^P, \boldsymbol{\rho}, c^P) d\nu^\pi - c^\kappa(\mathcal{S}^\pi, \boldsymbol{\pi}) = W(\mathcal{A}^P, \boldsymbol{\pi}, c^P),$$

and zero for everyone else. No other agent does any learning.

- (iii) At a history with an outstanding offer from owner κ to agent ℓ at terms \boldsymbol{t} under belief $\boldsymbol{\pi}$, trade happens whenever $W(\bar{\mathcal{A}}^P \cup \{\boldsymbol{t}\}, \boldsymbol{\pi}, c^\ell) \leq W(\bar{\mathcal{A}}^P, \boldsymbol{\pi}, c^\ell)$. By backwards induction, this results in a continuation payoff of $\boldsymbol{\pi} \cdot \boldsymbol{t}$ for κ , and $W(\mathcal{A}^P, \boldsymbol{\pi}, c^P) - \boldsymbol{\pi} \cdot \boldsymbol{t}$ for ℓ . Since this implies in particular that $\boldsymbol{\pi} \cdot \boldsymbol{t} \leq W(\bar{\mathcal{A}}^P, \boldsymbol{\pi}, c^\ell) \leq W(\bar{\mathcal{A}}^P, \boldsymbol{\pi}, c^P)$, the inductive hypothesis holds.

If the terms are so high that \boldsymbol{t} is the ignorance equivalent of RI problem $(\bar{\mathcal{A}}^P \cup \{\boldsymbol{t}\}, \boldsymbol{\pi}, c^\ell)$, no learning occurs and the game simply proceeds to history $(\boldsymbol{\pi}, \kappa)$ without an outstanding offer. Again, the inductive hypothesis holds.

At all other terms, agent ℓ draws a costly signal \mathcal{S}^ℓ that updates beliefs to the Bayes-plausible distribution ν^ℓ . Regardless of whether trade occurs or not, the sum of the continuation payoffs between agents ℓ and κ is bounded above by

$$\int_{\boldsymbol{\rho} \in \Delta \mathcal{I}} W(\mathcal{A}^P, \boldsymbol{\rho}, c^P) d\nu^\ell - c^\ell(\mathcal{S}^\ell, \boldsymbol{\pi}) \quad (24)$$

by backwards induction. Agent ℓ can guarantee himself a continuation payoff of zero by rejecting, so his optimal payoff is nonnegative. In turn, this bounds the continuation payoff of owner κ above by (24), which is at most $W(\mathcal{A}^P, \boldsymbol{\pi}, c^P)$, thus confirming the inductive hypothesis.

- (iv) At a history $(\boldsymbol{\pi}, \kappa)$ without an outstanding offer and with $\kappa \notin \mathcal{K}^\pi$, the suggested terms are on $\bar{\mathcal{A}}^P$ and represent an expected transfer of $\boldsymbol{\pi} \cdot \boldsymbol{t} = W(\mathcal{A}^P, \boldsymbol{\pi}, c^P)$ by

Theorem 2. By backwards induction, this results in payoffs $\boldsymbol{\pi} \cdot \mathbf{t} = W(\mathcal{A}^P, \boldsymbol{\pi}, c^P)$ for agent κ and $W(\mathcal{A}^P, \boldsymbol{\pi}, c^P) - \boldsymbol{\pi} \cdot \mathbf{t} = 0$ for agent ℓ . All other agents earn zero.

Optimality. From these subgame payoffs, it also follows that no agent has a profitable deviation. The owner earns $\int_{\boldsymbol{\rho} \in \Delta \mathcal{I}} W(\mathcal{A}^P, \boldsymbol{\rho}, c^P) d\nu - c^\kappa(\mathcal{S}, \boldsymbol{\pi})$ by deviating to a learning strategy \mathcal{S} with associated belief distribution ν , $\boldsymbol{\pi} \cdot \mathbf{a}$ by deviating to a game-ending move of implementing action \mathbf{a} , and at most $W(\mathcal{A}^P, \boldsymbol{\pi}, c^P)$ by extending any other offer. All are bounded above by his subgame payoff of $W(\mathcal{A}^P, \boldsymbol{\pi}, c^P)$. The acceptance strategy in (iii) is already chosen to be optimal across agent ℓ 's feasible subgame payoffs. And any other agent who would engage in learning will face only costs but no benefits.

Moreover, agents emulate the first-best strategy on the equilibrium path, and thus reap social surplus Δ .