Rational Inattention via Ignorance Equivalence

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

November 9, 2021

Abstract

We introduce the concept of the *ignorance equivalent* to effectively summarize the payoff possibilities in a finite Rational Inattention problem. The ignorance equivalent is a unique fictitious action that is weakly preferable to all existing learning strategies and yet generates no new profitable learning opportunities when added to the menu of choices. We fully characterize the relationship between the ignorance equivalent and the optimal learning strategies. Agents with heterogeneous priors self-select their own ignorance equivalent, which gives rise to an expected-utility analogue of the Rational Inattention problem. The approach provides new insights for menu expansion, the formation of consideration sets, the value of information, and belief elicitation. In a strategic game of contract choice, the ignorance equivalent emerges naturally in equilibrium.

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, 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, processing or "learning" that information. 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. 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 preserve key properties of the original problem but reduce its complexity by abstracting from learning and risk, respectively. Both remove the importance of (possibly subjective) beliefs about the state: Regardless of their 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 posterior beliefs about the state. Quite naturally, this is sufficient to rule out any new profitable learning opportunities and, more surprisingly, it is also necessary due to 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 of those is quite immediate now: Agents with different priors self-select into their appropriate ignorance equivalent.

By drawing on the self-selection property, one can transform any finite RI problem into a standard expected utility maximization. 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 previous paragraph, 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 it is always optimal to choose the agent's own ignorance equivalent, the RI problem effectively turns into a standard expected utility maximization over the learning-proof menu.

The expected-utility representation of the RI problem is 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 is 'actionable' given the menu at the agent's disposal. Even an informative signal may not induce any changes in the ignorance equivalent and thus may have no effect on her optimal behavior.

The learning-proof menu also simplifies the comparative statics of menu expansion. First, a researcher may want to know whether an agent will adjust her learning strategy if a new action becomes available. To do so, it is sufficient to consider a simplified menu consisting of the original ignorance equivalent and the new action—resulting in a much smaller set of learning strategies to evaluate. Second, a newly available action may increase the appeal of a previously unchosen action due to 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 for all expanded menus.

In addition, the learning-proof menu has relevance for experimental design. For instance, it is helpful to elicit private beliefs without generating belief-distorting learning incentives. Offering actions in the learning-proof menu represents the least-cost way for increasing the reporting accuracy of prior beliefs in the presence of agent learning.

Strategic menu design also occurs in contracting games. We study a formal outsourcing game between two rationally inattentive agents with exogenous information shocks and show that the learning-proof menu naturally emerges in the Perfect Bayesian Equilibrium. We provide here a simplified example of the game to highlight the strategic role of the ignorance equivalent.

Example 1. 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.

Consider an asset manager (he) who is free to design a fund $\boldsymbol{\alpha}$ that delivers return α_i in state i. The asset manager seeks the investor's business and has no information costs. When designing the fund, the asset manager realizes that the investor may first learn some information about the state and then decide whether to invest in the offered fund – which could lead the asset manager to miss out on some of the investor's business and create adverse selection issues. The manager also realizes that the investor might learn from the design of the fund itself if he tailors the payouts to his information about the state – which would increase the investor's information rents.

The ignorance equivalent is the answer to all the manager's problems. It does not reveal any free information and ensures that the investor willingly forgoes any learning and participates unconditionally — thus enabling the manager to extract the maximal information rents. ◇

Related literature. Rational inattention was first introduced into economics by Sims [2003], deploying the ideas of information theory to a model of learning.[1] Rational inattention models rapidly found their way into a variety of fields, from finance to monetary economics.[2]

While the basic idea of the RI model is elegant and simple, the model typically does not admit an analytic solution. As a result, Mutual Information (based on Shannon entropy) has become the de-facto cost structure for its tractability, with most applied research relying on special cases (e.g., a Linear-Quadratic Gaussian setup) or approximations. In Armenter et al. [2021], we propose a geometric approach to finite RI problems with Mutual Information costs that is well suited for numerical solution methods and computing the ignorance equivalent.

On a more conceptual level, theory developments have characterized optimal RI behavior in various ways. Caplin and Dean [2013] and Caplin et al. [2018a] solve the general finite model using a "posterior-based" approach; Matějka and McKay [2015] highlight the structural similarity with multinomial logit models; Caplin and Dean [2015] broaden the class of RI cost functions beyond Shannon entropy and conduct an empirical exploration of their validity; and Maćkowiak et al. [2018a] explore dynamic learning in an RI context. Alongside, a rich literature discusses the relative strengths of a variety of cost functions, including Bloedel and Zhong [2020], Denti et al. [2019], Hébert and Woodford [2020], Mensch [2018] and Pomatto et al. [2018], highlighting in particular that perceptual distance may make some states more difficult to disinguish than others. Our cost assumptions are compatible with some, but not all, of these parametrizations, and rely heavily on Bloedel and Zhong [2020]'s idea to study costs that are sequential learning proof. All of these papers focus directly on specific learning strategies rather than condensing the 'payoff possibilities' into a single vector in the spirit of our ignorance-equivalent approach.

RI models are also helpful to study the formation of endogenous *consideration sets*, which contain actions that are chosen with positive probability. Caplin et al.

[1] Of course, information theory traces back to the groundbreaking work of Claude Shannon, and information economics to George J. Stigler [Stigler, 1961].

[2] We cannot hope to properly review what is by now a large literature. See Maćkowiak et al. [2018b] for a survey of both theoretical and applied work with rational inattention models.

[2018b] characterize the optimal consideration sets in the case of Mutual Information costs. There is a strong sense for why our learning-proof menu can be seen as a 'latent' consideration set: First, all actions of the optimal consideration set are part of the learning-proof menu. Second, all actions that are in the learning-proof menu become part of the optimal consideration set if either the agent's prior changes, or the right new action is added to the menu. As such, our results provide novel insight regarding the comparative statics of these consideration sets for more general costs.

Finally, our contract-design game is related to Bayesian Persuasion with costly information acquisition [Gentzkow and Kamenica, 2014] where a sender strategically provides information to influence the receiver's choice. Instead of a sender-receiver structure, we focus on outsourcing where a manager offers state-contingent payments to an agent in return for getting access to the choice problem itself. Whereas the potential of receiver learning pushes the sender to reveal information up front [Matyskova, 2018], our manager avoids early information revelation by offering the same (learning-proof) terms unconditionally.[3] However, both their sender and our manager seek to undercut the incentives for receiver/agent learning as this would reduce their payoff.

Paper structure. We describe the standard Rational Inattention problem in Section 2 and discuss what cost functions are compatible with our framework. We then introduce the ignorance equivalent in Section 3, along with its key properties. In Section 4, we define the learning-proof menu and discuss its relevance for determining welfare-enhancing menu expansions, the value of information and robust belief-elicitation. In Section 5, we describe a strategic two-player game between two RI agents and show that the ignorance equivalent and learning-proof menu naturally emerge in the Perfect Bayesian Equilibrium. Section 6 concludes. With the exception of some immediate corollaries, all proofs are in the Appendix.

[3] In that sense, our unconditional contract offers remind of Myerson [1983]'s 'inscrutable' mechanism choices.

2 Rational Inattention Model

The rationally inattentive decision maker has to implement an action from the finite menu $\mathcal{A}$.[4] Payoffs of each action depend on an unknown state of the world $i \in \mathcal{I} := \{1, ..., I\}$. The *prior* probabilities for each state are strictly positive, $\boldsymbol{\pi} \in \mathsf{int}(\Delta\mathcal{I})$.[5] No two actions are payoff equivalent, and we identify an action $\boldsymbol{a} \in \mathcal{A}$ by its state-dependent payoffs $(a_1, ..., a_I) \in \mathbb{R}^I$.

The agent can condition her choice on the outcome of a costly signal $\mathcal{S} = \langle S, \boldsymbol{q}\rangle$, where S refers to a finite signal realization space and $\boldsymbol{q} \in (\Delta S)^{\mathcal{I}}$ denotes the conditional probabilities $q_i(s)$ of realization s in state i. We denote the cost of signal $\mathcal{S}$ under belief $\boldsymbol{\rho} \in \Delta\mathcal{I}$ by $c(\mathcal{S}, \boldsymbol{\rho}) \geq 0$. This cost is finite for all signals except possibly for those that rule out some state i with certainty, i.e. when $q_i(s) = 0$ for some $s \in \mathsf{support}(\boldsymbol{q})$. 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 $\arg\max_{\boldsymbol{a}\in\mathcal{A}} \boldsymbol{a} \cdot \boldsymbol{\pi}^s$.

For each choice problem $(\mathcal{A}, \boldsymbol{\pi}, c)$, we denote the maximal welfare by

$$W(\mathcal{A}, \boldsymbol{\pi}, c) = \max_{\mathcal{S}=\langle S,\boldsymbol{q}\rangle} \sum_{s\in S} \left(\sum_{i=1}^{I} q_i(s)\pi_i \right) \left(\max_{\boldsymbol{a}\in\mathcal{A}} \sum_{i=1}^{I} a_i \pi_i^s \right) - c(\mathcal{S}, \boldsymbol{\pi}). \tag{RI}$$

Because of the obedience principle, it is without loss of generality to restrict attention to signals where the signal returns a 'recommended action', $S = \mathcal{A}$, and the agent implements the signal recommendation.[6] For simplicity, we use the term 'learning strategy' to refer to situations where the agent follows such a signal. We say that "unconditional implementation of $\boldsymbol{a}$ is optimal" if the degenerate strategy with $\boldsymbol{q}(\boldsymbol{a}) = \mathbf{1}$ is optimal, and that "it is optimal to implement $\boldsymbol{a}$ with positive probability" if at least one strategy with $\boldsymbol{q}(\boldsymbol{a}) > \mathbf{0}$ is optimal. Since adding additional actions only increases the space of learning strategies, welfare is menu-monotone, $W(\mathcal{A}', \boldsymbol{\pi}, c) \geq W(\mathcal{A}, \boldsymbol{\pi}, c)$ whenever $\mathcal{A}' \supseteq \mathcal{A}$.

[4] Our results carry over to compact menus, as long as it is without loss of generality to assume that the optimal signal is finite. In particular, we will use the same notation when we discuss the learning-proof menu in Section 4.

[5] The open simplex can be written as $\mathsf{int}(\Delta\mathcal{I}) := \left\{\boldsymbol{\pi} \in \mathbb{R}^I \mid \boldsymbol{\pi} \gg \mathbf{0} \text{ and } \boldsymbol{\pi} \cdot \mathbf{1} = 1\right\}$. Throughout, we use the convention that $\boldsymbol{v} \geq \boldsymbol{w}$ if and only if $v_i \geq w_i\ \forall i$, that $\boldsymbol{v} > \boldsymbol{w}$ if and only if $\boldsymbol{v} \geq \boldsymbol{w}$ and $\boldsymbol{v} \neq \boldsymbol{w}$, and that $\boldsymbol{v} \gg \boldsymbol{w}$ if and only if $v_i > w_i\ \forall i$.

[6] Lemma A.1 in Appendix A.1 provides a formal statement of this standard result in the information design literature.

Admissible costs. A common family of cost functions that is compatible with our model is known as *uniformly posterior-separable costs* [Caplin et al., 2022], where

$$c(\mathcal{S},\boldsymbol{\pi}) = \sum_{s\in S}(\boldsymbol{\pi}\cdot\boldsymbol{q}(s))\phi(\boldsymbol{\pi}^s) - \phi(\boldsymbol{\pi}), \tag{UPS}$$

for any potential function $\phi : \Delta\mathcal{I} \to \mathbb{R}\cup\{\infty\}$ that assigns finite values to a convex open set $\mathcal{J} \supseteq \mathsf{int}(\Delta\mathcal{I})$, is twice differentiable and convex over $\mathcal{J}$, and ensures that c is weakly concave in the prior.[7] 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].

In order to highlight the specific properties that drive our results, we characterize the (possibly larger) set of admissible cost functions with five specific conditions. For each property, we start with the formal definitions, then provide the intuition behind each assumption and discuss some key implications that are proven formally in Appendix A.1. We finish the section by formally establishing that (UPS) costs satisfy all conditions. The reader who already has one of the above-mentioned cost functions in mind may feel free to jump ahead to Section 3.

(C1) The cost function is continuous wherever finite: For any prior $\boldsymbol{\pi}$ and signal $\mathcal{S} = \langle S, \boldsymbol{q}\rangle$ with $c(\mathcal{S},\boldsymbol{\pi}) < \infty$ and for any $\varepsilon > 0$, there exists $\delta > 0$ such that $|c(\langle S, \tilde{\boldsymbol{q}}\rangle, \tilde{\boldsymbol{\pi}}) - c(\mathcal{S},\boldsymbol{\pi})| < \varepsilon$ whenever $\|\tilde{\boldsymbol{q}} - \boldsymbol{q}\| + \|\boldsymbol{\pi} - \tilde{\boldsymbol{\pi}}\| < \delta$ under the Euclidean norm $\|\cdot\|$.

Continuity rules out any sudden cost jumps due to marginal changes in either the prior or the signal. It ensures that the choice problem (RI) locally satisfies all the conditions of Berge's Theorem of the maximum, so that the optimum exists and indirect utility is locally continuous in the prior belief (Lemma A.2).

(C2) The agent can freely dispose of information: Cost function $c(\cdot,\boldsymbol{\pi})$ is non-decreasing in the Blackwell order[8] and $c(\mathcal{S},\cdot)$ is weakly concave in the prior for all $\mathcal{S}$.

[7] Bloedel and Zhong [2020] provide a check for prior-concavity based on the Hessian of any twice-differentiable potential ϕ.

[8] Formally, signal $\mathcal{S} = \langle S, \boldsymbol{q}\rangle$ is Blackwell more informative than signal $\tilde{\mathcal{S}} = \langle \tilde{S}, \tilde{\boldsymbol{q}}\rangle$ if for each $s \in S$, there exists a lottery $m^s \in \Delta\tilde{S}$ such that $\tilde{\boldsymbol{q}} = \sum_{s\in S} m^s\boldsymbol{q}(s)$. Equivalently, the distribution over posteriors generated by $\mathcal{S}$ forms a mean-preserving spread of that generated by $\tilde{\mathcal{S}}$.

Free disposal of information (C2) is best explained by introducing an 'assistant' to the decision maker. For Blackwell monotonicity, suppose the assistant draws a signal $\langle S, \boldsymbol{q}\rangle$ and then, depending on its realization $s \in S$, communicates a garbled message $m^s \in \Delta S'$ 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. 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. Prior-concavity implies in particular that welfare W is convex in the prior belief (Lemma A.3).

(C3) Sequential information acquisition brings no cost savings: For any contingency plan where the agent first draws $\mathcal{S} = \langle S, \boldsymbol{q}\rangle$ and upon observing $s \in S$ draws signal $\mathcal{S}^s = \langle S^s, q^s\rangle$, a one-shot implementation of the same process, $\tilde{\mathcal{S}} = \langle S \times \bigcup_{s\in S} S^s, \tilde{\boldsymbol{q}}\rangle$ with $\tilde{q}_i(s, \tilde{s}) = q_i(s) q_i^s(\tilde{s})$, costs no more than the expected cost of the contingency plan, $c(\tilde{\mathcal{S}}, \boldsymbol{\pi}) \leq c(\mathcal{S}, \boldsymbol{\pi}) + \sum_{s\in S} (\boldsymbol{\pi} \cdot \boldsymbol{q}(s))\, c(\mathcal{S}^s, \boldsymbol{\pi}^s)$.

As argued by Bloedel and Zhong [2020], an optimizing agent without delay costs would always exploit any cost savings associated with sequential information acquisition. Condition (C3) simply assumes that $c(\mathcal{S}, \boldsymbol{\pi})$ is already the cost-minimizing envelope over all sequential information strategies, avoiding more cumbersome notation. Condition (C3) also implies that if the agent implements $\boldsymbol{a} \in \mathcal{A}$ at some posterior $\boldsymbol{\rho}$, then she would do so unconditionally if her prior is set to $\boldsymbol{\rho}$ (Lemma A.4).

The next condition requires that the marginal cost of the first bit of information is zero, ruling out fixed costs of learning. We define the ε-precision dilution of a signal $\mathcal{S} = \langle S, \boldsymbol{q}\rangle$ at prior $\boldsymbol{\pi} \in \Delta\mathcal{I}$ as $\mathcal{S}^{(\varepsilon,\boldsymbol{\pi})} = \langle S, \boldsymbol{q}^{(\varepsilon,\boldsymbol{\pi})}\rangle$ with $q_i^{(\varepsilon,\boldsymbol{\pi})}(s) = \varepsilon q_i(s) + (1 - \varepsilon)(\boldsymbol{q}(s) \cdot \boldsymbol{\pi})$. Intuitively, the ε-precision dilution is obtained by having an assistant draw signal $\mathcal{S}$ with probability ε and otherwise default to a pure noise signal with probabilities equal to $\mathcal{S}$'s marginals $\boldsymbol{\pi} \cdot \boldsymbol{q}$. The assistant only communicates the signal realization $s \in S$, and does not reveal whether it stems from the pure noise signal. Thus, if $\boldsymbol{\pi}^s$ denotes the agent's posterior upon observing $s \in S$ from signal $\mathcal{S}$, the same observation from $\mathcal{S}^{(\varepsilon,\boldsymbol{\pi})}$ leads to posterior $\varepsilon\boldsymbol{\pi}^s + (1 - \varepsilon)\boldsymbol{\pi}$.

(C4) Learning costs are negligible at zero information: For any signal $\mathcal{S} = \langle S, \boldsymbol{q}\rangle$, belief $\boldsymbol{\pi} \in \Delta\mathcal{I}$, and constant $\bar{c} > 0$, there exists $\varepsilon \in (0, 1]$ such that the cost of the ε-precision dilution of signal $\mathcal{S}$ is bounded above by $\varepsilon\bar{c}$, $c(\mathcal{S}^{(\varepsilon,\boldsymbol{\pi})}, \boldsymbol{\pi}) < \varepsilon\bar{c}$.

Condition (C4) has several important implications: Pure noise is always free (Lemma A.5) and the agent can randomize between strategies at no additional cost. And if two actions achieve the same expected consumption utility, condition (C4) guarantees that there exists a noisy enough signal whose expected benefits outweigh the cost by any arbitrary factor (Lemma A.6). In combination with condition (C3), this implies in particular that under an optimal signal, all posterior-optimal actions are welfare equivalent with probability one (Lemma A.7), or she could improve upon her payoff with a sequential strategy that draws another signal.

Henceforth, whenever we talk of an RI problem $(\mathcal{A}, \boldsymbol{\pi}, c)$, we implicitly assume that $\mathcal{A} \subset \mathbb{R}^I$, $\boldsymbol{\pi} \in \mathsf{int}(\Delta\mathcal{I})$ and c satisfies conditions (C1) to (C4).

When sequential information acquisition is neither more nor less costly than one-shot information acquisition, the ignorance equivalent has additional sufficiency properties when it comes to menu expansion. We impose this additional condition only when we talk about menu expansion in Section 4, but list it here for easy reference.

(C5) Sequential information acquisition incurs no extra costs: For any contingency plan where the agent first draws $\mathcal{S} = \langle S, \boldsymbol{q} \rangle$ and upon observing $s \in S$ draws signal $\mathcal{S}^s = \langle S^s, q^s \rangle$, a one-shot implementation of the same process, $\tilde{\mathcal{S}} = \langle S \times \bigcup_{s \in S} S^s, \tilde{\boldsymbol{q}} \rangle$ with $\tilde{q}_i(s, \tilde{s}) = q_i(s) q_i^s(\tilde{s})$, costs no less than the expected cost of the contingency plan, $c(\tilde{\mathcal{S}}, \boldsymbol{\pi}) \geq c(\mathcal{S}, \boldsymbol{\pi}) + \sum_{s \in S} (\boldsymbol{\pi} \cdot \boldsymbol{q}(s))\, c(\mathcal{S}^s, \boldsymbol{\pi}^s)$.

Together with (C3), this condition implies that the agent is exactly indifferent across all Blackwell-equivalent contingency plans.

Some of our conditions on the cost function are more prevalent in the literature than others. Continuity (C1) and Blackwell monotonicity (C2) are satisfied by virtually all commonly used cost functions. Prior-concavity (C2) is less often mentioned, but satisfied in particular by all cost functions that are prior-free [Denti et al., 2019, Gentzkow and Kamenica, 2014, Mensch, 2018, Pomatto et al., 2018]. Sequential learning-proofness (C3) has been introduced and characterized by Bloedel and Zhong [2020] and is linked to the dynamic information-sampling formulation of Hébert and Woodford [2019]. In addition, Frankel and Kamenica [2019] show that Conditions (C3) and (C5) jointly restrict costs exactly to the uniformly posterior-separable family. While we are not aware of other papers that explicitly impose Condition (C4), all of our conditions are jointly satisfied by the class of prior-concave and smooth uniformly posterior-separable functions according to (UPS).

Lemma 1. *Any uniformly posterior-separable cost within the family defined by Formulation* (UPS) *satisfies Conditions (C1) to (C5).*

The proof draws heavily on Bloedel and Zhong [2020]'s characterization of sequential learning-proof cost functions, and refers to them verbatim for properties (C2), (C3) and (C5). The remaining properties follow from the differentiability of the cost potential function.

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 RI problem $(\mathcal{A}, \boldsymbol{\pi}, c)$ if and only if

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

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 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

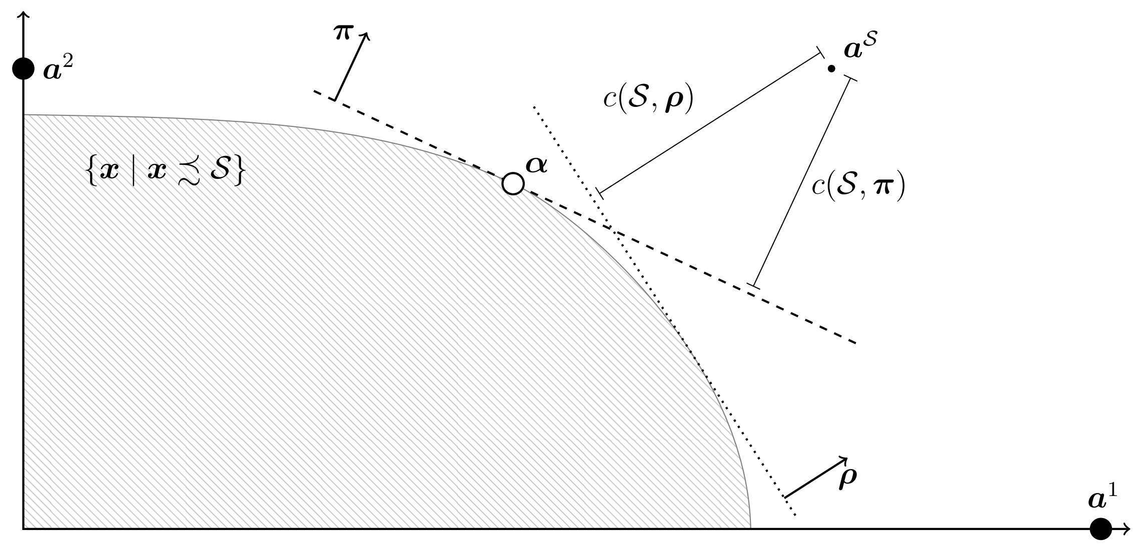

Figure 1: Dominated payoff vectors and construction of the ignorance equivalent.

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 $\boldsymbol{\rho} \in \Delta\mathcal{I}$ obtains expected utility $\boldsymbol{\rho} \cdot \boldsymbol{x}$ from implementing action $\boldsymbol{x}$ unconditionally, and incurs no learning costs. If instead the agent follows the learning strategy $\mathcal{S} = \langle \mathcal{A}, \boldsymbol{q} \rangle$, she achieves expected consumption utility $\boldsymbol{a}^{\mathcal{S}}_i = \sum_{\boldsymbol{a} \in \mathcal{A}} q_i(\boldsymbol{a}) a_i$ 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 $\boldsymbol{x}$ is dominated by learning strategy $\mathcal{S}$ if the latter yields weakly larger welfare for any belief $\boldsymbol{\rho}$. 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 $\boldsymbol{x}$ to the menu.

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

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

Figure 1 sketches a sample RI problem $(\mathcal{A}, \boldsymbol{\pi}, c)$ with two states (on either axis) and two actions ($\boldsymbol{a}^1$ and $\boldsymbol{a}^2$) to illustrate the concept of dominance. The learning strategy $\mathcal{S}$ implements action $\boldsymbol{a}^2$ with probability 1/4 in state 1 and with probability 1 in state 2, leading to expected consumption utility $a^{\mathcal{S}}_1 = \frac{1}{4}a^2_1 + \frac{3}{4}a^3_1$ and $a^{\mathcal{S}}_2 = a^2_2$. 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 $\boldsymbol{a}^{\mathcal{S}}$.

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.

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,* $\arg\max_{\boldsymbol{x} \precsim \mathcal{S}} \boldsymbol{\pi} \cdot \boldsymbol{x}$.

Four arguments are key to the result:

First, an optimal strategy $\mathcal{S}$ exists by continuity of the cost (C1). Second, unconditional implementation of the most attractive $\mathcal{S}$-dominated payoff vector $\boldsymbol{\alpha}$ achieves the same welfare as $\mathcal{S}$, establishing that $W(\{\boldsymbol{\alpha}\}, \boldsymbol{\pi}, c) = \boldsymbol{\pi} \cdot \boldsymbol{a}^{\mathcal{S}} - c(\mathcal{S}, \boldsymbol{\pi}) = W(\mathcal{A}, \boldsymbol{\pi}, c)$. These steps are mostly technical. Continuity of the cost (C1) ensures existence of an optimal signal by a standard application of Berge's Theorem, spelled out in Lemma A.2 in the appendix. Prior-concavity of the cost (C2) ensures that for any prior and signal, the inequality in Definition 2 is binding for at least one dominated payoff vector. We prove this by relying on the finite intersection property from topology. Applied to $\boldsymbol{\pi}$ and $\mathcal{S}$, this implies that the dashed line in Figure 1, which describes the net utility from implementing $\mathcal{S}$ under prior $\boldsymbol{\pi}$, touches the set of dominated payoff vectors $\{\boldsymbol{x} \mid \boldsymbol{x} \precsim \mathcal{S}\}$. Equivalently, unconditional implementation of $\boldsymbol{\alpha}$ achieves welfare $W(\mathcal{A}, \boldsymbol{\pi}, c)$.

Third, dominance $\boldsymbol{\alpha} \succsim \mathcal{S}$ 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 continuing with strategy $\mathcal{S}$ instead and relying only on actions in menu $\mathcal{A}$. By sequential optimality (C3), a one-shot implementation of this two-step strategy is weakly cheaper and the agent can thus achieve at least as much welfare by restricting attention to menu $\mathcal{A}$, $W(\mathcal{A} \cup \{\boldsymbol{\alpha}\}, \boldsymbol{\pi}, c) = W(\mathcal{A}, \boldsymbol{\pi}, c)$. 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 on cost conditions (C2) to (C4): Proceeding by contradiction, we assume that there are two distinct candidate ignorance equivalents $\boldsymbol{\alpha}^1 \neq \boldsymbol{\alpha}^2$ and construct a two-step strategy where a binary signal helps the agent choose between strategy $\mathcal{S}$ or unconditional implementation of an ignorance equivalent. We show that the welfare gains are at least linear in precision and thus initially outweigh the extra costs by (C4), implying that the agent can improve upon the previously optimal strategy. Indeed, the two candidate payoff vectors must achieve the same expected utility $W(\mathcal{A}, \boldsymbol{\pi}, c)$ at prior $\boldsymbol{\pi}$ by Definition 1, yet differ in expected utility for some other belief $\boldsymbol{\rho} \neq \boldsymbol{\pi}$. Suppose that the agent initially draws a binary signal $\mathcal{S}^0$ that updates her belief either towards or away from $\boldsymbol{\rho}$ with small precision ε. If the agent implements the more attractive of $\boldsymbol{\alpha}^1$ and $\boldsymbol{\alpha}^2$ after each realization of signal $\mathcal{S}^0$, consumption utility increases linearly with precision and the agent incurs no costs in the second step. If the agent follows strategy $\mathcal{S}$ after each realization, consumption utility remains $\boldsymbol{\pi} \cdot \boldsymbol{a}^{\mathcal{S}}$ but the expected cost of $\mathcal{S}$ decreases with the access to extra information by (C2). Ultimately, we are interested in sequential learning strategies that rely on menu $\mathcal{A} \cup \{\boldsymbol{\alpha}^k\}$ for one of the two candidates, with the agent implementing the more favorable of $\boldsymbol{\alpha}^k$ or $\mathcal{S}$ after each realization. By the above, at least one of these strategies (say $k = 1$) brings second-step welfare benefits that scale linearly with precision ε, at a first-step signal cost that is negligible next to ε by (C4). The one-shot implementation of the same strategy is no more costly by (C3), and thus indicates a profitable learning opportunity that arises from adding $\boldsymbol{\alpha}^1$ to the menu, reaching a contradiction with Definition 1.

Necessity of dominance. Theorem 1 implies that dominance is not just sufficient to rule out learning opportunities – it is also necessary. This may surprise at first; after all, it requires that an optimal signal needs to be preferable to the unconditional implementation of $\boldsymbol{\alpha}$ for *all* beliefs.[9] However, if there are any welfare gains $\Delta > 0$ from implementing $\boldsymbol{\alpha}$ rather than $\mathcal{S}$ at some posterior $\boldsymbol{\rho}$, the local perturbation $\boldsymbol{\rho}^\varepsilon := (1-\varepsilon)\boldsymbol{\pi} + \varepsilon\boldsymbol{\rho}$ also yields first-order welfare gains by prior-concavity of the cost (C2),

$$\boldsymbol{\rho}^\varepsilon \cdot \boldsymbol{a}^{\mathcal{S}} - c(\boldsymbol{a}^{\mathcal{S}}, \boldsymbol{\rho}^\varepsilon) \leq (1-\varepsilon) \underbrace{[\boldsymbol{\pi} \cdot \boldsymbol{a}^{\mathcal{S}} - c(\boldsymbol{a}^{\mathcal{S}}, \boldsymbol{\pi})]}_{=\boldsymbol{\pi}\cdot\boldsymbol{\alpha} \text{ by Definition 1}} + \varepsilon \underbrace{[\boldsymbol{\rho} \cdot \boldsymbol{a}^{\mathcal{S}} - c(\boldsymbol{a}^{\mathcal{S}}, \boldsymbol{\rho})]}_{=\boldsymbol{\rho}\cdot\boldsymbol{\alpha}+\Delta} = \boldsymbol{\rho}^\varepsilon \cdot \boldsymbol{\alpha} + \varepsilon\Delta.$$

[9] In situations with multiple optimal signals, Theorem 1 requires that *each* of them dominates $\boldsymbol{\alpha}$.

In the same way as in the uniqueness argument above, the agent could then initially draw a noisy signal to help her choose between $\boldsymbol{\alpha}$ or $\mathcal{S}$ by updating her belief either towards or away from $\boldsymbol{\pi}^{\varepsilon}$. The marginal cost of this additional signal vanishes with ε by (C4) and is thus eventually outweighed by the gains. A one-shot implementation would do even weakly better by (C3) and would thus improve upon the welfare of $\mathcal{S}$.

Connection with optimal strategies. While the ignorance equivalent is always unique for any RI problem $(\mathcal{A}, \boldsymbol{\pi}, c)$, there may exist multiple optimal learning strategies. Nevertheless, there is a tight connection thanks to Theorem 1: From any optimal learning strategy $\mathcal{S}$, the ignorance equivalent is obtained as the most attractive dominated payoff vector, and given an ignorance equivalent $\boldsymbol{\alpha}$, the set of optimal signals are exactly those that dominate $\boldsymbol{\alpha}$. Since the dominance relationship does not depend on the prior $\boldsymbol{\pi}$, any two RI problems $(\mathcal{A}, \boldsymbol{\pi}, c)$ and $(\mathcal{A}, \boldsymbol{\pi}', c)$ that share the same ignorance equivalent also share all optimal learning strategies.

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

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.2), the ignorance equivalent is continuous in the prior.

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

Self-selection. By definition, the ignorance equivalent generates no additional learning under the agent's prior. Theorem 1 further implies that the ignorance equivalent musts be dominated, and thus it generates no additional learning opportunities under any prior, whether it is the one it was designed for or not. 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 their respective ignorance equivalents to the menu would not be welfare-enhancing for either agent, yet each would be willing to implement their appropriate ignorance equivalent unconditionally.

Corollary 3. *Let* $\boldsymbol{\alpha}$ *denote the ignorance equivalent of RI problem* $(\mathcal{A}, \boldsymbol{\pi}, c)$*. For any belief* $\boldsymbol{\rho} \in \Delta\mathcal{I}$*,*

$$W(\mathcal{A}, \boldsymbol{\rho}, c) = W(\mathcal{A} \cup \{\boldsymbol{\alpha}\}, \boldsymbol{\rho}, c).$$

Moreover, the ignorance equivalent of $(\mathcal{A}, \boldsymbol{\rho}, c)$ *is equal to that of* $(\mathcal{A} \cup \{\boldsymbol{\alpha}\}, \boldsymbol{\rho}, c)$ *for any prior* $\boldsymbol{\rho} \in \mathsf{int}(\Delta\mathcal{I})$*.*

4 Learning-Proof Menu

By design, adding the ignorance equivalent does not generate any profitable learning opportunities but makes it optimal to forgo learning. By the self-selection property, the same is true if we include ignorance equivalents across *all* priors. The result is a reformulation of the agent's RI problem as a standard expected-utility maximization problem over a modified menu. In essence, the agent's capacity to learn is *as if* she had access to additional, fictitious actions – and including these actions renders her learning capacity obsolete. Thanks to the reduced complexity, this perspective allows for clearer intuition on several topics, including the value of exogenous information and eliciting beliefs.

Formally, we define the *learning-proof menu* as the set of all such fictitious actions.

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

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

The menu $\bar{\mathcal{A}}$ is 'learning-proof' because an RI agent, when faced with $\bar{\mathcal{A}}$, would forgo learning no matter her prior, and it shares key properties with the original menu because their ignorance equivalents always agree.

Theorem 2. *For any menu* $\mathcal{A}$ *and cost function* c*, the learning-proof menu* $\bar{\mathcal{A}}$ *is the smallest set such that*

(a) Ignorance is always an optimal strategy in $(\bar{\mathcal{A}}, \boldsymbol{\pi}, c)$ *for any* $\boldsymbol{\pi} \in \mathsf{int}(\Delta\mathcal{I})$*,*

$$W(\bar{\mathcal{A}}, \boldsymbol{\pi}, c) = \max_{\boldsymbol{a} \in \bar{\mathcal{A}}} \boldsymbol{\pi} \cdot \boldsymbol{a}.$$

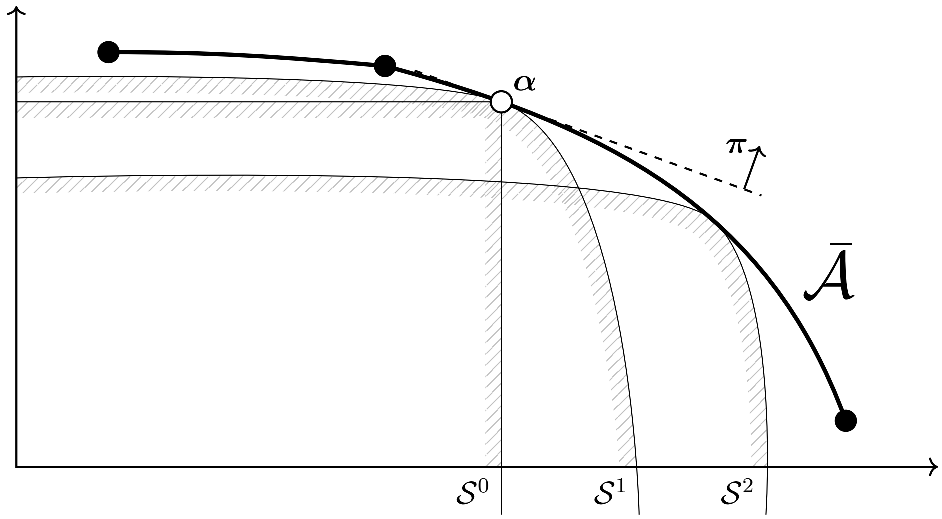

Figure 2: Construction of the Learning-Proof Menu $\bar{\mathcal{A}}$ (thick solid line) for an RI problem with two states and three actions (action payoffs are indicated as black dots).

(b) The menus $\mathcal{A}$ and $\bar{\mathcal{A}}$ share the same ignorance equivalent at any prior $\boldsymbol{\pi} \in \mathsf{int}(\Delta\mathcal{I})$, $\boldsymbol{\alpha}^{(\mathcal{A},\boldsymbol{\pi},c)} = \boldsymbol{\alpha}^{(\bar{\mathcal{A}},\boldsymbol{\pi},c)}$.

Because each ignorance equivalent is uniquely maximal in some direction $\boldsymbol{\pi} \in \mathsf{int}(\Delta\mathcal{I})$, the learning-proof menu represents the upper boundary[10] of the strictly convex set

$$\bigcap_{\boldsymbol{\pi}\in\mathsf{int}(\Delta\mathcal{I})} \left\{\boldsymbol{x} \in \mathbb{R}^I \mid \boldsymbol{\pi}\cdot\boldsymbol{x} \leq W(\mathcal{A},\boldsymbol{\pi},c)\right\}. \tag{1}$$

Any change in the problem parameters that raise agent welfare W across all priors will move this boundary outwards. Examples of such changes include the addition of new actions to the menu $\mathcal{A}$, or reductions in the learning cost c.

Similarly, 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,

$$\bigcup_{\boldsymbol{q}\in(\Delta\mathcal{A})^{\mathcal{I}}} \left\{\boldsymbol{x} \in \mathbb{R}^I \mid \boldsymbol{x} \precsim \langle\mathcal{A},\boldsymbol{q}\rangle\right\} \tag{2}$$

is thus also equal to the learning-proof menu $\bar{\mathcal{A}}$. Figure 2 illustrates this construction

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

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 upper boundary of all dominated payoff vectors $x \precsim \mathcal{S}^k$. Strategies $\mathcal{S}^1$ and $\mathcal{S}^2$ rely only on actions in the original menu and are feasible in RI problem $(\mathcal{A}, \boldsymbol{\pi}, c)$. They are therefore included in the union of Equation (2). Strategy $\mathcal{S}^1$ is optimal at $\boldsymbol{\pi}$ 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, the set of dominated payoff vectors lies strictly below $\bar{\mathcal{A}}$. Strategy $\mathcal{S}^0$ implements $\boldsymbol{\alpha}$ unconditionally at no cost. It is an example of a new learning strategy that becomes feasible in menu $\bar{\mathcal{A}}$. Strategy $\mathcal{S}^0$, too, is optimal at prior $\boldsymbol{\pi}$. Yet by the self-selection property (Corollary 3), none of the newly feasible learning strategies dominate additional payoffs vectors.

Expected-utility approach. Conceptually, the learning-proof menu allows us to identify the solution to the original (RI) problem in two steps by first locating the ignorance equivalent $\boldsymbol{\alpha}$ through the expected utility maximization problem $\arg\max_{\boldsymbol{a} \in \bar{\mathcal{A}}} \boldsymbol{\pi} \cdot \boldsymbol{a}$. Corollary 1 then identifies the full set of optimal learning strategies in a way that is independent of the prior belief.

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 'conceptual' 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 $\boldsymbol{a} \in \mathcal{A}$, the following are equivalent:*

(a) Action $\boldsymbol{a}$ is an anchor of the learning-proof menu, $\boldsymbol{a} \in \bar{\mathcal{A}}$.

(b) It is optimal to implement $\boldsymbol{a}$ unconditionally for some prior $\boldsymbol{\pi} \in \mathsf{int}(\Delta\mathcal{I})$.

(c) It is optimal to implement $\boldsymbol{a}$ with positive probability for some prior $\boldsymbol{\pi} \in \mathsf{int}(\Delta\mathcal{I})$.

(d) There exists no prior $\boldsymbol{\rho} \in \mathsf{int}(\Delta\mathcal{I})$ such that the ignorance equivalent of $(\mathcal{A}, \boldsymbol{\rho}, c)$ dominates $\boldsymbol{a}$ statewise, $\boldsymbol{\alpha}^{(\mathcal{A},\boldsymbol{\rho},c)} > \boldsymbol{a}$.

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., 2018b] 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.

Characterization (d) describes in what sense non-anchor actions $\boldsymbol{a} \in \mathcal{A} \setminus \bar{\mathcal{A}}$ 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 $\boldsymbol{\alpha}^{(\mathcal{A},\boldsymbol{\rho},c)}$ that dominates $\boldsymbol{a}$ statewise. By Theorem 1, this also implies that there exists a specific learning strategy (any one that is optimal under prior $\boldsymbol{\rho}$) which all agents, irrespective of their prior, strictly prefer to implementing $\boldsymbol{a}$.

4.1 Menu Expansion

When new actions are added to the menu, the RI agent re-calibrates her entire learning strategy. This can lead to patterns of behavior that are absent in fixed-information models: Matějka and McKay [2015] show by example that adding a new action to the menu may 'activate' a previously unchosen action which now is implemented with positive probability. This suggests that the comparative statics of the consideration set can depend in complex ways on the geometry of the full menu $\mathcal{A}$ and its interaction with the cost function.

Fortunately, the ignorance equivalent and the learning-proof menu bring structure to menu expansion, as long as the agent is indifferent across all sequential implementations of a given signal. To ensure this, we here impose condition (C5) in addition to (C1) to (C4).

If one is interested purely in whether the new action is implemented with positive probability, it is without loss of generality to replace the menu with its ignorance equivalent.[11] Since the ignorance equivalent can be derived from the optimal learning strategy (Theorem 1), the result also means that unchosen actions do not affect *whether or not* a new action is attractive. However, the diversity of action payoffs in the full menu $\mathcal{A}$ presents more opportunities for profitable learning, and so the

[11]In the case of Shannon entropy costs, this result is mathematically related to the 'market entry condition' in Caplin and Dean [2015].

absolute welfare gains from the new addition can be larger in the full menu. As such, previously unchosen actions may affect *how often* and in which contingencies the new action is implemented.

Theorem 3. *Let* $\boldsymbol{\alpha}$ *denote the ignorance equivalent of RI problem* $(\mathcal{A}, \boldsymbol{\pi}, c)$ *that satisfies (C5). The following hold for any payoff vector* $\boldsymbol{a}^+ \in \mathbb{R}^I$*:*

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

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

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

Corollary 5. *Consider an RI problem* $(\mathcal{A}, \boldsymbol{\pi}, c)$ *that satisfies (C5). An available payoff vector* $\boldsymbol{a} \in \mathcal{A}$ *is implemented with probability* 1 *if and only if*

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

Proof. The core of the proof is an inductive application of Theorem 3(a), as detailed in Appendix A.3. □

If an available action $\boldsymbol{a} \in \mathcal{A}$ is implemented without learning, then it is also the ignorance equivalent of the RI problem; conversely, if the ignorance equivalent belongs to the original menu $\mathcal{A}$, then no learning is optimal.

There is another sense in which anchor actions describe a 'latent' consideration set. In Corollary 4(c), we pointed out that appropriate changes in the prior can induce the agent to implement any anchor action with positive probability. The same is true if we keep the prior fixed and instead introduce a new action to the menu.

Theorem 4. *Fix any RI problem* $(\mathcal{A}, \boldsymbol{\pi}, c)$ *that satisfies (C5) and consider anchor action* $\boldsymbol{a} \in \mathcal{A} \cap \bar{\mathcal{A}}$*. There exists a payoff vector* $\boldsymbol{a}^+ \in \mathbb{R}^I$ *such that it is optimal to implement* $\boldsymbol{a}$ *with positive probability in RI problem* $(\mathcal{A} \cup \{\boldsymbol{a}^+\}, \boldsymbol{\pi}, c)$.

By Corollary 4(d), each non-anchor action is dominated by a learning strategy across all beliefs. This implies that the converse to Theorem 4 also holds, and it is exactly the anchor actions that can be activated through menu expansion in the sense of Matějka and McKay [2015].

It is worth highlighting that it is too vague to ask whether a novel action $\boldsymbol{a}^+$ would be attractive to an agent facing RI problem $(\mathcal{A}, \boldsymbol{\pi}, c)$: The answer may be negative if only $\boldsymbol{a}^+$ is added to the menu, but positive if $\boldsymbol{a}^+$ is added alongside other actions. Indeed, learning introduces strong complementarities between actions and the attractiveness of an individual action depends on the agent's other options. If $\boldsymbol{a}^+$ offers higher payoff under some extreme belief than all existing options in $\mathcal{A}$, it increases the consumption gain of sufficiently informative signals. Still, the cost of these signals may remain prohibitive until the agent also receives 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 $\mathcal{A}' \subseteq \mathcal{A}$, the learning-proof menu is the answer: If $\boldsymbol{a}^+$ is located on or above $\bar{\mathcal{A}}$, the action becomes an anchor in $\mathcal{A} \cup \{\boldsymbol{a}^+\}$, and as such can be activated by the simultaneous addition of some complementary action. Conversely, if $\boldsymbol{a}^+$ is located below $\bar{\mathcal{A}}$, activation is impossible.

4.2 Value of information

Since the learning-proof menu transforms the RI model into a standard expected utility maximization problem, it is particularly well suited to study the welfare impacts of exogenous information.

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

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

for access to a signal $\mathcal{S} = \langle S, \boldsymbol{q} \rangle$ that does not definitively rule out any state, so that the posterior after observing each $s \in S$ has full support, $\boldsymbol{\pi}^s \in \mathsf{int}(\Delta\mathcal{I})$.

In particular, this means that at any prior $\boldsymbol{\pi}$, the local geometry of the learning-proof menu determines whether the agent benefits from free access to noisy information. The ignorance equivalent is located at the intersection of $\bar{\mathcal{A}}$ with the supporting hyperplane orthogonal to $\boldsymbol{\pi}$, and the tangent space of the surface $\bar{\mathcal{A}}$ at this point, 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 for small enough changes in the prior, the supporting hyperplane merely rotates around V but remains tangent at $\boldsymbol{\alpha}$. In other words, 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 implementation strategy. In particular, if this is true for all signal realizations, then the agent does not benefit from the noisy information, nor is she persuaded to behave any differently; her conditional choices remain exactly as they were before.

4.3 Eliciting beliefs

The learning-proof menu can also offer insight for an analyst who seeks to identify an agent's unknown prior belief. For example, a marketing company may want to gauge the subjective beliefs about a new product within a representative 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. Tsakas [2020] shows that in order to obtain truth-telling but discourage learning, the analyst needs to tailor the state-dependent payments to the agent's information cost function.

This design problem is tightly connected to the learning-proof menu. For concreteness, suppose that the analyst is familiar with the agent's information cost (c) and her state-dependent payoffs from a set of 'outside options' ($\mathcal{A}$). By offering the agent the full learning-proof menu $\bar{\mathcal{A}}$, the analyst removes all learning incentives.[12] Moreover, the analyst's task is much simpler, as the agent's choice now depends on her prior in the same way it does in a standard expected utility maximization problem. Figure 2 can be read this way: If the agent can select any point on the solid black line, her unconditional choice of $\boldsymbol{\alpha}$ reveals that her prior is $\boldsymbol{\pi}$, orthogonal to the unique tangent hyperplane. More generally, the agent's prior is uniquely pinned

[12] We here assume that whenever the agent is exactly indifferent across multiple learning strategies, she picks the least informative one — for example, if the agent is exactly indifferent between unconditional implementation and learning, she picks ignorance. The analyst can always place payoffs slightly above the surface of the learning-proof menu to favor no learning.

down when there is a unique tangent hyperplane at the chosen action, and with some pooling at the kinks (or points of non-differentiability) of $\bar{\mathcal{A}}$. It is particularly convenient that the analyst needs only a single observation to identify the set of priors consistent with the agent's choices.

In practice, the challenges of belief elicitation are many and beyond the scope of this paper. However, the learning-proof menu does prove useful even if the analyst may not have a perfect grasp of the agent's state-dependent payoffs or learning costs, or if design constraints limit the number of additional options that can be offered to the agent. Starting with the latter, suppose that the agent's original options correspond to the three actions marked as black dots in Figure 2, and the analyst can add just one intermediate option. What conditional payoffs should he choose? Clearly, payoffs below the learning-proof menu are useless because the agent would never choose such an option by Corollary 4(d). Payoffs within the learning-proof menu, say $\boldsymbol{\alpha}$ as depicted in Figure 2, are a least-cost way to increase the accuracy of the belief-elicitation because agents with priors close to $\boldsymbol{\pi}$ would now mostly select the new option. Arguably though, this experimental design is not very robust if the analyst is uncertain about the agent's payoff, as he then cannot pinpoint the exact location of the actions (black dots). Similarly, uncertainty about the agent's learning costs affects the hashed sets that mark dominated payoff vectors. Payoffs like $\boldsymbol{\alpha}$ can easily fall below the learning-proof menu even for small perturbations of the black dots or hashed sets. To gain robustness, the analyst can make the new option slightly more attractive; this increases costs but also dampens the impact of model misspecification. Yet, the analyst wants to exercise moderation: If he makes the new option too attractive, he risks removing one of the outside options as an anchor and thus rendering it useless for belief elicitation.

5 Application to Bilateral Contracting

The ignorance equivalent naturally emerges in strategic interactions between risk-neutral RI agents. Generalizing Example 1 from the introduction, we here define a two-player contracting game between two RI agents called Abigail and Bertrand.

Abigail is an RI decision maker with cost function c who has access to a menu $\mathcal{A} \subset \mathbb{R}^I$ of outside options, each $\boldsymbol{a} \in \mathcal{A}$ representing a vector of state-dependent monetary payoffs. If left to her own devices, Abigail thus faces the standard RI problem $(\mathcal{A}, \boldsymbol{\pi}, c)$. Bertrand on the other hand is a manager with weakly lower cost function $\tilde{c} \leq c$ who, if given a mandate by Abigail, can select from a weakly larger menu $\widetilde{\mathcal{A}} \supseteq \mathcal{A}$. Both agents initially share a common prior $\boldsymbol{\pi}$ about the state of the world. The setup fits a broad range of applications where a potential manager has an operational or informational advantage. In finance, a fund manager may invest in funds that are inaccessible to a retail consumer like Abigail; in production management, a major supplier may have access to technology that is out of reach for smaller businesses. In both situations, the manager may also have an informational advantage due to his connections and his experience in interpreting relevant evidence.

We want to identify the contract terms that emerge if the manager can make Abigail a take-it-or-leave-it offer, taking into account that the manager is free to acquire information both before designing contract terms and after obtaining the mandate, and that Abigail can condition her acceptance on privately acquired information.

Basic setup. We formally define this *outsourcing game* as follows:

- At time zero, manager Bertrand privately draws a signal $\mathcal{S}^0 = \langle S_0, \boldsymbol{q}^0 \rangle$ at cost $\tilde{c}(\mathcal{S}^0, \boldsymbol{\pi})$ and, upon observing its realization $s \in S_0$, publishes a menu of contracts $\mathcal{B}_s \subset \mathbb{R}^I$. Each $\boldsymbol{b} \in \mathcal{B}_s$ represents an agreement whereby Abigail gives the choice mandate to Bertrand in return for receiving a state-dependent transfer $\boldsymbol{b}$. We denote the set of all such menus by $\bar{B} = \{\mathcal{B}_s \mid s \in S_0\}$.

- At time one, Abigail observes a menu of contract terms $\mathcal{B} = \mathcal{B}_s \in \bar{\mathcal{B}}$, updates her belief to $\boldsymbol{\pi}^{\mathcal{B}}$ according to Bayes' rule, and follows a learning strategy $\mathcal{S}^{\mathcal{B}} = \langle \mathcal{A} \cup \mathcal{B}, \boldsymbol{q}^{\mathcal{B}} \rangle$ at cost $c(\mathcal{S}^{\mathcal{B}}, \boldsymbol{\pi}^{\mathcal{B}})$. If Abigail selects an outside option $\boldsymbol{a} \in \mathcal{A}$, the game ends and Bertrand receives a payoff of zero. If Abigail retains Bertrand's services by selecting $\boldsymbol{b} \in \mathcal{B}$, the game proceeds to time two.

- At time two, Bertrand updates his belief to $\boldsymbol{\pi}^{(s,\boldsymbol{b})}$ based on his time-zero draw $s \in S_0$ and Abigail's selection $\boldsymbol{b} \in \mathcal{B}_s$, and follows a learning strategy $\mathcal{S}^{(s,\boldsymbol{b})} = \langle \widetilde{\mathcal{A}}, \boldsymbol{q}^{(s,\boldsymbol{b})} \rangle$ at cost $\tilde{c}(\mathcal{S}^{(s,\boldsymbol{b})}, \boldsymbol{\pi}^{(s,\boldsymbol{b})})$. Once the state i is revealed, Bertrand transfers b_i to Abigail.

By unconditionally offering the ignorance equivalent $\boldsymbol{\alpha}$ of $(\mathcal{A}, \boldsymbol{\pi}, c)$, Bertrand incurs no initial learning costs and ensures that he always receives the mandate, since unconditional selection of $\boldsymbol{\alpha}$ is optimal in RI problem $(\mathcal{A} \cup \{\boldsymbol{\alpha}\}, \boldsymbol{\pi}, c)$. At time two, Bertrand then faces the RI problem $(\widetilde{\mathcal{A}} - \boldsymbol{\alpha}, \boldsymbol{\pi}, \tilde{c})$, and achieves an overall payoff of $W(\widetilde{\mathcal{A}}, \boldsymbol{\pi}, \tilde{c}) - \boldsymbol{\pi} \cdot \boldsymbol{\alpha}$. The next result shows unconditional agreement on a contract with terms $\boldsymbol{\alpha}$ constitutes a Perfect Bayesian Equilibrium of the game.

Theorem 5.A. *The outsourcing game admits a manager-preferred Perfect Bayesian Equilibrium where the ignorance equivalent* $\boldsymbol{\alpha}$ *is offered and accepted unconditionally.*

Proof. See Online Appendix B. □

The result relies on three key factors: First, offering Abigail the ignorance equivalent transfer is the least-cost way for Bertrand to get hired unconditionally. Second, any tailoring of his menu offer only serves as a free signal to Abigail and increases the required transfers. Third, Bertrand indirectly pays for any information revealed by Abigail's contract choice, as she only acquires additional information if the extra transfers raise her consumption utility. Since it is weakly cheaper for him to obtain the same information directly, contracting on the ignorance equivalent is Bertrand's preferred equilibrium.

Information shocks. We now generalize the game by allowing for free public information to arrive at the beginning of each time period, and refer to the resulting game as *outsourcing with information shocks*. Some shocks affect the equilibrium more than others: Any information that arrives at time zero merely alters the initial prior $\boldsymbol{\pi}$. Similarly, any information that arrives at time two merely increases Bertrand's net payoff, but does not affect the expected transfer between the agents or any strategic decision in earlier rounds. The most interesting case is information that arrives at time one, after contract terms have been offered but before Abigail accepts. If the free information is precise enough, Abigail no longer unconditionally accepts the ignorance-equivalent contract offer, leading to an adverse-selection problem for Bertrand.

Luckily, Bertrand can hedge against such information shocks by unconditionally offering the entire learning-proof menu $\bar{\mathcal{A}}$ for menu $\mathcal{A}$ and cost c in period zero. At time one, Abigail observes a draw from the free public signal, updates her belief to $\boldsymbol{\rho}$ and then self-selects the terms corresponding to the ignorance equivalent of $(\mathcal{A}, \boldsymbol{\rho}, c)$. As such, this menu offer ensures that Abigail forgoes all costly information acquisition and Bertrand always receives the mandate.

Theorem 5.B. *The outsourcing game with information shocks admits a manager-preferred Perfect Bayesian Equilibrium where the learning-proof menu $\bar{A}$ is offered unconditionally and contracting happens at the ignorance equivalent under the public time-one posterior.*

Proof. See Online Appendix B. □

The key ideas are the same as above: If Bertrand leaves any learning up to Abigail, he will indirectly pay for her information-acquisition costs either by missing out on a contract or by paying higher transfers. Offering the learning-proof menu lets Bertrand do any learning 'in house' while paying only minimal rents to Abigail.

It is worth noting that our analysis assumes that agents fully incorporate strategic information in the game. If signal costs are interpreted purely as information-acquisition costs, it is natural to assume that agents incorporate all information that can be inferred from previous choices in the game. On the other hand, if signals costs are at least partly capturing information-processing constraints, agents may not fully incorporate this freely available information. If they do not, this generates a difference in beliefs for the same publicly observable information, and then there exist terms that both Abigail and Bertrand consider strictly preferable to the ignorance equivalent $\boldsymbol{\alpha}$.

6 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 strategic games where one agent designs the menu of the other, this ignorance-equivalent action often emerges as an actual option. The same is true in experimental settings where 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 apply standard expected-utility techniques to problems with learning.

A Proofs

A.1 Properties of the cost function

Lemma A.1. *Under (C2), it is without loss of generality in RI problem* (RI) *to restrict attention to learning strategies $\langle \mathcal{A}, \boldsymbol{q} \rangle$.*

Proof. For any signal $\mathcal{S} = \langle S, \boldsymbol{q} \rangle$ and conditional selection $\boldsymbol{a}^s \in \mathcal{A}$ for each $s \in S$, we can define the learning strategy $\hat{\mathcal{S}} = \langle \mathcal{A}, \hat{\boldsymbol{q}} \rangle$ with $\hat{q}_i(\boldsymbol{a}) = \sum_{s \in S: \boldsymbol{a}^s = \boldsymbol{a}} q_i(s)$ for each $\boldsymbol{a} \in \mathcal{A}$. 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.2. *Under (C1) and (C2), the indirect utility W is continuous in the prior belief at any interior prior $\boldsymbol{\pi}$. Moreover, there exists an upper hemicontinuous correspondence $Q^* : \mathsf{int}(\Delta \mathcal{I}) \to (\Delta \mathcal{A})^I$ with nonempty and compact values, such that a learning strategy $\mathcal{S} = \langle \mathcal{A}, \boldsymbol{q} \rangle$ is optimal in RI problem $(\mathcal{A}, \boldsymbol{\pi}, c)$ if and only if $\boldsymbol{q} \in Q^*(\boldsymbol{\pi})$.*

Proof. Since $\boldsymbol{\pi}$ is interior, there exists $\delta > 0$ such that the closed ball $\bar{B}_\delta(\boldsymbol{\pi})$ is strictly in the interior of the simplex. Pick any action $\boldsymbol{a} \in \mathcal{A}$ and consider the straightforward strategy where the agent picks an uninformative signal $\mathcal{S}^0 = \langle \mathcal{A}, \delta_{\boldsymbol{a}} \mathbf{1} \rangle$ that recommends unconditional implementation of $\boldsymbol{a}$. For any $\boldsymbol{\pi}' \in B_\delta(\boldsymbol{\pi})$, the achieved welfare $\boldsymbol{\pi}' \cdot \boldsymbol{a} - c(\mathcal{S}^0, \boldsymbol{\pi}')$ yields a lower bound on the optimal welfare. At the same time, the consumption utility of the agent is bounded above in each state i by $\bar{a}_i = \max_{\boldsymbol{a} \in \mathcal{A}} a_i$. As a consequence, $\mathcal{S}^0$ is strictly preferable to any signal $\mathcal{S}$ with cost $c(\mathcal{S}, \boldsymbol{\pi}') > \boldsymbol{\pi}' \cdot (\bar{\boldsymbol{a}} - \boldsymbol{a}) + c(\mathcal{S}^0, \boldsymbol{\pi}')$.

Since it is without loss of generality to restrict attention to learning strategies (Lemma A.1), this allows us to restate the optimization problem locally as a choice over marginal probabilities

$$W(\mathcal{A}, \boldsymbol{\pi}', c) = \max_{\boldsymbol{q} \in Q(\boldsymbol{\pi}')} \sum_{\boldsymbol{a} \in \mathcal{A}} \sum_{i \in \mathcal{I}} \pi'_i q_i(\boldsymbol{a}) a_i - c(\langle \mathcal{A}, \boldsymbol{q} \rangle, \boldsymbol{\pi}') \qquad \forall \boldsymbol{\pi}' \in \bar{B}_\delta(\boldsymbol{\pi})$$

over a domain $Q(\boldsymbol{\pi}') = \left\{ \boldsymbol{q} \in (\Delta \mathcal{A})^I \mid c(\langle \mathcal{A}, \boldsymbol{q} \rangle, \boldsymbol{\pi}') \leq \boldsymbol{\pi}' \cdot (\bar{\boldsymbol{a}} - \boldsymbol{a}) + c(\mathcal{S}^0, \boldsymbol{\pi}') \right\}$ that is non-empty and compact. By continuity of the cost function (C1), the objective function is continuous and the correspondence Q is continuous with nonempty and compact values. The claim then follows by Berge's Theorem of the Maximum. $\square$

Lemma A.3. *Under (C2), indirect utility 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 direct signal 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.

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

Lemma A.4. *Suppose (C3) holds. Consider an optimal learning strategy $\mathcal{S} = \langle \mathcal{A}, \boldsymbol{q} \rangle$ to RI problem $(\mathcal{A}, \boldsymbol{\pi}, c)$. Let $\boldsymbol{a} \in \mathsf{support}(\boldsymbol{q})$ be an action that is implemented with positive probability, and $\boldsymbol{\pi}^{\boldsymbol{a}}$ the associated posterior belief, $\pi_i^{\boldsymbol{a}} = \frac{\pi_i q_i(\boldsymbol{a})}{\boldsymbol{\pi} \cdot \boldsymbol{q}(\boldsymbol{a})}$. In RI problem $(\mathcal{A}, \boldsymbol{\pi}^{\boldsymbol{a}}, c)$, unconditional implementation of $\boldsymbol{a}$ is optimal, $W(\mathcal{A}, \boldsymbol{\pi}^{\boldsymbol{a}}, c) = \boldsymbol{\pi}^{\boldsymbol{a}} \cdot \boldsymbol{a}$.*

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

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

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

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

which implies that this sequential strategy achieves strictly higher utility than signal

$\mathcal{S}$ alone. By (C3), the same is true for its one-shot equivalent, contradicting the optimality of $\mathcal{S}$ in $(\mathcal{A}, \boldsymbol{\pi}, c)$. $\square$

Lemma A.5. *Assumption (C4) implies that any pure-noise signal $\mathcal{S} = \langle S, \boldsymbol{q}\rangle$ with $q_i(s) \equiv q_j(s)\ \forall i, j \in \mathcal{I}$ and $s \in S$ is free.*

Proof. Since the marginal and conditional probabilities for each realization are the same across states, $\mathcal{S} = \mathcal{S}^{(\varepsilon,\boldsymbol{\pi})}$ for any ε. Consequently, condition (C4) requires that for any $\bar{c} > 0$, there exists $\varepsilon \in (0, 1)$ such that $c(\mathcal{S}, \boldsymbol{\pi}) \leq \bar{c}\varepsilon < \bar{c}$, and so by taking the limit $\bar{c} \to 0$, $c(\mathcal{S}, \boldsymbol{\pi}) = 0$. $\square$

Lemma A.6. *Consider any belief $\boldsymbol{\pi} \in \Delta\mathcal{I}$ and any two actions with the same expected utility, $\boldsymbol{a} \cdot \boldsymbol{\pi} = \boldsymbol{a}' \cdot \boldsymbol{\pi}$, that differ in at least one positive probability state. Under (C4), and any $L > 0$, there exists a learning strategy $\langle \{\boldsymbol{a}, \boldsymbol{a}'\}, \boldsymbol{q}\rangle$ such that the change in consumption utility outweighs the signal cost by more than a factor L,*

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

Proof. Consider a signal $\mathcal{S} = \langle \{\boldsymbol{a}, \boldsymbol{a}'\}, \boldsymbol{q}\rangle$ with $q_i(\boldsymbol{a}) = 1 - q_i(\boldsymbol{a}') \equiv \frac{1}{2} + \delta(a_i - a_i')$ for some $\delta > 0$ small enough such that all probabilities are nondegenerate. Now suppose that the agent follows the ε-precision dilution $\mathcal{S}^{(\varepsilon,\boldsymbol{\pi})}$. The change in expected consumption utility is linear in ε,

$$\begin{aligned}\sum_{i=1}^{I} \pi_i \left[q_i^{(\varepsilon,\boldsymbol{\pi})}(\boldsymbol{a})a_i + q_i^{(\varepsilon,\boldsymbol{\pi})}(\boldsymbol{a}')a_i'\right] - u &= \varepsilon \left[\sum_{i=1}^{I} \pi_i \left[q_i(\boldsymbol{a})a_i + q_i(\boldsymbol{a}')a_i'\right] - u\right] \\ &= \varepsilon \left[\delta \sum_{i=1}^{I} \pi_i (a_i - a_i')^2\right] > 0.\end{aligned}$$

Letting $\bar{c} := \frac{\delta}{L}\sum_{i=1}^{I} \pi_i(a_i - a_i')^2$, (C4) implies that for some $\varepsilon \in (0, 1)$, the benefits outweigh the cost by more than a factor L since $c(\mathcal{S}^{(\varepsilon,\boldsymbol{\pi})}, \boldsymbol{\pi}) < \varepsilon\bar{c}$. $\square$

Lemma A.7. *Under any optimal posterior $\boldsymbol{\pi}^s$, all optimal actions are payoff-equivalent with probability one, $\boldsymbol{a}^1, \boldsymbol{a}^2 \in \arg\max_{\boldsymbol{a}\in\mathcal{A}} \boldsymbol{\pi}^s \cdot \boldsymbol{a}$ if and only if $a_i^1 = a_i^2$ whenever $\pi_i^s > 0$.*

Proof. By contradiction, suppose that there exist two actions $\boldsymbol{a}, \boldsymbol{a}' \in \mathcal{A}$ that both maximize expected utility under posterior $\boldsymbol{\pi}^s$ but differ in at least one positive-probability state. Let $\mathcal{S}$ be the learning strategy that satisfies the conditions of

Lemma A.6 for $L = 1$. Now suppose that the agent, after observing s, follows $\mathcal{S}$ rather than implementing one of the available actions. This sequential approach strictly improves the agent's welfare, contradicting the optimality of $\boldsymbol{\pi}^s$ by (C3). $\square$

Proof of Lemma 1: Continuity of c (C1) follows directly from that of the potential ϕ since

$$|c(\langle S, \tilde{\boldsymbol{q}}\rangle, \tilde{\boldsymbol{\pi}}) - c(\mathcal{S}, \boldsymbol{\pi})| \leq \sum_{s \in S} |(\boldsymbol{\pi} \cdot \tilde{\boldsymbol{q}}(s))\phi(\tilde{\boldsymbol{\pi}}^s) - (\boldsymbol{\pi} \cdot \boldsymbol{q}(s))\phi(\boldsymbol{\pi}^s)| + |\phi(\boldsymbol{\pi}) - \phi(\tilde{\boldsymbol{\pi}})|.$$

Bloedel and Zhong [2020] establish Blackwell monotonicity and prior-concavity (C2) for this class of cost functions, and show that any (UPS) cost satisfies indifference to sequential learning, implying that the inequalities in (C3) and (C5) both hold.

For the zero-marginal cost condition (C4), let $\boldsymbol{\pi}^s$ denote the posteriors after drawing s from $\mathcal{S} = \langle S, \boldsymbol{q}\rangle$ under prior $\boldsymbol{\pi}$. Observing s from the noisy signal $\mathcal{S}^{(\varepsilon,\boldsymbol{\pi})}$ leads to posterior $\varepsilon\boldsymbol{\pi}^s + (1-\varepsilon)\boldsymbol{\pi}$, but the marginal likelihood for each outcome s is unchanged. And since

$$\begin{aligned}
\lim_{\varepsilon \to 0} \frac{1}{\varepsilon} c(\mathcal{S}^{(\varepsilon,\boldsymbol{\pi})}, \boldsymbol{\pi}) &= \sum_{s \in S} (\boldsymbol{\pi} \cdot \boldsymbol{q}(s)) \lim_{\varepsilon \to 0} \frac{\phi(\varepsilon\boldsymbol{\pi}^s + (1-\varepsilon)\boldsymbol{\pi}) - \phi(\boldsymbol{\pi})}{\varepsilon} \\
&= \nabla\phi(\boldsymbol{\pi}) \cdot \left[\sum_{s \in S} (\boldsymbol{\pi} \cdot \boldsymbol{q}(s))(\boldsymbol{\pi}^s - \boldsymbol{\pi})\right] = \nabla\phi(\boldsymbol{\pi}) \cdot \mathbf{0} = 0,
\end{aligned}$$

for any $\bar{c} > 0$, $\exists \varepsilon > 0$ such that $c(\mathcal{S}^{(\varepsilon,\boldsymbol{\pi})}, \boldsymbol{\pi}) < \varepsilon\bar{c}$. $\square$

A.2 Ignorance equivalent

Preliminaries. The following mathematical property is key to many of our core results. It states that if the achievable expected payoff is limited by a convex function of the belief, it is without loss of generality to impose limits across all beliefs at once: Doing so does not reduce the achievable payoff under any prior.

Lemma A.8. *Let $\Phi : \Delta\mathcal{I} \to \mathbb{R} \cup \{\infty\}$ denote a convex function that is finite-valued over $\mathsf{int}(\Delta\mathcal{I})$. Then for any $\boldsymbol{\pi} \in \mathsf{int}(\Delta\mathcal{I})$, the optimization problem*

$$\arg\max \{\boldsymbol{\pi} \cdot \boldsymbol{x} \mid \boldsymbol{\rho} \cdot \boldsymbol{x} \leq \Phi(\boldsymbol{\rho}) \ \forall \boldsymbol{\rho} \in \Delta\mathcal{I}\}$$

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

Proof. Key to the proof is the finite intersection property, which states that a collection of subsets on a compact space has nonempty intersection if and only if the intersection of every finite subcollection is nonempty.

To apply the finite intersection property, we define first the compact set $C \subseteq \mathbb{R}^I$ as follows: Start with a finite set of priors $Q \subset \mathsf{int}(\Delta\mathcal{I})$ whose convex cone contains a closed ball around $\boldsymbol{\pi}$, $\bar{B}_\delta(\boldsymbol{\pi}) \subset \{\sum_{\boldsymbol{\rho}\in Q} z(\boldsymbol{\rho})\boldsymbol{\rho} \mid z : Q \to [0,\infty)\}$ for some $\delta > 0$. Let C be defined as

$$C := \left\{\boldsymbol{x} \in \mathbb{R}^I \mid \boldsymbol{\pi}\cdot\boldsymbol{x} \geq \Phi(\boldsymbol{\pi}) \text{ and } \boldsymbol{\rho}\cdot\boldsymbol{x} \leq \Phi(\boldsymbol{\rho})\ \forall \boldsymbol{\rho} \in Q\right\}.$$

As the intersection of closed half-spaces, C convex and closed. Further, a known result in geometry states that any such intersection can be written as the Minkowski sum of a convex bounded polytope and a cone [e.g. Theorem 1.2 in Ziegler, 2012]. This implies that C is compact if and only if there does not exist a point $\boldsymbol{x}^0 \in C$ and a direction $\boldsymbol{v} \in \mathbb{R}^I \setminus \{\boldsymbol{0}\}$ such that the ray $\{\boldsymbol{x}^0 + z\boldsymbol{v} \mid z \geq 0\}$ is entirely contained in C. There exists no such direction here: Without loss of generality, assume that $\|\boldsymbol{v}\| = \delta$. If $\boldsymbol{\rho}\cdot\boldsymbol{v} > 0$ for any $\boldsymbol{\rho} \in Q$, the constraint $\boldsymbol{\rho}\cdot(\boldsymbol{x}^0 + z\boldsymbol{v}) \leq \Phi(\boldsymbol{\rho})$ is violated for z large enough. Conversely, if $\boldsymbol{\rho}\cdot\boldsymbol{v} \leq 0$ for all $\boldsymbol{\rho} \in Q$, the same holds for any $\boldsymbol{\rho} \in \bar{B}_\delta(\boldsymbol{\pi})$ and in particular for $\boldsymbol{\pi} + \boldsymbol{v}$. From there, it then follows that $\boldsymbol{\pi}\cdot\boldsymbol{v} = (\boldsymbol{\pi}+\boldsymbol{v})\cdot\boldsymbol{v} - \boldsymbol{v}\cdot\boldsymbol{v} \leq 0 - \boldsymbol{v}\cdot\boldsymbol{v} < 0$. This in turn implies that the constraint $\boldsymbol{\pi}\cdot(\boldsymbol{x}^0 + z\boldsymbol{v}) \geq \Phi(\boldsymbol{\pi})$ is violated for z large enough. Thus, C is compact.

Consider now any arbitrary finite set of priors $R \subseteq \Delta\mathcal{I}$, and denote any intersection of lower half-spaces as

$$X_R := \{\boldsymbol{x} \mid \boldsymbol{\rho}\cdot\boldsymbol{x} \leq \Phi(\boldsymbol{\rho})\ \forall \boldsymbol{\rho} \in R\}.$$

To establish that the intersection $C \cap X_R$ is nonempty, we look at the larger cone $X_{R\cup Q}$. This closed cone contains the point $z\boldsymbol{1}$ for $z \leq \min_{\boldsymbol{\rho}\in R\cup Q} \Phi(\boldsymbol{\rho})$ large enough, and thus is nonempty. Moreover, since $\boldsymbol{\pi}$ can be written as a convex combination over Q, $\boldsymbol{\pi} = \sum_{\boldsymbol{\rho}\in Q} z(\boldsymbol{\rho})\boldsymbol{\rho}$, the set is bounded above in direction $\boldsymbol{\pi}$ by $\sum_{\boldsymbol{\rho}\in Q} z(\boldsymbol{\rho})\Phi(\boldsymbol{\rho})$. As a result, the set contains an extremal point $\boldsymbol{x}^0 \in \arg\max_{\boldsymbol{x}\in X_{R\cup Q}} \boldsymbol{\pi}\cdot\boldsymbol{x}$. Let R^0 identify which half-spaces are binding at $\boldsymbol{x}^0$, $R^0 := \{\boldsymbol{\rho} \in R \cup Q \mid \boldsymbol{\rho}\cdot\boldsymbol{x}^0 = \Phi(\boldsymbol{\rho})\}$. To show that $\boldsymbol{x}^0 \in C$, we establish two properties:

(i) $\boldsymbol{\pi}$ is contained in the convex hull of R^0.

By contradiction, suppose $\{\boldsymbol{\pi}\} \cap \mathsf{conv.hull}(R^0) = \emptyset$. The separating hyperplane theorem then implies the existence of a nonzero vector $\boldsymbol{v}$ such that $\boldsymbol{v} \cdot \boldsymbol{\pi} > 0$ and $\boldsymbol{v} \cdot \boldsymbol{\rho} < 0$ for all $\boldsymbol{\rho} \in R^0$, or equivalently, $\boldsymbol{\rho} \cdot (\boldsymbol{x}^0 + \varepsilon \boldsymbol{v}) \leq \Phi(\boldsymbol{\rho})$ for all $\boldsymbol{\rho} \in X_{R^0}$ and all $\varepsilon \geq 0$.[13] The finitely many non-binding inequalities $\boldsymbol{\rho} \cdot (\boldsymbol{x}^0 + \varepsilon \boldsymbol{v}) < \Phi(\boldsymbol{\rho})$ for $\boldsymbol{\rho} \in (R \cup Q) \setminus R^0$ are all maintained for $\varepsilon > 0$ small enough. In other words, the point $\boldsymbol{x}^0 + \varepsilon \boldsymbol{v}$ is contained in $X_{R \cup Q}$ yet achieves strictly higher expected utility than $\boldsymbol{x}^0$ under $\boldsymbol{\pi}$, contradicting the optimality of $\boldsymbol{x}^0$.

(ii) $\boldsymbol{\pi} \cdot \boldsymbol{x}^0 \geq \Phi(\boldsymbol{\pi})$.

Write $\boldsymbol{\pi}$ as a convex combination $\sum_{\boldsymbol{\rho} \in R^0} m(\boldsymbol{\rho})\boldsymbol{\rho}$, and note that

$$\boldsymbol{\pi} \cdot \boldsymbol{x}^0 = \sum_{\boldsymbol{\rho} \in R^0} m(\boldsymbol{\rho})(\boldsymbol{\rho} \cdot \boldsymbol{x}^0) = \sum_{\boldsymbol{\rho} \in R^0} m(\boldsymbol{\rho})\Phi(\boldsymbol{\rho}) \geq \Phi(\boldsymbol{\pi})$$

by definition of R^0 and convexity of Φ.

From the last observation, it follows that $C \cap X_R = \{\boldsymbol{x} \mid \boldsymbol{\pi} \cdot \boldsymbol{x} \geq \Phi(x)\} \cap X_{R \cup Q}$ contains $\boldsymbol{x}^0$ and thus is nonempty.

Applying now the finite intersection property, there exists a point $\boldsymbol{x}^{(\Phi,\boldsymbol{\pi})}$ in the intersection $C \cap X_{\Delta \mathcal{I}}$. By virtue of belonging to C, $\boldsymbol{\pi} \cdot \boldsymbol{x}^{(\Phi,\boldsymbol{\pi})}$ is weakly larger than $\Phi(\boldsymbol{\pi})$, and by virtue of belonging to $X_{\{\boldsymbol{\pi}\}} \supseteq X_{\Delta \mathcal{I}}$, it is weakly smaller. Together, the two inequalities imply that the two are equal and $\boldsymbol{x}^{(\Phi,\boldsymbol{\pi})}$ is maximal in direction $\boldsymbol{\pi}$ over $X_{\Delta \mathcal{I}}$. $\square$

Existence and Uniqueness. By following a learning strategy $\mathcal{S} = \langle \mathcal{A}, \boldsymbol{q} \rangle$, the agent achieves expected consumption utility

$$a_i^{\mathcal{S}} := \sum_{\boldsymbol{a} \in \mathcal{A}} q_i(\boldsymbol{a}) a_i \tag{3}$$

conditional on state i, irrespective of her prior belief. We define the set of $\mathcal{S}$-dominated payoff vectors as

$$\bar{\mathcal{A}}^{\mathcal{S}} := \bigcap_{\boldsymbol{\rho} \in \Delta \mathcal{I}} \left\{ \boldsymbol{x} \in \mathbb{R}^I \mid \boldsymbol{\rho} \cdot \boldsymbol{x} \leq \boldsymbol{\rho} \cdot \boldsymbol{a}^{\mathcal{S}} - c(\mathcal{S}, \boldsymbol{\rho}) \right\}. \tag{4}$$

[13] Formally, the separating hyperplane theorem assures that there exists a vector $\boldsymbol{w} \in \mathbb{R}^I$ and real numbers $\underline{c} < \bar{c}$ that separate the two disjoint nonempty compact convex sets, $\boldsymbol{w} \cdot \boldsymbol{\pi} \geq \bar{c}$ and $\boldsymbol{w} \cdot \boldsymbol{\rho} \leq \underline{c}$ for all $\boldsymbol{\rho} \in R^0$. Letting $\boldsymbol{v} = \boldsymbol{w} - \frac{\underline{c}+\bar{c}}{2}\mathbf{1}$, note that $\boldsymbol{v} \cdot \boldsymbol{\pi} = \boldsymbol{w} \cdot \boldsymbol{\pi} - \frac{\underline{c}+\bar{c}}{2} > 0$ and $\boldsymbol{v} \cdot \boldsymbol{\rho} = \boldsymbol{w} \cdot \boldsymbol{\rho} - \frac{\underline{c}+\bar{c}}{2} < 0$.

The decision maker would always weakly prefer to follow the costly advice of signal $\mathcal{S}$ rather than unconditionally implement $\boldsymbol{a} \in \bar{\mathcal{A}}^{\mathcal{S}}$, no matter her belief. Still, the previous result implies that for every belief, one of these payoff vectors leaves the agent exactly indifferent. Indeed, the limit function $\Phi(\boldsymbol{\rho}) = \boldsymbol{\rho} \cdot \boldsymbol{a}^{\mathcal{S}} - c(\mathcal{S}, \boldsymbol{\rho})$ is convex by prior-concavity of the cost function (C2). Lemma A.8 thus ensures that no matter the agent's prior, there always exists a $\mathcal{S}$-dominated payoff vector that, if implemented unconditionally, achieves the same welfare as $\mathcal{S}$.

Corollary 7. *Under (C2), for any learning strategy $\mathcal{S}$ and any interior belief $\boldsymbol{\pi} \in \mathit{int}(\Delta\mathcal{I})$, the payoff vector $\boldsymbol{a}^{(\mathcal{S},\boldsymbol{\pi})} = \arg\max_{\boldsymbol{x} \in \bar{\mathcal{A}}^{\mathcal{S}}} \boldsymbol{\pi} \cdot \boldsymbol{x}$ exists and satisfies*

$$\boldsymbol{\pi} \cdot \boldsymbol{a}^{(\mathcal{S},\boldsymbol{\pi})} = \boldsymbol{\pi} \cdot \boldsymbol{a}^{\mathcal{S}} - c(\mathcal{S}, \boldsymbol{\pi}). \tag{5}$$

Of particular interest is the case where $\mathcal{S}$ is optimal under a specific prior $\boldsymbol{\pi}$, in which case Corollary 7 asserts that there exists a point $\boldsymbol{\alpha} \in \bar{\mathcal{A}}^{\mathcal{S}}$ 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.2). By Corollary 7, there exists a point $\boldsymbol{\alpha} \in \bar{\mathcal{A}}^{\mathcal{S}}$ such that $\boldsymbol{\pi} \cdot \boldsymbol{\alpha} = W(\mathcal{A}, \boldsymbol{\pi}, c)$. (In other words, the inequality in Definition 2 holds and binds at prior $\boldsymbol{\pi}$.) We now show that $\boldsymbol{\alpha}$ is an ignorance equivalent of $(\mathcal{A}, \boldsymbol{\pi}, c)$.

(i) $W(\{\boldsymbol{\alpha}\}, \boldsymbol{\pi}, c) \geq W(\mathcal{A}, \boldsymbol{\pi}, c)$.

This follows simply because an agent faced with RI problem $(\{\boldsymbol{\alpha}\}, \boldsymbol{\pi}, c)$ can achieve utility $\boldsymbol{\pi} \cdot \boldsymbol{\alpha} = W(\mathcal{A}, \boldsymbol{\pi}, c)$ by implementing $\boldsymbol{\alpha}$ unconditionally.

(ii) $W(A \cup \{\boldsymbol{\alpha}\}, \boldsymbol{\pi}, c) \leq W(\mathcal{A}, \boldsymbol{\pi}, c)$.

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 $\boldsymbol{q}(\boldsymbol{\alpha}) = \mathbf{0}$, then the strategy does not rely on the presence of $\boldsymbol{\alpha}$, and is thus also feasible in problem $(\mathcal{A}, \boldsymbol{\pi}, c)$. If $\boldsymbol{q}(\boldsymbol{\alpha}) > \mathbf{0}$, let $\boldsymbol{\rho}$ denote the agent's posterior belief upon observing realization $\boldsymbol{\alpha}$. Consider the sequential strategy where the agent draws $\tilde{\mathcal{S}}$ and follows its recommendation except for realization $\boldsymbol{\alpha}$, when she instead draws and follows $\mathcal{S}$. The only change in the agent's payoff occurs conditional on realization $\boldsymbol{\alpha}$, when she achieves expected utility $\boldsymbol{\rho} \cdot \boldsymbol{a}^{\mathcal{S}} - c(\mathcal{S}, \boldsymbol{\rho})$ rather than $\boldsymbol{\rho} \cdot \boldsymbol{\alpha}$. Since $\boldsymbol{\alpha} \in \bar{\mathcal{A}}^{\mathcal{S}}$ and hence $\boldsymbol{\rho} \cdot \boldsymbol{\alpha} \leq \boldsymbol{\rho} \cdot \boldsymbol{a}^{\mathcal{S}} - c(\mathcal{S}, \boldsymbol{\rho})$, the sequential strategy weakly increases welfare. A one-shot implementation of this same strategy is weakly cheaper by (C3) and thus forms a lower bound for $W(\mathcal{A}, \boldsymbol{\pi}, c)$ that weakly exceeds $W(\mathcal{A} \cup \{\boldsymbol{\alpha}\}, \boldsymbol{\pi}, c)$.

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

$$u := W(\{\boldsymbol{\alpha}^1\}, \boldsymbol{\pi}, c) = W(\mathcal{A}, \boldsymbol{\pi}, c) = W(\{\boldsymbol{\alpha}^2\}, \boldsymbol{\pi}, c), \tag{6}$$

the two payoff vectors achieve the same expected utility u.

By Lemma A.6, there exists a signal $\mathcal{S}^0 = \langle \{\boldsymbol{\alpha}^1, \boldsymbol{\alpha}^2\}, \boldsymbol{q}^0 \rangle$ such that

$$\sum_{i=1}^{I} \pi_i (q_i^0(\boldsymbol{\alpha}^1)\alpha_i^1 + q_i^0(\boldsymbol{\alpha}^2)\alpha_i^2) - u > 2c(\mathcal{S}^0, \boldsymbol{\pi}). \tag{7}$$

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

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

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

$$\begin{aligned} V^1 + V^2 = & \sum_{i=1}^{I} \pi_i \left[q_i^0(\boldsymbol{\alpha}^1)\alpha_i^1 + q_i^0(\boldsymbol{\alpha}^2)\alpha_i^2\right] - 2c(\mathcal{S}^0, \boldsymbol{\pi}) \\ & + (\boldsymbol{q}^0(\boldsymbol{\alpha}^1) \cdot \boldsymbol{\pi}) W(\mathcal{A}, \boldsymbol{\pi}^{\boldsymbol{\alpha}^1}, c) + (\boldsymbol{q}^0(\boldsymbol{\alpha}^2) \cdot \boldsymbol{\pi}) W(\mathcal{A}, \boldsymbol{\pi}^{\boldsymbol{\alpha}^2}, c). \end{aligned} \tag{8}$$

The first line is strictly larger than $W(\mathcal{A}, \boldsymbol{\pi}, c)$ by Equation (7), and the second is weakly larger than that by prior-convexity of W (Lemma A.3). 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. □

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 $\boldsymbol{\alpha} \precsim \mathcal{S}$ 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. □

Proof of Corollary 2: Consider a sequence of beliefs $\{\boldsymbol{\pi}^n\}_{n=0}^{\infty}$ that converges to a prior $\boldsymbol{\pi}^0 \in \mathsf{int}(\Delta\mathcal{I})$, 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.2, 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 $\mathcal{S}^0 = \lim_{k\to\infty} \mathcal{S}^{n_k}$ 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,[14] and hence $\{\boldsymbol{\alpha}^n\}_{n=0}^{\infty}$ itself, converges to $\boldsymbol{\alpha}^0$. □

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

[14] By 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 subsubsequence, but its limit payoff vector must be different from $\boldsymbol{\alpha}^0$.

relying on $\mathcal{S}$ instead. The one-shot implementation $\mathcal{S}^{\boldsymbol{\rho}}$ of this strategy is admissible in $(\mathcal{A}, \boldsymbol{\rho}, c)$, yet achieves weakly higher welfare than $\mathcal{S}^+$ by (C3). This implies

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

with the opposite inequality binding by menu-monotonicity. In particular, $\mathcal{S}^{\boldsymbol{\rho}}$ 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 equivalent for both problems is thus equal to the unique $\arg\max_{\boldsymbol{x} \precsim \mathcal{S}^{\boldsymbol{\rho}}} \boldsymbol{\pi} \cdot \boldsymbol{x}$. □

A.3 Learning-Proof Menu

Proof of Theorem 2: Under any prior $\boldsymbol{\pi} \in \mathsf{int}(\Delta\mathcal{I})$, $\bar{\mathcal{A}}$ needs to contain its own ignorance equivalents $(\bar{\mathcal{A}}, \boldsymbol{\pi}, c)$ by property (a). Moreover, this ignorance equivalent has to be equal to that of $(\mathcal{A}, \boldsymbol{\pi}, c)$ by property (b). As such, the elements included in $\bar{\mathcal{A}}$ per Definition 3 are jointly required by the two conditions, and no smaller set can satisfy both.

Conversely, fix any interior belief $\boldsymbol{\pi}$ and note that unconditional implementation of $\boldsymbol{\alpha}^{(\mathcal{A}, \boldsymbol{\pi}, c)}$ is available to the agent and achieves welfare $W(\mathcal{A}, \boldsymbol{\pi}, c)$. Suppose by contradiction that the agent can do strictly better by relying on a learning strategy $\mathcal{S}^*$ that recommends actions $\boldsymbol{a} \in \bar{\mathcal{A}}$. At any realized posterior $\boldsymbol{\rho}$, unconditional implementation of $\boldsymbol{\alpha}^{(\mathcal{A}, \boldsymbol{\rho}, c)}$ achieves weakly higher expected utility than any other available action $\boldsymbol{a} \in \bar{\mathcal{A}}$ by Corollary 3, and this same expected utility can be achieved by instead drawing and following an optimal learning strategy for RI problem $(\mathcal{A}, \boldsymbol{\rho}, c)$ by Definition 1. By implementing the one-shot version of this learning strategy, the agent implements only actions from menu $\mathcal{A}$ yet obtains a welfare above $W(\mathcal{A}, \boldsymbol{\pi}, c)$, contradicting the optimality of W. In other words, unconditional implementation of $\boldsymbol{\alpha}^{(\mathcal{A}, \boldsymbol{\pi}, c)} \in \bar{\mathcal{A}}$ is optimal in RI problem $(\bar{\mathcal{A}}, \boldsymbol{\pi}, c)$, establishing that the ignorance equivalent of $(\mathcal{A}, \boldsymbol{\pi}, c)$ also represents the unique ignorance equivalent of $(\bar{\mathcal{A}}, \boldsymbol{\pi}, c)$. □

Proof of Corollary 4: We proceed in steps.

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

(b) $\Leftrightarrow$ (c): Since unconditional implementation implies that $\boldsymbol{a}$ is chosen with positive probability, (b) trivially implies (c). Lemma A.4 formalizes the converse claim.

(b) $\Rightarrow$ (d): Proving the contrapositive claim, assume also that there exists $\boldsymbol{\rho} \in \mathsf{int}(\Delta\mathcal{I})$ such that $\boldsymbol{\alpha}^{(\mathcal{A},\boldsymbol{\rho},c)} > \boldsymbol{a}$. In particular, for any prior $\boldsymbol{\pi}$, we have $\boldsymbol{\pi} \cdot \boldsymbol{a} < \boldsymbol{\pi} \cdot \boldsymbol{\alpha}^{(\mathcal{A},\boldsymbol{\rho},c)}$. By self-selection of the ignorance equivalent (Corollary 3), adding $\boldsymbol{\alpha}^{(\mathcal{A},\boldsymbol{\rho},c)}$ 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 $\boldsymbol{\pi}$,

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

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

(d) $\Rightarrow$ (a): Proving the contrapositive claim, assume that $\boldsymbol{a}$ 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 (1), $\boldsymbol{a}$ lies strictly below each individual half-space. As a consequence, the ray $\{\boldsymbol{a} + t\mathbf{1} \mid t \geq 0\}$ 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. $\square$

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

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

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

$$\bar{\mathcal{A}}^{\boldsymbol{\pi}} = \left\{\boldsymbol{x} \in \mathbb{R}^I \mid W(\mathcal{A} \cup \{\boldsymbol{x}\}, \boldsymbol{\pi}, c) \leq W(\mathcal{A}, \boldsymbol{\pi}, c)\right\} \tag{9}$$

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 $\bar{\mathcal{A}}^{\boldsymbol{\pi}}$ 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 $(\{\boldsymbol{\alpha}, \boldsymbol{a}^+\}, \boldsymbol{\pi}, c)$ 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 $\mathcal{S}^1$. By the now-familiar argument,[15] this weakly enhances welfare, and implies $W(\mathcal{A} \cup \{\boldsymbol{a}^+\}, \boldsymbol{\pi}, c) \geq W(\{\boldsymbol{\alpha}, \boldsymbol{a}^+\}, \boldsymbol{\pi}, c)$.

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

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

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,

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

To do so, let $\mathcal{S}^+ = \langle \mathcal{A} \cup \{\boldsymbol{a}^+\}, \boldsymbol{q}^+ \rangle$ denote an optimal learning strategy for RI problem $(\mathcal{A} \cup \{\boldsymbol{a}^+\}, \boldsymbol{\pi}, c)$, and $\mathcal{S} = \langle \mathcal{A}, \boldsymbol{q} \rangle$ one for $(\mathcal{A}, \boldsymbol{\pi}, c)$. Let $\Delta = W(\mathcal{A} \cup \{\boldsymbol{a}^+\}, \boldsymbol{\pi}, c) - W(\mathcal{A}, \boldsymbol{\pi}, c) > 0$ 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 $\boldsymbol{a}^+$ 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 $\mathcal{S}^0_\varepsilon = \langle \{-, +\}, \boldsymbol{q}^0_\varepsilon \rangle$ which pools all realizations other than $\boldsymbol{a}^+$ by returning $-$ with probability $\boldsymbol{q}^0_\varepsilon(-) = 1 - \varepsilon \boldsymbol{q}^+(\boldsymbol{a}^+)$. If $\mathcal{S}^0_\varepsilon$ evaluates to $+$, the agent implements $\boldsymbol{a}^+$ at the same posterior $\boldsymbol{\pi}^+$ as above. If $\mathcal{S}^0_\varepsilon$ evaluates to $-$, the agent updates her belief to $\boldsymbol{\pi}^-_\varepsilon$ and draws $\mathcal{S}^1_\varepsilon$, which conditions on 'not a^+' by returning $\boldsymbol{a} \in \mathcal{A}$ with probability $\frac{\varepsilon q_i^+(\boldsymbol{a}) + (1-\varepsilon) q_i(\boldsymbol{a})}{1 - \varepsilon q_i^+(\boldsymbol{a})}$ in state i. Since the agent is indifferent across all Blackwell-equivalent sequential information strategies under (C3) and (C5), this strategy too yields welfare $W(\mathcal{A}, \boldsymbol{\pi}, c) + \varepsilon\Delta$. Welfare can only increase if we replace signal $\mathcal{S}^1_\varepsilon$ by the ignorance equivalent $\boldsymbol{\alpha}^\varepsilon$ of the corresponding RI problem $(\mathcal{A}, \boldsymbol{\pi}^\varepsilon_-, c)$, hence

$$(\boldsymbol{\pi} \cdot \boldsymbol{q}^0_\varepsilon(-))(\boldsymbol{\pi}^-_\varepsilon \cdot \boldsymbol{\alpha}^\varepsilon) + (\boldsymbol{\pi} \cdot \boldsymbol{q}^0_\varepsilon(+))(\boldsymbol{\pi}^+ \cdot \boldsymbol{a}^+) - c(\mathcal{S}^0_\varepsilon, \boldsymbol{\pi}) \geq W(\mathcal{A}, \boldsymbol{\pi}, c) + \varepsilon\Delta. \quad (10)$$

[15] See the proofs of Theorem 1 or Corollary 3.

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 $\boldsymbol{\pi}^+$,

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

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

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

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

$$(\boldsymbol{\pi} \cdot \boldsymbol{q}_\varepsilon^0(-))\boldsymbol{\pi}_\varepsilon^- \cdot (\boldsymbol{\alpha} - \boldsymbol{\alpha}^\varepsilon) = \boldsymbol{\pi} \cdot (\boldsymbol{\alpha} - \boldsymbol{\alpha}^\varepsilon) - \varepsilon(\boldsymbol{\pi} \cdot \boldsymbol{q}^+(\boldsymbol{a}^+))\boldsymbol{\pi}^+ \cdot (\boldsymbol{\alpha} - \boldsymbol{\alpha}^\varepsilon) > -\varepsilon\frac{\Delta}{2}. \tag{11}$$

Taken together, Equations (10) and (11) imply that the agent can achieve welfare

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

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

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 $\boldsymbol{a} \in \mathcal{A}$ is optimal in RI problem $(\mathcal{A}, \boldsymbol{\pi}, c)$. By optimality, $W(\{\boldsymbol{a}\}, \boldsymbol{\pi}, c) = W(\mathcal{A}, \boldsymbol{\pi}, c)$, and by menu-monotonicity, the latter is weakly larger than $W(\{\boldsymbol{a}, \boldsymbol{a}'\}, \boldsymbol{\pi}, c)$ for each $\boldsymbol{a}' \in \mathcal{A}$.

Conversely, suppose condition $W(\{\boldsymbol{a}, \boldsymbol{a}'\}, \boldsymbol{\pi}, c) \geq W(\{\boldsymbol{a}\}, \boldsymbol{\pi}, c)$ hold for each $\boldsymbol{a}' \in \mathcal{A}$. We prove by induction that unconditional implementation of $\boldsymbol{a}$ is optimal in RI problem $(A, \boldsymbol{\pi}, c)$, starting with the trivial case $A = \{\boldsymbol{a}\}$ and adding actions one-by-one. For the inductive step, assume $W(\{\boldsymbol{a}\}, \boldsymbol{\pi}, c) = W(A, \boldsymbol{\pi}, c)$ for some subset

$A \subseteq \mathcal{A}$, and consider what happens when $\boldsymbol{a}' \in \mathcal{A} \setminus A$ is added to the menu. Since $\boldsymbol{a} \in A$, the assumption satisfies all the conditions of Definition 1, and $\boldsymbol{a}$ denotes the ignorance equivalent of $(A, \boldsymbol{\pi}, c)$. By Theorem 3(a), the condition $W(\{\boldsymbol{a}, \boldsymbol{a}'\}, \boldsymbol{\pi}, c) \leq W(\{\boldsymbol{a}\}, \boldsymbol{\pi}, c)$ then implies that $W(A \cup \{\boldsymbol{a}'\}, \boldsymbol{\pi}, c) \leq W(A, \boldsymbol{\pi}, c)$. Menu-monotonicity implies the opposite inequality and thus

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

We now turn back to our definition of $\boldsymbol{\pi}$-dominated payoff vectors in Equation (9), and rewrite it in a way that fits Lemma A.8. To do so, we let $\mathcal{S}_{\boldsymbol{\pi}}^{(k,\boldsymbol{\sigma})}$ 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 $\Phi_{\boldsymbol{\pi}} : \Delta\mathcal{I} \to \mathbb{R}$ as

$$\Phi_{\boldsymbol{\pi}}(\boldsymbol{\sigma}) := \boldsymbol{\sigma} \cdot \boldsymbol{\alpha}^{(\mathcal{A},\boldsymbol{\pi},c)} + \lim_{k \downarrow 0} \frac{1}{k} c(\mathcal{S}_{\boldsymbol{\pi}}^{(k,\boldsymbol{\sigma})}, \boldsymbol{\pi}), \tag{12}$$

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

Lemma A.9. *The function $\Phi_{\boldsymbol{\pi}}$ given in Equation (12) is convex and finite-valued, and $\bar{\mathcal{A}}^{\boldsymbol{\pi}} = \left\{\boldsymbol{x} \in \mathbb{R}^I \,\middle|\, \boldsymbol{\sigma} \cdot \boldsymbol{x} \leq \Phi_{\boldsymbol{\pi}}(\boldsymbol{\sigma}) \ \forall \boldsymbol{\sigma} \in \Delta\mathcal{I}\right\}$.*

Proof. By Theorem 3(a), $\bar{\mathcal{A}}^{\boldsymbol{\pi}}$ can be stated with reference to the ignorance equivalent $\boldsymbol{\alpha} = \boldsymbol{\alpha}^{(\mathcal{A},\boldsymbol{\pi},c)}$ alone, $\bar{\mathcal{A}}^{\boldsymbol{\pi}} = \left\{\boldsymbol{x} \in \mathbb{R}^I \mid W(\{\boldsymbol{\alpha}, \boldsymbol{x}\}, \boldsymbol{\pi}, c) \leq W(\{\boldsymbol{\alpha}\}, \boldsymbol{\pi}, c)\right\}$. And since the menu $\{\boldsymbol{\alpha}, \boldsymbol{x}\}$ is binary, any feasible strategy can be described as drawing $\mathcal{S}_{\boldsymbol{\pi}}^{(k,\boldsymbol{\sigma})}$ 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

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

which has the same sign as $\boldsymbol{\sigma} \cdot (\boldsymbol{x} - \boldsymbol{\alpha}) - \frac{1}{k} c(\mathcal{S}_{\boldsymbol{\pi}}^{(k,\boldsymbol{\sigma})}, \boldsymbol{\pi})$. As k converges to zero from above, $k \downarrow 0$, the weighted cost term shrinks. This is because for any $t \in [0,1]$, $\mathcal{S}_{\boldsymbol{\pi}}^{(tk,\boldsymbol{\sigma})}$ can be implemented in two steps by an uninformative coin flip that triggers a draw of $\mathcal{S}_{\boldsymbol{\pi}}^{(k,\boldsymbol{\sigma})}$ with probability t, and otherwise recommends action $\boldsymbol{\alpha}$. By Blackwell monotonicity (C2) and sequential learning-proofness (C3), this implies the monotonic ranking $\frac{1}{tk}c(\mathcal{S}_{\boldsymbol{\pi}}^{(tk,\boldsymbol{\sigma})},\boldsymbol{\pi}) \leq \frac{1}{k}c(\mathcal{S}_{\boldsymbol{\pi}}^{(k,\boldsymbol{\sigma})},\boldsymbol{\pi})$. In other words, the signal is welfare-enhancing for any $k>0$ if and only if

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

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

It remains to show that the function $\Phi_{\boldsymbol{\pi}}:\Delta\mathcal{I}\to\mathbb{R}$ is convex. Consider a two-stage process that first flips an uninformative coin that triggers a draw of $\mathcal{S}_{\boldsymbol{\pi}}^{(k,\boldsymbol{\sigma})}$ with probability t, and otherwise triggers a draw of $\mathcal{S}_{\boldsymbol{\pi}}^{(k,\boldsymbol{\sigma}')}$. Overall, this process recommends action $\boldsymbol{x}$ at posterior $t\boldsymbol{\sigma}+(1-t)\boldsymbol{\sigma}'$ with marginal likelihood k. Again, the one-shot implementation is weakly cheaper by (C2) and (C3), $c(\mathcal{S}^{(k,t\boldsymbol{\sigma}+(1-t)\boldsymbol{\sigma}')},\boldsymbol{\pi}) \leq tc(\mathcal{S}_{\boldsymbol{\pi}}^{(k,\boldsymbol{\sigma})},\boldsymbol{\pi}) + (1-t)c(\mathcal{S}_{\boldsymbol{\pi}}^{(k,\boldsymbol{\sigma}')},\boldsymbol{\pi})$. Convexity follows since

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

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

Lemma A.10. *For any prior $\boldsymbol{\rho}\in\mathsf{int}(\Delta\mathcal{I})$, unconditional implementation of the payoff vector $\boldsymbol{a}^{\boldsymbol{\rho}} = \arg\max\left\{\boldsymbol{\rho}\cdot\boldsymbol{x} \mid x\in\bar{\mathcal{A}}^{\boldsymbol{\pi}}\right\}$ is optimal in any RI problem $(\mathcal{A}',\boldsymbol{\rho},c)$ where the finite menu $\mathcal{A}'\subseteq\mathcal{A}\cup\bar{\mathcal{A}}^{\boldsymbol{\pi}}$ contains $\boldsymbol{a}^{\boldsymbol{\rho}}$.*

Proof. By Lemmas A.8 and A.9, the point $\boldsymbol{a}^{\boldsymbol{\rho}}$ satisfies $\boldsymbol{\rho}\cdot\boldsymbol{a}^{\boldsymbol{\rho}} = \Phi_{\boldsymbol{\pi}}(\boldsymbol{\rho})$. By contradiction, assume that there exists a learning strategy $\mathcal{S}=\langle\mathcal{A}',\boldsymbol{q}\rangle$ that achieves higher welfare than unconditional implementation of $\boldsymbol{a}^{\boldsymbol{\rho}}$ under prior $\boldsymbol{\rho}$,

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

By definition of $\Phi_{\boldsymbol{\pi}}$, this implies that there exists $k>0$ small enough such that

$\sum_{i\in\mathcal{I}} \rho_i \sum_{\boldsymbol{a}\in\mathcal{A}'} q_i(\boldsymbol{a})a_i - c(\mathcal{S},\boldsymbol{\rho}) > \boldsymbol{\rho}\cdot\boldsymbol{\alpha} + \frac{1}{k}c(\mathcal{S}_{\boldsymbol{\pi}}^{(k,\boldsymbol{\rho})},\boldsymbol{\pi})$. Rearranging terms, we obtain

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

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 $\mathcal{S}_{\boldsymbol{\pi}}^{(k,\boldsymbol{\rho})}$ is drawn first, and upon realization $\boldsymbol{x}$, the agent draws and follows signal $\mathcal{S}$, otherwise she implements $\boldsymbol{\alpha}$. This contradicts the binary characterization for optimality of ignorance (Corollary 5) since by definition of $\bar{\mathcal{A}}^{\boldsymbol{\pi}}$, none of the actions $\boldsymbol{a}\in\mathcal{A}'\subseteq\mathcal{A}\cup\bar{\mathcal{A}}^{\boldsymbol{\pi}}$ yields a welfare improvement by itself. □

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: Since anchor $\boldsymbol{a}$ is part of the learning-proof menu, there exists a prior $\boldsymbol{\rho}\in\mathsf{int}(\Delta\mathcal{I})$ such that $\boldsymbol{a}$ is implemented unconditionally in $(\mathcal{A},\boldsymbol{\rho},c)$, and thus denotes the ignorance equivalent. Fix any $k\in(0,1)$ small enough such that the belief $\boldsymbol{\rho}^+ = \frac{1}{1-k}\boldsymbol{\pi} - \frac{k}{1-k}\boldsymbol{\rho}$ has full support. By Lemmas A.8 and A.9, there exist payoff vectors $\boldsymbol{\alpha},\boldsymbol{a}^+\in\bar{\mathcal{A}}^{\boldsymbol{\rho}}$ such that $\boldsymbol{\pi}\cdot\boldsymbol{\alpha} = \Phi_{\boldsymbol{\rho}}(\boldsymbol{\pi})$ and $\boldsymbol{\rho}^+\cdot\boldsymbol{a}^+ = \Phi_{\boldsymbol{\rho}}(\boldsymbol{\rho}^+)$. Moreover, by Lemma A.10, unconditional implementation of $\boldsymbol{\alpha}$ and $\boldsymbol{a}^+$ is optimal in RI problems $(\mathcal{A}\cup\{\boldsymbol{a}^+,\boldsymbol{\alpha}\},\boldsymbol{\pi},c)$ and $(\mathcal{A}\cup\{\boldsymbol{a}^+,\boldsymbol{\alpha}\},\boldsymbol{\rho}^+,c)$ respectively. In particular,

$$W(\mathcal{A}\cup\left\{\boldsymbol{a}^+\right\},\boldsymbol{\pi},c) \le W(\mathcal{A}\cup\left\{\boldsymbol{a}^+,\boldsymbol{\alpha}\right\},\boldsymbol{\pi},c) = \boldsymbol{\pi}\cdot\boldsymbol{\alpha}.$$

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

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 $\boldsymbol{\rho}$:

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

Since sequential implementation yields neither gains nor losses by (C3) and (C5), both strategies have the same expected cost,

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

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

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

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

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

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

References

Roc Armenter, Michèle Müller-Itten, 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.

Alexander W. Bloedel and Weijie Zhong. The cost of optimally acquired information. *Unpublished working paper*, 2020. Available at https://web.stanford.edu/~abloedel/AWB_JMP.pdf.

Andrew Caplin and Mark Dean. Behavioral Implications of Rational Inattention with Shannon Entropy. *NBER Working Paper 19318*, 2013.

Andrew Caplin and Mark Dean. Revealed Preference, Rational Inattention, and Costly Information Acquisition. *American Economic Review*, 105(7):2183–2203, July 2015. doi: 10.1257/aer.20140117. URL http://www.aeaweb.org/articles?id=10.1257/aer.20140117.

Andrew Caplin, Dániel Csaba, John Leahy, and Oded Nov. Rational inattention, competitive supply, and psychometrics. Working Paper 25224, National Bureau of Economic Research, November 2018a. URL http://www.nber.org/papers/w25224.

Andrew Caplin, Mark Dean, and John Leahy. Rational inattention, optimal consideration sets, and stochastic choice. *Review of Economic Studies*, 86(3):1061–1094, 07 2018b. 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*, forthcoming, 2022.

Tommaso Denti, Massimo Marinacci, and Aldo Rustichini. Experimental cost of information. *Unpublished working paper*, 2019.

Alexander Frankel and Emir Kamenica. Quantifying information and uncertainty. *American Economic Review*, 109(10):3650–80, 2019.

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.

Bartosz Maćkowiak, Filip Matějka, and Mirko Wiederholt. Dynamic rational inattention: Analytical results. *Journal of Economic Theory*, 176:650 – 692, 2018a. ISSN 0022-0531. doi: https://doi.org/10.1016/j.jet.2018.05.001. URL http://www.sciencedirect.com/science/article/pii/S002205311830139X.

Bartosz Maćkowiak, Filip Matějka, and Mirko Wiederholt. Survey: Rational inattention, a disciplined behavioral model. Working paper, Center for Economic and Policy Research, October 2018b.

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.

Ludmila Matyskova. Bayesian persuasion with costly information acquisition. *CERGE-EI Working Paper Series*, 614, 2018.

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.

Roger B Myerson. Mechanism design by an informed principal. *Econometrica: Journal of the Econometric Society*, pages 1767–1797, 1983.

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.

George J. Stigler. The economics of information. *Journal of Political Economy*, 69(3): 213–225, 1961. ISSN 00223808, 1537534X. URL http://www.jstor.org/stable/1829263.

Elias Tsakas. Robust scoring rules. *Theoretical Economics*, 15(3):955–987, 2020. URL https://econtheory.org/ojs/index.php/te/article/view/20200955/0.

Günter M. Ziegler. *Lectures on polytopes*, volume 152. Springer Science & Business Media, 2012.

B Online Appendix

Outsourcing Games

Proof of Theorem 5.A:. Let $\boldsymbol{\alpha}$ denote the ignorance equivalent of RI problem $(\mathcal{A}, \boldsymbol{\pi}, c)$. If a contract with these terms is offered and accepted unconditionally, the resulting equilibrium payoff for manager Bertrand is equal to $W(\widetilde{\mathcal{A}}, \boldsymbol{\pi}, \tilde{c}) - \boldsymbol{\pi} \cdot \boldsymbol{\alpha}$, and that for agent Abigail is equal to $\boldsymbol{\pi} \cdot \boldsymbol{\alpha} = W(\mathcal{A}, \boldsymbol{\pi}, c)$. We consider the following candidate equilibrium, and then verify that these strategies and beliefs form a Perfect Bayesian Equilibrium (PBE).

- At time zero, Bertrand draws a free uninformative signal and unconditionally offers the singleton menu $\{\boldsymbol{\alpha}\}$.
- At time one on the equilibrium path, Abigail maintains the prior belief $\boldsymbol{\pi}$ and selects the contract $\boldsymbol{\alpha}$ unconditionally. For off-equilibrium menu offers $\mathcal{B} \neq \{\boldsymbol{\alpha}\}$, we distinguish two cases: If there exists a prior $\boldsymbol{\rho} \in \mathsf{int}(\Delta\mathcal{I})$ such that $W(\mathcal{A} \cup \{\boldsymbol{\alpha}\} \cup \mathcal{B}, \boldsymbol{\rho}, c) = W(\mathcal{A}, \boldsymbol{\rho}, c)$, Abigail updates her belief[16] to one such $\boldsymbol{\rho}$ and proceeds with an optimal strategy for RI problem $(\mathcal{A}, \boldsymbol{\rho}, c)$. In doing so, she effectively ends the game. If no such belief exists, then it is in particular never optimal to implement $\boldsymbol{\alpha}$ unconditionally, and by Corollary 4(b) and (d), there exists some belief $\boldsymbol{\sigma}$ such that the ignorance equivalent $\boldsymbol{\alpha}^{(\mathcal{A} \cup \{\boldsymbol{\alpha}\} \cup \mathcal{B}, \boldsymbol{\sigma}, c)}$ dominates $\boldsymbol{\alpha}$ statewise. Abigail then updates her belief to $\boldsymbol{\sigma}$ and proceeds with an optimal strategy for RI problem $(\mathcal{B}, \boldsymbol{\sigma}, c)$.
- At time two, regardless if on or off the equilibrium path, Bertrand maintains the prior belief $\boldsymbol{\pi}$ and follows an optimal learning strategy for RI problem $(\widetilde{\mathcal{A}}, \boldsymbol{\pi}, \tilde{c})$. Once the state realizes, he executes the required transfer to Abigail.

Bayes-Plausibility. On the equilibrium path, no information is obtained until the last round and each agent selects a degenerate strategy. By endowing each agent with the belief $\boldsymbol{\pi}$, we trivially ensure that Bayes rule holds.

Sequential Rationality. We treat each time period in a separate paragraph.

[16] We directly assign a likelihood to each state i, but note that any such interior belief $\boldsymbol{\rho}$ can be generated within the game. One way is to assume that Bertrand draws and follows a binary signal $\langle \{\mathcal{B}, \{\boldsymbol{\alpha}\}\}, \boldsymbol{q}^0 \rangle$ with $q_i^0(\mathcal{B}) = \varepsilon \rho_i / \pi_i$ for $\varepsilon \in (0,1)$ small enough to ensure all probabilities are interior. Bayes rule then dictates that the resulting likelihood for state i is equal to $\rho_i = \frac{\pi_i q_i^0(\mathcal{B})}{\boldsymbol{\pi} \cdot \boldsymbol{q}^0(\mathcal{B})}$.

At time two, Bertrand incurs the additive transfer irrespective of his choice. Since $W(\widetilde{\mathcal{A}} - \{\boldsymbol{b}\}, \boldsymbol{\pi}, \tilde{c}) = W(\widetilde{\mathcal{A}}, \boldsymbol{\pi}, \tilde{c}) - \boldsymbol{\pi} \cdot \boldsymbol{b}$, the suggested strategy is optimal.

At time one, Abigail's choice directly determines her final payoff and in that sense, she too is facing a standard RI problem. If her belief $\boldsymbol{\rho}$ is such that $W(\mathcal{A} \cup \{\boldsymbol{\alpha}\} \cup \mathcal{B}, \boldsymbol{\rho}, c) = W(\mathcal{A}, \boldsymbol{\rho}, c)$ then any solution to the latter RI problem is clearly optimal. If no such belief exists, then not only is unconditional implementation of $\boldsymbol{\alpha}$ never optimal, the same can be said of any $\boldsymbol{a} \in \mathcal{A}$. By Corollary 4(b) and (c), none of these actions is chosen in any optimal solution to RI problem $(\mathcal{A} \cup \{\boldsymbol{\alpha}\} \cup \mathcal{B}, \boldsymbol{\sigma}, c)$. In turn, this implies that Abigail is just as well off if she restricts attention to submenu $\mathcal{B}$ and proceeds with an optimal RI strategy for $(\mathcal{B}, \boldsymbol{\sigma}, c)$ as the candidate equilibrium suggests. In passing, note that this also implies that $\boldsymbol{\alpha}^{(\mathcal{A} \cup \{\boldsymbol{\alpha}\} \cup \mathcal{B}, \boldsymbol{\sigma}, c)} = \boldsymbol{\alpha}^{(\mathcal{B}, \boldsymbol{\sigma}, c)}$.

At time zero, Bertrand guarantees himself an expected payoff of

$$U_{\mathsf{B}} = W(\widetilde{\mathcal{A}}, \boldsymbol{\pi}, \tilde{c}) - \boldsymbol{\pi} \cdot \boldsymbol{\alpha}$$

by following the suggested strategy. If he instead deviates by drawing a signal $\mathcal{S}^0 = \langle S, \boldsymbol{q}^0 \rangle$ and offering $\mathcal{B}^s$ conditional on observing $s \in S$, one of two things happens: Either Abigail ends the game or $W(\mathcal{A} \cup \{\boldsymbol{\alpha}\} \cup \mathcal{B}^s, \boldsymbol{\rho}, c)$ strictly exceeds $W(\mathcal{A}, \boldsymbol{\rho}, c)$ for all priors $\boldsymbol{\rho}$. We will argue that no matter what, Bertrand is just as well off by offering $\{\boldsymbol{\alpha}\}$ instead of $\mathcal{B}^s$. In the former case, offering $\mathcal{B}^s$ yields a payoff of zero, whereas offering $\{\boldsymbol{\alpha}\}$ ensures that Bertrand obtains the mandate unconditionally and then achieves a net payoff of $W(\widetilde{\mathcal{A}}, \boldsymbol{\pi}^s, \tilde{c}) - \boldsymbol{\pi}^s \cdot \boldsymbol{\alpha}$. By self-selection (Corollary 3), unconditional implementation of $\boldsymbol{\alpha}$ yields at most $W(\mathcal{A}, \boldsymbol{\pi}^s, c)$, imposing the lower bound

$$W(\widetilde{\mathcal{A}}, \boldsymbol{\pi}^s, \tilde{c}) - \boldsymbol{\pi}^s \cdot \boldsymbol{\alpha} \geq W(\widetilde{\mathcal{A}}, \boldsymbol{\pi}^s, \tilde{c}) - W(\mathcal{A}, \boldsymbol{\pi}^s, c)$$

which in turn is nonnegative by menu- and cost-monotonicity of W. In the latter case, Abigail follows a learning strategy $\mathcal{S}^s = \langle \mathcal{B}^s, \boldsymbol{q}^1 \rangle$ and hires Bertrand at different terms depending on the signal realization. While $\boldsymbol{\alpha} < \boldsymbol{\alpha}^{(\mathcal{B}^s, \boldsymbol{\sigma}, c)}$ trivially implies that Bertrand could secure the mandate more cheaply with menu $\{\boldsymbol{\alpha}\}$ rather than $\mathcal{B}^s$, the latter includes the draw of $\mathcal{S}^s$ at no additional cost to Bertrand. The main difficulty here is that Abigail evaluates the cost of $\mathcal{S}^s$ at a different prior than Bertrand, $\boldsymbol{\mathcal{S}} \neq \boldsymbol{\pi}^s$, which complicates a direct comparison with Bertrand's own 'in-house' cost of learning. However, Corollary 7 implies that, as long as $\boldsymbol{\pi}^s$ has full support, there exist contract terms $\boldsymbol{x}^s = \max_{\boldsymbol{x} \in \bar{\mathcal{A}}^{\mathcal{S}^s}} \boldsymbol{\pi}^s \cdot \boldsymbol{x}$ such that the expected transfer under terms $\boldsymbol{x}^s$ equals the expected transfer under $\mathcal{S}^s$ net of learning costs,

$$\boldsymbol{\pi}^s \cdot \boldsymbol{x}^s = \boldsymbol{\pi}^s \cdot \boldsymbol{a}^{\mathcal{S}^s} - c(\mathcal{S}^s, \boldsymbol{\pi}^s),$$

where $\bar{a}_i^{\mathcal{S}^s} = \sum_{\boldsymbol{b}\in\mathcal{B}} q_i^1(\boldsymbol{b}) b_i$ denotes the average transfer in state i under $\mathcal{S}^s$. The set $\bar{\mathcal{A}}^{\mathcal{S}^s}$ also includes $\boldsymbol{\alpha}^{(\mathcal{B}^s,\boldsymbol{\sigma},c)}$ by optimality of $\mathcal{S}^s$ (Theorem 1). And since it is only bounded above, statewise dominance $\boldsymbol{\alpha} < \boldsymbol{\alpha}^{(\mathcal{B}^s,\boldsymbol{\sigma},c)}$ also implies that $\boldsymbol{\alpha} \in \bar{\mathcal{A}}^{\mathcal{S}^s}$. In particular, $\boldsymbol{\pi}^s \cdot \boldsymbol{x} \leq \boldsymbol{\pi}^s \cdot \boldsymbol{x}$. Together, this allows us to conclude that

$$\boldsymbol{\pi}^s \cdot \boldsymbol{\alpha} + \tilde{c}(\mathcal{S}^s, \boldsymbol{\pi}^s) \leq \boldsymbol{\pi}^s \cdot \boldsymbol{x} + \tilde{c}(\mathcal{S}^s, \boldsymbol{\pi}^s) = \boldsymbol{\pi}^s \cdot \boldsymbol{a}^{\mathcal{S}^s} - c(\mathcal{S}^s, \boldsymbol{\pi}^s) + \tilde{c}(\mathcal{S}^s, \boldsymbol{\pi}^s) \leq \boldsymbol{\pi}^s \cdot \boldsymbol{a}^{\mathcal{S}^s},$$

and so Bertrand would be better off offering only terms $\boldsymbol{\alpha}$ and acquiring signal $\mathcal{S}^s$ by himself rather than incurring expected transfer $\boldsymbol{\pi}^s \cdot \boldsymbol{a}^{S^1}$ and obtaining $\mathcal{S}^s$ for free. By continuity, the same holds even if $\boldsymbol{\pi}^s$ does not have full support. Taken together, we have shown that Bertrand is weakly better off by drawing $\mathcal{S}^0$, unconditionally offering contract menu $\{\boldsymbol{\alpha}\}$ and then potentially drawing additional signals $\mathcal{S}^s$ before solving the time-two RI problem $(\widetilde{\mathcal{A}}, \boldsymbol{\pi}^{(s,\boldsymbol{b})}, \tilde{c})$ under some updated belief $\boldsymbol{\pi}^{(s,\boldsymbol{b})}$. Since it is weakly cheaper to acquire all information in one shot by (C3), and transfers are now always equal to $\boldsymbol{\alpha}$, the net payoff for Bertrand is thus weakly smaller than under the candidate equilibrium strategy.

Together, Bayes-plausibility and sequential rationality establish the suggested strategies as a PBE. To show that no other PBE achieves a higher expected payoff for Bertrand, we refer the reader to the more general version in the next proof. $\square$

Proof of Theorem 5.B:. Let $\bar{\mathcal{A}}$ denote the ignorance equivalent of RI problem $(\mathcal{A}, \boldsymbol{\pi}, c)$. Let further $\mathcal{F}^t = \langle F^t, \boldsymbol{r}^t \rangle$ denote the free public signal at the beginning of each round $t = 0, 1, 2$. We denote the history dependence of the public beliefs and the optimal strategies as $\boldsymbol{\pi}^{\tilde{h},h}$ and $\mathcal{S}^{\tilde{h},h}$ respectively, where history $\tilde{h}$ refers to the strategic choices observed in earlier rounds, and h to the public signals drawn at the beginning of each round. Since each draw from $\mathcal{F}^0$ essentially defines a different game, we simplify notation by fixing an arbitrary draw f_0 and dropping the dependence superscript — so $\boldsymbol{\pi}$ henceforth refers to the updated belief $\boldsymbol{\pi}^{f_0}$.

Our candidate equilibrium is as follows:

- At time zero, Bertrand draws a free uninformative signal and unconditionally offers the contract menu $\bar{\mathcal{A}}$.

- At time one on the equilibrium path, Abigail updates her belief to $\boldsymbol{\pi}^{f_1}$ and unconditionally accepts the contract with terms equal to the ignorance equivalent of RI problem $(\mathcal{A}, \boldsymbol{\pi}^{(\bar{A}, f_1)}, c)$. For off-equilibrium menu offers $\mathcal{B} \neq \mathcal{A}$, Abigail's beliefs and actions are as in the equilibrium without shocks.

- At time two, regardless if on or off equilibrium path, Bertrand's belief equals $\boldsymbol{\pi}^{(f_1,f_2)}$ and he draws and follows an optimal learning strategy for RI problem $(\widetilde{\mathcal{A}}, \boldsymbol{\pi}^{(f_1,f_2)}, \tilde{c})$. Once the state realizes, he executes the required transfer to Abigail.

We now verify that these strategies and beliefs form a Perfect Bayesian Equilibrium.

Bayes-Plausibility. On the equilibrium path, only public information is revealed until the last round and each agent selects a degenerate strategy at each information set. By endowing each agent with the public belief at each point in time, we trivially ensure that Bayes rule holds.

Sequential Rationality. We treat each time period in a separate paragraph.

At time two, Bertrand incurs the additive transfer irrespective of his choice. Since $W(\widetilde{\mathcal{A}} - \{\boldsymbol{b}\}, \boldsymbol{\pi}^{(f_1,f_2)}, \tilde{c}) = W(\widetilde{\mathcal{A}}, \boldsymbol{\pi}^{(f_1,f_2)}, \tilde{c}) - \boldsymbol{\pi}^{(f_1,f_2)} \cdot \boldsymbol{b}$, the suggested strategy is optimal.

At time one, Abigail's choice directly determines her final payoff and in that sense, she too is facing a standard RI problem. On the equilibrium path, picking the appropriate ignorance equivalent is optimal by Theorem 2(b). Off the equilibrium path, the same arguments as in the previous proof establish optimality of her strategy.

At time zero, Bertrand can guarantee himself an expected payoff of

$$U_{\mathsf{B}} = \sum_{i \in \mathcal{I}} \pi_i \sum_{f_1 \in F_1} r_i^1(f_1) \left[\sum_{f_2 \in F_2} r_i^2(f_2) W(\widetilde{\mathcal{A}}, \boldsymbol{\pi}^{(f_1,f_2)}, \tilde{c}) - W(\mathcal{A}, \boldsymbol{\pi}^{f_1}, c) \right]$$

by following the suggested strategy. If he instead deviates by drawing a signal $\mathcal{S}^0 = \langle S, \boldsymbol{q}^0 \rangle$ and offering $\mathcal{B}^s$ conditional on observing $s \in S$, one of two things happens: Either Abigail ends the game or $W(\mathcal{A} \cup \{\boldsymbol{\alpha}\} \cup \mathcal{B}^s, \boldsymbol{\rho}, c)$ strictly exceeds $W(\mathcal{A}, \boldsymbol{\rho}, c)$ for all priors $\boldsymbol{\rho}$. We will argue that no matter what, Bertrand is just as well off by offering $\bar{\mathcal{A}}$ instead of $\mathcal{B}^s$. In the former case, offering $\mathcal{B}^s$ yields a payoff of zero, whereas offering $\bar{\mathcal{A}}$ ensures that for any f_1 draw, Bertrand obtains the mandate at terms $\boldsymbol{\alpha}^{(\mathcal{A},\boldsymbol{\pi}^{f_1},c)}$ and achieves a conditional payoff of

$$\sum_{f_2\in F_2} r_i^2(f_2)W(\widetilde{\mathcal{A}},\boldsymbol{\pi}^{(f_1,f_2)},\tilde{c}) - \boldsymbol{\pi}^{(s,f_1)}\cdot\boldsymbol{\alpha}^{(\mathcal{A},\boldsymbol{\pi}^{f_1},c)}.$$

By convexity of the welfare function (Lemma A.3), the former term is weakly larger than $W(\widetilde{\mathcal{A}},\boldsymbol{\pi}^{(f_1,f_2)},\tilde{c})$, and by self-selection (Corollary 3), the latter term is weakly smaller than $W(\mathcal{A},\boldsymbol{\pi}^{(s,f_1)},c)$. The difference is thus nonnegative by menu- and cost-monotonicity of W. In the latter case, Abigail follows a learning strategy $\mathcal{S}^s=\langle\mathcal{B}^s,\boldsymbol{q}^1\rangle$ and hires Bertrand at different terms depending on the signal realization. By the same argument as in the previous proof, Bertrand is weakly better off by contracting at terms $\boldsymbol{\alpha}^{(\mathcal{A},\boldsymbol{\pi}^{f_1},c)}$ unconditionally and acquiring $\mathcal{S}^s$ by himself. Offering $\bar{\mathcal{A}}$ achieves exactly that. Taken together, we have thus shown that Bertrand is weakly better off by drawing $\mathcal{S}^0$, unconditionally offering contract menu $\mathcal{A}$ and then potentially drawing additional signals $\mathcal{S}^s$ before solving the time-two RI problem $(\widetilde{\mathcal{A}},\boldsymbol{\pi}^{(s,\boldsymbol{b},f_1,f_2)},\tilde{c})$ under some updated belief $\boldsymbol{\pi}^{(s,\boldsymbol{b},f_1,f_2)}$. Since it is weakly cheaper to acquire all information in one shot at time two by (C2) and (C3) and $\bar{\mathcal{A}}$ is offered unconditionally, the net payoff for Bertrand is thus weakly smaller than under the candidate equilibrium strategy.

We now show that no other PBE achieves a higher expected payoff for Bertrand. For this, we only have to consider on-path strategies and payoffs. We consider an arbitrary time-zero strategy $(\mathcal{S}^0,\{\mathcal{B}_s\mid s\in S_0\})$ for Bertrand. At time one, Abigail observes the public draw f_1 and a contract offer $\mathcal{B}\in\bar{B}$ and updates her belief to $\boldsymbol{\pi}^{(\mathcal{B},f_1)}$ according to Bayes rule. Her optimal equilibrium strategy $\mathcal{S}^{(\mathcal{B},f_1)}$ is determined by RI-problem $(\mathcal{A}\cup\mathcal{B},\boldsymbol{\pi}^{(\mathcal{B},f_1)},c)$, and by menu-monotonicity, Abigail expects a payoff of at least $W(\mathcal{A},\boldsymbol{\pi}^{(\mathcal{B},f_1)},c)$,

$$\sum_{i\in\mathcal{I}}\pi_i^{(\mathcal{B},f_1)}\left[\sum_{\boldsymbol{a}\in\mathcal{A}\cup\mathcal{B}} q_i^{(\mathcal{B},f_1)}(\boldsymbol{a})a_i\right] - c(\mathcal{S}^{(\mathcal{B},f_1)},\boldsymbol{\pi}^{(\mathcal{B},f_1)}) \geq W(\mathcal{A},\boldsymbol{\pi}^{(\mathcal{B},f_1)},c) \tag{13}$$

for all $\mathcal{B}\in\bar{B}$ and $f_1\in F_1$. If the game proceeds to time two, Bertrand's optimal strategy is determined by RI problem $(\widetilde{\mathcal{A}},\boldsymbol{\pi}^{(s,\boldsymbol{b},f_1,f_2)},\tilde{c})$ for each $s\in S_0$ and $\boldsymbol{b}\in\mathcal{B}_s$ because the additive transfer $-\boldsymbol{b}$ is incurred irrespective of his choice.

Overall, Bertrand obtains an equilibrium payoff of

$$U'_{\mathsf{B}} = \sum_{(i,s,f_1,f_2)} \pi_i q_i^0(s) r_i^1(f_1) r_i^2(f_2) \sum_{\boldsymbol{b}\in\mathcal{B}_s} q_i^{(\mathcal{B}_s,f_1)}(\boldsymbol{b}) \Big[W(\widetilde{\mathcal{A}}, \boldsymbol{\pi}^{(s,\boldsymbol{b},f_1,f_2)}, \tilde{c}) - b_i \Big] - \tilde{c}(\mathcal{S}^0, \boldsymbol{\pi}), \quad (14)$$

where the first sum is taken over the Cartesian product $\mathcal{I} \times S_0 \times F_1 \times F_2$. We compare this payoff to a situation where Bertrand personally draws all the signals at time two. More specifically, he first obtains the mandate by offering $\mathcal{A}$ unconditionally. He then first draws $\mathcal{S}^0$ and observes its realization s, for a cost of $\tilde{c}(\mathcal{S}^0, \boldsymbol{\pi}^{(f_1,f_2)})$. Next, Bertrand draws $\mathcal{S}^{(\mathcal{B}_s,f_1)}$ and observes its realization $\boldsymbol{a} \in \mathcal{A} \cup \mathcal{B}_s$, for a cost of $\tilde{c}(\mathcal{S}^{(\mathcal{B}_s,f_1)}, \boldsymbol{\pi}^{(s,f_1,f_2)})$. If $\boldsymbol{a} \in \mathcal{A}$, Bertrand follows its recommendation, and otherwise solves the RI problem $(\widetilde{\mathcal{A}}, \boldsymbol{\pi}^{(s,\boldsymbol{a},f_1,f_2)}, \tilde{c})$. He then pays Abigail the required transfer $\boldsymbol{\alpha}^{(\mathcal{A},\boldsymbol{\pi}^{f_1},c)}$. A one-shot implementation of this information strategy costs weakly less by (C3) and is admissible in RI problem $(\widetilde{\mathcal{A}}, \boldsymbol{\pi}^{(f_1,f_2)}, \tilde{c})$, yielding the lower bound

$$\begin{aligned} U_{\mathsf{B}} \geq \sum_{(i,s,f_1,f_2)} \pi_i q_i^0(s) r_i^1(f_1) r_i^2(f_2) \Bigg[& \sum_{\boldsymbol{a}\in\mathcal{A}} q_i^{(\mathcal{B}_s,f_1)}(\boldsymbol{a}) a_i + \sum_{\boldsymbol{a}\in\mathcal{B}_s} q_i^{(\mathcal{B}_s,f_1)}(\boldsymbol{a}) W(\widetilde{\mathcal{A}}, \boldsymbol{\pi}^{(s,\boldsymbol{a},f_1,f_2)}, \tilde{c}) \\ & - \tilde{c}(\mathcal{S}^0, \boldsymbol{\pi}^{(f_1,f_2)}) - \tilde{c}(\mathcal{S}^{(\mathcal{B}_s,f_1)}, \boldsymbol{\pi}^{(s,f_1,f_2)}) - \alpha_i^{(\mathcal{A},\boldsymbol{\pi}^{f_1},c)} \Bigg]. \end{aligned}$$

In particular, the difference in payoff across the two equilibrium payoffs is at least

$$\begin{aligned} U_{\mathsf{B}} - U'_{\mathsf{B}} \geq \sum_{(i,s,f_1,f_2)} \pi_i q_i^0(s) r_i^1(f_1) r_i^2(f_2) \Bigg[& \sum_{\boldsymbol{a}\in\mathcal{A}} q_i^{(\mathcal{B}_s,f_1)}(\boldsymbol{a}) a_i + \sum_{\boldsymbol{b}\in\mathcal{B}_s} q_i^{(\mathcal{B}_s,f_1)}(\boldsymbol{b}) b_i - \alpha_i^{(\mathcal{A},\boldsymbol{\pi}^{f_1},c)} \\ & + \tilde{c}(\mathcal{S}^0, \boldsymbol{\pi}) - \tilde{c}(\mathcal{S}^0, \boldsymbol{\pi}^{(f_1,f_2)}) - \tilde{c}(\mathcal{S}^{(\mathcal{B}_s,f_1)}, \boldsymbol{\pi}^{(s,f_1,f_2)}) \Bigg]. \end{aligned}$$

We can get rid of the dependence on s by letting $Q_i^0(\mathcal{B}) = \sum_{s\in S_0:\mathcal{B}_s=\mathcal{B}} q_i^0(s)$ denote the likelihood of menu offer $\mathcal{B}$ in state i. By prior-concavity of the cost function (C2), acquiring the same signal under better information is less costly on average. As a consequence, the lower bound gets weakly smaller when we replace $\tilde{c}(\mathcal{S}^0, \boldsymbol{\pi}^{(f_1,f_2)})$ with $\tilde{c}(\mathcal{S}^0, \boldsymbol{\pi})$, and $\tilde{c}(\mathcal{S}^{(\mathcal{B}_s,f_1)}, \boldsymbol{\pi}^{(s,f_1,f_2)})$ with $\tilde{c}(\mathcal{S}^{(\mathcal{B}_s,f_1)}, \boldsymbol{\pi}^{(\mathcal{B}_s,f_1)})$. And since $\tilde{c} \leq c$, replacing any of the negative cost terms with c only lowers the bound further and drops the dependence on f_2,

$$U_{\mathsf{B}} - U'_{\mathsf{B}} \geq \sum_{(i,\mathcal{B},f_1)} \pi_i Q_i^0(\mathcal{B}) r_i^1(f_1) \Bigg[\sum_{\boldsymbol{a}\in\mathcal{A}\cup\mathcal{B}} q_i^{(\mathcal{B},f_1)}(\boldsymbol{a}) a_i - \alpha_i^{(\mathcal{A},\boldsymbol{\pi}^{f_1},c)} - c(\mathcal{S}^{(\mathcal{B},f_1)}, \boldsymbol{\pi}^{(\mathcal{B},f_1)}) \Bigg]. \quad (15)$$

Optimality of Abigail's strategy $\mathcal{S}^{\mathcal{B}}$ implies that this bound can further be relaxed by Equation (13),

$$U_{\mathsf{B}} - U'_{\mathsf{B}} \geq \sum_{(i,\mathcal{B},f_1)} \pi_i Q_i^0(\mathcal{B}) r_i^1(f_1) \Bigg[W(\mathcal{A}, \boldsymbol{\pi}^{(\mathcal{B},f_1)}, c) - \alpha_i^{(\mathcal{A},\boldsymbol{\pi}^{f_1},c)} \Bigg].$$

By convexity of W, this in turn gets weakly smaller if we drop the reference to $\mathcal{B}$,

$$\begin{aligned} U_{\mathsf{B}} - U'_{\mathsf{B}} &\geq \sum_{(i,f_1)} \pi_i r_i^1(f_1) \Bigg[W(\mathcal{A}, \boldsymbol{\pi}^{f_1}, c) - \alpha_i^{(\mathcal{A},\boldsymbol{\pi}^{f_1},c)} \Bigg] \\ &= (\boldsymbol{\pi} \cdot \boldsymbol{r}^1(f_1)) \Bigg[W(\mathcal{A}, \boldsymbol{\pi}^{f_1}, c) - \boldsymbol{\pi}^{f_1} \cdot \boldsymbol{\alpha}^{(\mathcal{A},\boldsymbol{\pi}^{f_1},c)} \Bigg] = 0. \end{aligned}$$

In other words, there exists no Perfect Bayesian Equilibrium where Bertrand achieves a higher payoff than in the one described at the beginning of this proof. $\square$