Rational Inattention via Ignorance Equivalence

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

January 17, 2020

Abstract

We present a novel approach to finite Rational Inattention (RI) models based on the *ignorance equivalent*, a fictitious action with state-dependent payoffs that effectively summarizes the optimal learning and conditional choices. The ignorance equivalent allows us to recast the RI problem as a standard expected utility maximization over an augmented choice set called the *learning-proof menu*, yielding new insights regarding the behavioral implications of RI, in particular as new actions are added to the menu. Our geometric approach is also well suited to numerical methods, outperforming existing techniques both in terms of speed and accuracy, and offering robust predictions on the most frequently implemented actions.

^{*}Federal Reserve Bank of Philadelphia: roc.armenter@phil.frb.org

[†]University of Notre Dame: michele.muller@nd.edu

[‡]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 Isaac Baley, John Leahy, 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.

"It's much easier not to know things sometimes."

– Stephen Chbosky, The Perks of Being a Wallflower

1 Introduction

Economists have long sought to model decisions in the face of risk and incomplete information. In such environments, agents typically seek and acquire information, and thus effectively shape the uncertainty that they face — but there are few models describing how they do so.

Introduced by [Sims, 2003], Rational Inattention (RI) has gained increasing acceptance as a model of information acquisition and processing, particularly in macroeconomics and finance. An agent, facing a menu of actions with uncertain payoffs, can condition her actions upon any arbitrary signal of the state of nature, but more informative signals are more costly. In line with the literature [Caplin et al., 2018b, Matějka and McKay, 2015, Sims, 2003], we assume that signal costs are proportional to the average reduction in Shannon entropy. The result is an endogenous information structure that responds to incentives and changes in the environment and has been documented to reproduce empirical regularities in a variety of contexts, from portfolio design to price setting.

However, by and large we still do not know how effective RI models are at explaining real-world choice data. As noted by Gabaix [2014], the limited scope of existing applied work is partly due to the conceptual and computational complexity of RI optimization problems. Outside a handful of special cases, RI models do not admit a closed form solution. Existing numerical methods are often computationally intensive and may suffer from accuracy problems. The sheer size of the information structure — the joint probability distribution of actions and states — often proves to be an impediment to analyzing and understanding the key findings and relevant comparative statics.

We present a novel approach to finite RI models that yields important insights and is highly conducive to numerical solution methods. A key concept throughout the paper is that of the *ignorance equivalent* (IE), which we define to be a fictitious action that, whether added to the original choice menu or implemented unconditionally, makes the agent no better or worse off. We show that the IE always exists and is unique. The IE is conceptually similar to the certainty equivalent for choice under risk, effectively summarizing the most pertinent features of the choice problem. By drawing on the concept of the IE, one can transform any finite RI problem into a standard expected utility maximization. We construct what we term the "learningproof menu" from the collection of IEs across all priors. The key observation is that adding an agent's IE to the menu never generates new learning opportunities for herself or any other agent with a possibly different prior. Any agent is thus indifferent between the original and the learning-proof menu — and since it is always optimal in the latter to choose the agent's own IE and forgo learning, the RI problem is effectively equivalent to a standard expected utility maximization over the learning-proof menu.

The IE and the learning-proof menu also bring structure to the comparative statics of menu expansion. A single new action is part of the optimal learning scheme if and only if it is also attractive in combination with the IE alone.¹ However, because actions can be complements when the agent can learn, it is possible that a larger menu increases the appeal of existing actions [Matějka and McKay, 2015]. As the menu expands, actions may thus move in and out of the consideration set, i.e. the support of the optimal choice. We show that actions that are interior to the learningproof menu will never be implemented under any menu expansion. Conversely, those outside the learning-proof menu will be implemented with positive probability when the right new action is added to the menu.

At the heart of our results is the equivalence between a finite RI problem and a much simpler convex optimization problem via an explicit transformation of the payoffs. Intuitively, the transformation accentuates payoff differences across actions when information is cheap, and attenuates them when information is costly, embodying how the agent spends his limited attention resources. The new formulation, which we call the Geometric Attention Problem (GAP), operates on a reduced dimensionality, is scale invariant, and separates the role of the prior beliefs over states from the set of feasible choices. The GAP also provides a tight bound on the number of actions played with strictly positive probability.

We lay out several practical implications from our theoretical results, with an eye on numerical methods. In particular, we show how to use noisy estimates of the optimum to construct *partial covers*, which are subsets of the menu that include the consideration set with any specified probability, including 1. The method can be readily deployed to assess computational accuracy in an economically meaningful

¹While we express it in terms of the IE, this result is mathematically equivalent to the "market entry test" of Caplin et al. [2018b].

way, and to identify the robust behavioral features of the problem. We also present a notion of distance between choices based on the IE. We argue that the IE distance is both parsimonious and lends itself as a stopping criterion for numerical estimation, overcoming some of the shortcomings of alternative metrics.

Overall, the GAP is very well suited for numerical methods. There are plenty of known algorithms for convex problems that can be deployed for substantial gains in accuracy and computation time. We provide one such algorithm, adapted from standard sequential quadratic programming with active set methods.² Although our method is relatively unsophisticated, it performs favorably compared with other approaches, both in terms of speed and accuracy.

Finally, we provide three applications intended to illustrate some of the advantages of our approach. The first application is based on the price-setting problem of Matějka [2016]. We document that our algorithm is orders of magnitude faster than the Blahut-Arimoto algorithm suggested by Caplin et al. [2018b], and scales well with the size of the action and state spaces. Precision, rather than speed, is the focus of our second application, based on the two-dimensional portfolio design problem of Jung et al. [2019]. Our algorithm unveils some behavioral differences of importance. In both applications, we obtain robust predictions by deploying our characterization of partial covers and dominated actions, which both prove to be very tight estimates of the consideration set. Our third application is a novel task assignment problem, a complex but naturally finite RI problem that is the ideal scenario for the GAP approach.

Related Literature. Rational inattention was first introduced into economics by Sims [2003], deploying the ideas of information theory.³ Rational inattention models rapidly found their way in a variety of fields, from finance to monetary economics.⁴

Early work on RI models restricted their analysis to Linear-Quadratic Gaussian (LQG) frameworks, or assumed that the solution was Gaussian as an approximation, to obtain analytic results that can, in turn, be embedded in an equilibrium model, and have led to many insights for aggregate phenomena.⁵

²Source codes are available on GitHub at https://github.com/mmulleri/GAP-SQP.

³Of course, information theory traces back to the groundbreaking work of Claude Shannon and information economics does to [Stigler, 1961].

⁴We cannot hope to properly review what is by now a large literature. See Mackowiak et al. [2018] for a survey of both theoretical and applied work with rational inattention models.

⁵A necessarily incomplete list of examples is: Peng [2005], Peng and Xiong [2006] and Huang and

Sims [2006] exhorted researchers to go beyond the LQG case, providing a closedform solution for a particular non-LQG case as well. Our applications draw on two leading examples, Matějka [2016] and Jung et al. [2019], who analyze the full joint distribution of actions and states in more general settings. Gaglianone et al. [2019] estimate a dynamic RI model to understand the role of incentives on the accuracy of financial forecasts. Other researchers have instead sought to develop further the LQG framework to circumvent some of its shortcomings. Luo et al. [2017] apply Gaussian techniques with CARA preferences, allowing them to study the dynamics of consumption and wealth in general equilibrium. Mondria [2010] allows signals to be linear combinations of the underlying state of the economy, an approach that is also followed in Kacperczyk et al. [2016], among others. Miao et al. [2019] make further progress in multi-variate LQG environments.

Alongside applied work, there have been significant theoretical developments exploring and extending RI problems. Caplin and Dean [2013] and Caplin et al. [2018a] solve the general finite model using a 'posterior-based' approach and mirror the concavification procedure of Gentzkow and Kamenica [2014]; 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. [2018] explore dynamic learning in an RI context.

The recent paper by Caplin et al. [2018b] is closely related to our theory results. Caplin et al. [2018b] provide a set of necessary and sufficient first-order conditions, as well as a 'market-entry' test, which is a simple sufficient statistic that determines whether a new, currently unavailable action would be incorporated into the consideration set. We re-encounter both results in our formulation of the GAP, which leads to additional theoretical insights beyond their work, such the concept of the IE and the construction of the learning-proof menu, as well as the associated results on computation, convergence and solution bounds.

Paper structure. We start with a simple motivating example to fix ideas. In Section 3, we formally develop the ignorance equivalent approach and introduce our mathematical workhorse, the Geometric Attention Problem. In Section 4, we lay out

Liu [2007] for asset pricing; Maćkowiak and Wiederholt [2009] for monetary shocks; Van Nieuwerburgh and Veldkamp [2009] and Van Nieuwerburgh and Veldkamp [2010] for home bias and underdiversification in asset portfolios; or Dasgupta and Mondria [2018] for trade flows.



Figure 1: Introductory example with four actions (black dots) and two states (on either axis). The learning-proof menu is drawn in gray, and the ignorance equivalent of the RI problem, \tilde{a}^{RI} , is represented by a yellow dot.

the implications from our approach that that may be of more immediate use to applied researchers. Readers who are primarily interested in the practical takeaways of our approach may skip the theory developments from Subsection 3.3.1 onward. Section 5 illustrates the relevance of these new tools in three specific applications, and Section 6 concludes. With the exception of some immediate corollaries, all proofs are in the Appendix.

2 Illustrative Example

To illustrate the basic ideas behind rational inattention and visualize our ignorance equivalent approach, Figure 1 plots a simple choice problem with two states and a menu \mathcal{A} containing four actions.⁶ The coordinates of each black dot (a_1^k, a_2^k) report the state-specific payoffs from selecting action \mathbf{a}^k in state i = 1, 2. We first study three benchmark cases where the information technology is fixed.

A decision maker with no access to information can randomize over actions, but

⁶For this example, we use menu $\mathcal{A} = \{(0.2, 1.15), (1, 1.1), (1.5, 0.7), (1.7, 0.4)\}$, uniform prior $\pi = (0.5, 0.5)$ and information cost parameter $\lambda = 0.55$.

has to do so independent of the realized state. The yellow area highlights the feasible set of payoffs. The optimal choice is found via the supporting hyperplane that is perpendicular to her prior π . The action that generates the corresponding payoffs $(\tilde{a}^{NI} = a^3)$ is implemented unconditionally. Moreover, the decision maker is indifferent between the original menu \mathcal{A} and the singleton menu $\{\tilde{a}^{NI}\}$.

A decision maker with free access to full information would first learn the state and then select the payoff-maximizing action a^1 in state 1 and a^4 in state 2. Her expected utility is the same as she would get from unconditionally implementing the fictitious action $\tilde{a}^{\mathsf{FI}} = (a_1^1, a_2^4)$. Yet, adding \tilde{a}^{FI} to the menu does not make her any better off either, because it does not increase the achievable utility in either state. Taken together, these two observations imply that the decision maker is indifferent between the original menu \mathcal{A} , the singleton menu $\{\tilde{a}^{\mathsf{FI}}\}$ and the augmented menu $\mathcal{A} \cup \{\tilde{a}^{\mathsf{FI}}\}$. There is no other payoff vector with this property,⁷ and we refer to \tilde{a}^{FI} as the *ignorance equivalent* (IE) of the choice problem. We chose that name because of two properties: First, the decision maker would be willing to commit to blind implementation of \tilde{a}^{FI} (by picking the singleton menu $\{\tilde{a}^{\mathsf{FI}}\}$) rather than 'pick and choose' from the larger menu $\mathcal{A} \cup \{\tilde{a}^{\mathsf{FI}}\}$. We refer to this as 'voluntary ignorance'. Second, the decision maker is also willing to abandon action \tilde{a}^{FI} as long as he keeps access to the original menu \mathcal{A} (by selecting menu \mathcal{A} rather than $\{\tilde{a}^{\mathsf{FI}}\}$). We refer to this as a 'will to learn'.

The last benchmark case is a decision maker with free access to a partially informative binary signal. Upon observing realization $s \in \{1, 2\}$, the decision maker updates her belief to ρ^s , and selects the expected-utility maximizing action a^s . There is a unique payoff vector \tilde{a} that achieves the same expected utility overall and does no better under any posterior, and this fictitious action forms the IE of the choice problem. The constructive argument naturally extends to higher dimensions and any signal structure that generates I linearly independent posteriors.

Under rational inattention, the decision maker can choose her signal structure, but more informative signals are more costly. Under the parameters considered here, the optimal signal structure is the one we have considered last. Net of utility costs (indicated by the red arrow), it yields the same utility as the fictitious action \tilde{a}^{Rl} . We will show that \tilde{a}^{Rl} has the same intuitive properties over the choice problem with

⁷Formally, the payoffs are uniquely determined by the constraints $\boldsymbol{\pi} \cdot \tilde{\boldsymbol{a}}^{\mathsf{Fl}} = \sum_{i=1}^{I} \pi_i \max_{\boldsymbol{a} \in \mathcal{A}} a_i$ and $\tilde{a}_i^{\mathsf{Fl}} \leq \max_{\boldsymbol{a} \in \mathcal{A}} a_i \ \forall i$, for any number of states.

access to a continuum of costly signals as the IEs we have just introduced.

From IEs for various priors, one can construct a *learning proof menu* \mathcal{A} as drawn in gray. It contains the original menu, but we show that it makes the rationally inattentive decision maker no better off. No matter her prior, the decision maker would be willing to implement an action from $\overline{\mathcal{A}}$ unconditionally. As such, the learning proof menu essentially transforms the RI problem into a standard expected utility maximization problem.

3 Theory

We consider the standard Rational Inattention problem where an agent faces a finite menu of options with state-dependent payoffs and can condition her choice on arbitrary, but costly signals. More accurate signals are more costly, and we follow the literature [Caplin et al., 2018b, Matějka and McKay, 2015, Sims, 2003, 2006] in focusing on information-processing costs that are proportional to Shannon entropy.

3.1 Rational Inattention Problem

Formally, an agent has to implement an action from the finite **menu** \mathcal{A} . Payoffs from each action depend on an unknown state of the world. Each state $i \in \{1, ..., I\}$ occurs with positive **prior** probability $\pi_i > 0$. No two actions are payoff equivalent, and we identify an action $\boldsymbol{a} \in \mathcal{A}$ by its state-dependent payoffs $(a_1, ..., a_I) \in \mathbb{R}^I$.⁸ We denote the set of probability mass functions over the menu as $\mathcal{C}(\mathcal{A}) := \{p : \mathcal{A} \to [0, 1] \mid \sum_{\boldsymbol{a} \in \mathcal{A}} p(\boldsymbol{a}) = 1\}$. Building upon existing results [Matějka and McKay, 2015, Corollary 1], we do not model the details of the signal-generating process, and instead assume without loss of generality that the decision maker can directly select the conditional implementation probabilities $\boldsymbol{P} \in \mathcal{C}(\mathcal{A})^I$, where $P_i(\boldsymbol{a})$ denotes the probability of implementing action \boldsymbol{a} conditional on state i.

The optimal choice maximizes expected utility net of information processing costs, measured as the average reduction in entropy between prior and posterior. These costs are also known as the mutual information $\mathcal{I}(\boldsymbol{P}, \boldsymbol{\pi})$, and are equal to the expected Kullback-Leibler divergence between conditional and marginal choice probabilities

⁸Throughout, 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 $v_i > w_i \ \forall i$.

[Cover and Thomas, 2012],

$$\mathcal{I}(\boldsymbol{P}, \boldsymbol{\pi}) := \sum_{i=1}^{I} \sum_{\boldsymbol{a} \in A} \pi_i P_i(\boldsymbol{a}) \ln\left(\frac{P_i(\boldsymbol{a})}{p^{\boldsymbol{P}}(\boldsymbol{a})}\right),$$

where $p^{\mathbf{P}}(\mathbf{a}) = \sum_{i=1}^{I} \pi_i P_i(\mathbf{a})$ refers to the marginal implementation probability of \mathbf{a} .⁹ A proportionality constant $\lambda > 0$ translates the informational burden from nats into utils. Mathematically, the choice problem is parametrized by a triplet $(\mathcal{A}, \boldsymbol{\pi}, \lambda)$,

$$W(\mathcal{A}, \boldsymbol{\pi}, \lambda) = \max_{\boldsymbol{P} \in \mathcal{C}(\mathcal{A})^{I}} \sum_{i=1}^{I} \sum_{\boldsymbol{a} \in \mathcal{A}} \pi_{i} P_{i}(\boldsymbol{a}) a_{i} - \lambda \mathcal{I}(\boldsymbol{P}, \boldsymbol{\pi}).$$
(RI)

When π and λ are considered fixed, we refer to the value function as $W(\mathcal{A})$ for ease of notation. Clearly, the value function is nondecreasing under menu expansion, $W(\mathcal{A}) \leq W(\mathcal{A}')$ for all $\mathcal{A} \subseteq \mathcal{A}'$, since the agent can always restrict the support of \mathbf{P} to a subset of available actions at no cost.

3.2 Ignorance Equivalent

The central concept of our paper is the notion of an ignorance equivalent (IE) of a menu \mathcal{A} . To describe it, consider first a fictitious action with state-dependent payoffs $\boldsymbol{\alpha} \in \mathbb{R}^{I}$. If the agent is forced to implement this action blindly, his net utility equals $W(\{\boldsymbol{\alpha}\})$. If this option is added to his existing menu, his net utility equals $W(\mathcal{A} \cup \{\boldsymbol{\alpha}\})$. Voluntary ignorance occurs when $\boldsymbol{\alpha}$ is attractive enough so that $W(\{\boldsymbol{\alpha}\}) \geq W(\mathcal{A} \cup \{\boldsymbol{\alpha}\})$. Conversely, a will to learn is demonstrated when the agent weakly prefers the original menu over forced ignorance, $W(\mathcal{A}) \geq W(\{\boldsymbol{\alpha}\})$. The payoff vector is the 'ignorance equivalent' of menu \mathcal{A} if $\boldsymbol{\alpha}$ elicits both voluntary ignorance and a will to learn.

Definition 1. An ignorance equivalent of a menu \mathcal{A} under prior π and information cost λ is a payoff vector $\boldsymbol{\alpha} \in \mathbb{R}^{I}$ such that $W(\mathcal{A} \cup \{\boldsymbol{\alpha}\}) = W(\{\boldsymbol{\alpha}\}) = W(\mathcal{A})$.

The notion is analogous to the concept of a certainty equivalent for a lottery. The certainty equivalent for each lottery is unique, and it decreases with the agent's risk-aversion. Similarly, we will show that the IE is unique for any menu \mathcal{A} and its

⁹In line with conventional notation, we assume that $0 \ln 0 = 0$.

expected payoff decreases with the agent's information cost and weakly increases as new actions are added to \mathcal{A} (Corollary 1). It also satisfies a version of the *law of demand* with respect to changes in prior (Corollary 5). A simple example highlights the economic relevance of the IE.

Example 1. Abigail is looking to invest her wealth in one of \mathcal{A} different assets, which pay expected return \mathbf{a}_i in state of the world *i*. Being a rationally inattentive agent, Abigail will typically learn more about the state and improve her investment choices, but at a cost.

Consider an asset manager who is free to design a fund α that delivers return α_i in state *i*. The asset manager seeks Abigail's business, and has no information costs. When designing the fund, the asset manager realizes that Abigail may first learn some about the state, and then decide whether to invest in the offered fund—which could lead the asset manager to miss on some of Abigail's business and possibly be subject to adverse selection.

The IE is the answer to all the asset manager's problems. It ensures that Abigail wants to participate unconditionally and willingly forgoes any learning — thus enabling the asset manager to extract the maximal information rents. \diamond

The IE is the workhorse that ties together all the results of the paper. We use it as a "sufficient statistic" that condenses the information needed to identify the optimal choice and meaningfully compare across choices and menus, both from a theoretically perspective and for the purposes of numerical methods.

3.3 Geometric Attention Problem

The IE owes much of its power to convex geometry, and to the mathematical equivalence between (RI) and a simpler optimization problem which we call the Geometric Attention Problem (GAP). The starting point is a component-wise payoff transformation $\beta_i(\mathbf{a}) = e^{a_i/\lambda}$ that defines, for each action \mathbf{a} , an attention vector $\boldsymbol{\beta}(\mathbf{a}) \in (0, \infty)^I$. The mapping accentuates differences in payoffs when the information is cheap, and attenuates them when is costly.¹⁰ The convex hull over all attention vectors spanned

¹⁰Formally, consider two actions whose payoffs differ by a factor $k = \tilde{a}_i/a_i \in (0, 1)$ in state *i*. If instead we compare the relative size of the attention vectors, we observe that $\lim_{\lambda \to 0^+} e^{\tilde{a}_i/\lambda}/e^{a_i/\lambda} = 0$ and so action \tilde{a} attracts almost no attention relative to *a* in state *i*. For large information costs, $\lim_{\lambda \to \infty} e^{\tilde{a}_i/\lambda}/e^{a_i/\lambda} = 1$, and so they gather largely equal attention.



Figure 2: Visual representation of the Geometric Attention Problem. Attention vectors $\beta(\mathcal{A})$ are drawn as black dots, their convex hull \mathcal{B} is indicated in gray. The upper boundary $\partial^+ \mathcal{B}$ is drawn in blue. Indifference curves for w are drawn as dashed lines.

by \mathcal{A} ,

$$\mathcal{B} := \left\{ \sum_{\boldsymbol{a} \in \mathcal{A}} q(\boldsymbol{a}) \boldsymbol{\beta}(\boldsymbol{a}) \; \middle| \; q \in \mathcal{C}(\mathcal{A}) \right\} \subset \mathbb{R}_{+}^{I},$$

forms an "information possibilities set." The GAP simply selects the attention vector $\boldsymbol{b} \in \mathcal{B}$ that maximizes utility $w(\boldsymbol{b}) := \boldsymbol{\pi} \cdot \ln(\boldsymbol{b})$,

$$\max_{\boldsymbol{b}\in\mathcal{B}} w(\boldsymbol{b}). \tag{GAP}$$

Figure 2 attests to why we call the new problem "geometric." (GAP) transforms the menu into the convex polytope \mathcal{B} , which allows us to draw upon a vast literature within convex geometry. (GAP) also separates the role of prior beliefs from that of payoffs: Prior beliefs π determine the objective function. The attention cost parameter λ and the action payoffs a determine the feasible set \mathcal{B} . This allows us to isolate the geometric consequences of a change in one parameter.

3.3.1 Problem equivalence

Our first result makes the equivalence between (RI) and (GAP) formal and describes exactly how the optimal solutions between the two problems relate.

Theorem 1. If $P^* \in C(\mathcal{A})^I$ solves (*RI*), then

$$\boldsymbol{b}^* = \sum_{\boldsymbol{a} \in \mathcal{A}} \left(\sum_{i=1}^{I} \pi_i P_i(\boldsymbol{a}) \right) \boldsymbol{\beta}(\boldsymbol{a})$$

solves (GAP), with $w(\mathbf{b}^*) = W(\mathcal{A})/\lambda$. Conversely, if $\mathbf{b}^* = \sum_{\mathbf{a} \in \mathcal{A}} p(\mathbf{a})\beta(\mathbf{a})$ solves (GAP), then $\mathbf{P} \in \mathcal{C}(\mathcal{A})^I$ defined as

$$P_i(\boldsymbol{a}) = \frac{p(\boldsymbol{a})\beta_i(\boldsymbol{a})}{\sum_{\boldsymbol{a}'\in\mathcal{A}} p(\boldsymbol{a}')\beta_i(\boldsymbol{a}')}$$
(FONC)

solves (RI).

Proof. See Appendix A.

The conditional choice probabilities (FONC) are obtained from the first order conditions of the original problem, and have been reported previously [Matějka and McKay, 2015]. Plugging them back in to the objective function yields an optimization problem that depends only on the marginals p, but in its pure form restricts the choice of marginals to a finite subset of $C(\mathcal{A})$ – namely, those that are consistent with at least one set of conditionals satisfying Equation (FONC). We show that this constraint can be relaxed, and that the objective can be rewritten as $\pi \cdot \beta_i^{-1} (\sum_{a \in \mathcal{A}} q(a) \beta(a))$. To map this (GAP), we scale the objective by $1/\lambda$ and apply a simple change of variables.

Since (GAP) is strictly convex, it trivially admits a unique solution b^* , which fully characterizes the IE α .

Corollary 1 (Ignorance Equivalent). Each (*RI*) problem admits a unique IE, given by the pre-image of the corresponding (*GAP*) solution, $\boldsymbol{\alpha} = \boldsymbol{\beta}^{-1}(\boldsymbol{b}^*)$. The expected payoff $\boldsymbol{\pi} \cdot \boldsymbol{\alpha}$ weakly increases in the addition of new actions. It also weakly decreases in the information cost parameter λ , and strictly so whenever the solution to (*RI*) is non-degenerate.

Proof. The indifference conditions from Definition 1 require that $W(\mathcal{A}) = W(\{\alpha\})$ and $W(\mathcal{A} \cup \{\alpha\}) = W(\mathcal{A})$. Stated in terms of (GAP), the first requirement restricts $\beta(\alpha)$ to the indifference curve through b^* , while the second restricts $\beta(\alpha)$ to the separating hyperplane between the indifference curve and \mathcal{B} . The two sets intersect only at b^* , and hence the pre-image $\beta^{-1}(b^*)$ uniquely identifies the IE for any (RI) problem.

Since $W(\{\alpha\}) = \pi \cdot \alpha = W(\mathcal{A})$, the comparative statics follow from monotonicity of the original (RI) problem.

The solution to the original (RI) problem is unique as long as b^* admits a unique representation as a convex combination over points in $\beta(\mathcal{A})$.¹¹ The similarity of the attention vectors that span b^* captures the amount of learning that the decision maker undertakes. In the extreme case where the optimum is spanned by a single action, $b^* \in \beta(\mathcal{A})$, the decision maker forgoes learning altogether and blindly implements a single action. In all other cases, the optimal choice always involves learning. Learning is largest when $\beta(\alpha)$ is spanned by wildly different attention vectors, in which case the agent will closely tailor his action to the realized state. Lemma 2 in the appendix formalizes the link between the spanning surface and the amount of learning.

It is important to emphasize that merely because b^* can be written as a convex combination $\sum_{a \in \mathcal{A}} p(a)\beta(a)$ does *not* mean that the decision maker implements an unconditional lottery. Rather, the decision maker implements the optimal conditional probabilities P according to Equation (FONC).

3.3.2 Optimality conditions

The simple geometry of (GAP) allows us to succinctly characterize its optimum via linear inequality conditions.

Theorem 2. The solution to (GAP) is unique and fully characterized by either of the following two optimality conditions:

(a)
$$\nabla w(b^*) \cdot \beta(a) \leq 1$$
 for all $a \in \mathcal{A}$.

(b)
$$\nabla w(\mathbf{b}) \cdot \mathbf{b}^* \ge 1$$
 for all $\mathbf{b} \in \mathcal{B}$.

Proof. See Appendix A.

Figure 2 captures the geometric intuition for this result: Condition (a) says that \mathcal{B} lies weakly below the hyperplane that is tangent to the indifference curve at b^* . Condition (b) says that b^* lies above all hyperplanes that are tangent to the indifference curves at any suboptimal $b \in \mathcal{B}$. Both hold thanks to the convexity of (GAP).

¹¹This generates exactly the uniqueness conditions given by Matějka and McKay [2015].

Both types of optimality conditions are central to our paper. The first set is linear over the points in the convex hull \mathcal{B} . The condition identifies the optimum as the only point where $\nabla w(b^*)$ represents a supporting hyperplane for \mathcal{B} . Condition (a) has been stated before in terms of action payoffs, and forms the backbone for Caplin et al. [2018b].¹² One key observation is that the conditions jointly imply that inequality (a) binds for any attention vector that spans b^* .

To our knowledge, the second set of optimality conditions are new to the RI literature. Condition (b) is constructive in the sense that any feasible point $\boldsymbol{b} \in \mathcal{B}$ restricts the potential location of \boldsymbol{b}^* to a linear half-space dictated by the vector $\nabla w(\boldsymbol{b})$. A successive choice of feasible $\boldsymbol{b}^n \in \mathcal{B}$ then allows us to 'close in' on the optimum, and make precise statements about the true optimum based on numerical estimates.

3.3.3 Consideration set cardinality

In line with Caplin et al. [2018b], we refer to the set of actions that are chosen under an optimal model P^* as the agent's consideration set $\operatorname{support}(P^*) \subseteq \mathcal{A}$. Since \mathcal{B} is a finitely generated polytope, (GAP) readily generates an upper bound on the cardinality of the minimal consideration set. Indeed, Carathéodory's Theorem states that any point $\mathbf{b} \in \operatorname{conv.hull}(\mathcal{\beta}(\mathcal{A}))$ can be written as a convex combination of at most I+1 points in $\mathcal{\beta}(\mathcal{A})$.¹³ The minimal consideration set therefore contains at most I+1actions. This bound can be further strengthened by using the fact that \mathbf{b}^* is on the upper boundary $\partial^+\mathcal{B} = \{\mathbf{b} \in \mathcal{B} \mid \nexists \mathbf{b}' \in \mathcal{B} : \mathbf{b}' > \mathbf{b}\}$ by strict monotonicity of w. This simple intuition for finite type spaces complements the more general treatment of the cardinality of the consideration set in Jung et al. [2019].

Corollary 2. The minimal consideration set contains at most I actions.

Proof. See Appendix A.

3.3.4 Scale Invariance

The functional form of (GAP) has another separability feature that is particularly helpful for numerical evaluation: Scaling the feasible set \mathcal{B} with a positive constant along any dimension maintains optimality, even if the objective function is left intact.

¹²Necessity was highlighted previously by Matějka and McKay [2015].

¹³See for instance Eggleston [1958, Theorem 18].

Mathematically, component-wise scaling $\boldsymbol{b} \mapsto \boldsymbol{k} \otimes \boldsymbol{b} := (k_i b_i)_{i=1}^I$ merely offsets the objective value by a constant factor,

$$oldsymbol{\pi} \cdot \ln(oldsymbol{k} \otimes oldsymbol{b}) = \sum_{i=1}^{I} \pi_i \cdot \ln(k_i b_i) = oldsymbol{\pi} \cdot \ln(oldsymbol{b}) + oldsymbol{\pi} \cdot \ln(oldsymbol{k}).$$

As a consequence, the optimum scales by the same vector as the feasible set.

Corollary 3 (Axis Scaling). Consider any scaling vector $\mathbf{k} \in \mathbb{R}_+^I$. Attention vector \mathbf{b} solves (GAP) if and only if $\mathbf{k} \otimes \mathbf{b}$ solves $\max_{\mathbf{b} \in \mathbf{k} \otimes \mathcal{B}} w(\mathbf{b})$.

Proof. Since $w(\mathbf{k} \otimes \mathbf{b}) = w(\mathbf{k}) + w(\mathbf{b})$ for all $\mathbf{b} \in \mathcal{B}$, $w(\mathbf{b}^*) \ge w(\mathbf{b})$ for all $\mathbf{b} \in \mathcal{B}$ if and only if $w(\mathbf{k} \otimes \mathbf{b}^*) \ge w(\tilde{\mathbf{b}})$ for all $\tilde{\mathbf{b}} \in \mathbf{k} \otimes \mathcal{B}$.

Scalability greatly helps reduce floating point imprecision in our numerical algorithm. It also captures the fact that shifting a menu by a constant payoff vector, as in the Minkowski sum $\mathcal{A} + \{u\}$, does not affect the relative location of its IE. The IE shifts by the same vector, to $\boldsymbol{\alpha} + \boldsymbol{u}$. Indeed, the incentives for learning are unaffected because the payoff boost is independent from the action choice.

3.3.5 Relation to existing concavification approaches

Previous concavification procedures [Caplin et al., 2018b, Gentzkow and Kamenica, 2014, Kamenica and Gentzkow, 2011] express RI problems in terms of posterior beliefs. In the standard problem visualization over two states, the horizontal axis corresponds to posterior beliefs ρ in the likelihood of state 1, and each action defines a curve given by its expected payoff under posterior ρ , suitably corrected for the fact that extreme posteriors are more expensive to obtain. The optimal choice can then be characterized by the posteriors and action curves that locally support the concave upper envelope at the prior. Clearly, this approach is quite different from the one here, as best illustrated by Figure 2. Our axes correspond to states, not posteriors, and we represent actions by attention vectors, not curves.

Nevertheless, there is an intimate connection that arises by combining Corollary 3 and Theorem 2. Indeed, scaling by $\mathbf{k} = \nabla w(\mathbf{b}^*) \in \mathbb{R}^I_+$ moves the optimum to $\boldsymbol{\pi}$ by optimality condition (a) and allows us to express this optimal attention vector as a convex combination over points that correspond to the agent's posterior beliefs.¹⁴

¹⁴To see why, let b^* denote the optimum over \mathcal{B} . Note that $\nabla w(\pi) = 1$ and hence $\nabla w(\pi)$.

However, while scaling to the simplex offers intuition regarding the possible posterior beliefs, it has one obvious draw-back: One needs to know the optimal b^* in order to determine the scaling of \mathcal{B} . For conceptualization, this loop is not tragic – for computation, it is fatal. One fundamental advantage of our approach is that it yields a convex optimization problem where both the set of candidate points \mathcal{B} and the objective function w are explicitly defined.

3.4 Learning-Proof Menu

By drawing on the properties of (GAP), the IE approach can also be used to transform any (RI) problem into a standard expected utility maximization problem over a modified menu. We achieve this by forming a 'payoff possibilities frontier' by collecting the IEs across priors, and adding any statewise dominated payoffs.

Definition 2. The **learning-proof menu** generated by menu \mathcal{A} under information cost λ is equal to the set of payoff vectors

$$ar{\mathcal{A}} := \left\{ oldsymbol{a} \in \mathbb{R}^{I} \, \big| \, \exists oldsymbol{\pi} \in \Delta^{I-1} ext{ such that } oldsymbol{\pi} \gg oldsymbol{0} ext{ and } oldsymbol{lpha}^{oldsymbol{\pi}} \geq oldsymbol{a}
ight\},$$

where α^{π} denotes the IE of menu \mathcal{A} under prior π .

We show that the (RI) agent is always indifferent between \mathcal{A} and the larger menu $\overline{\mathcal{A}}$. Since the learning proof menu $\overline{\mathcal{A}}$ does not depend on the agent's prior, and always contains its own IE, the agent is essentially solving a standard expected utility maximization problem over $\overline{\mathcal{A}}$. In addition, the (GAP) geometry allows for a mathematically more concise definition of $\overline{\mathcal{A}}$, and the properties of the IEs generate intuitive comparative statics. We collect all these results in the following corollary.

Corollary 4 (Learning-proof menu). The learning-proof menu \overline{A} has the following properties:

(a) $\bar{\mathcal{A}}$ is equal to the Minkowski sum $\boldsymbol{\beta}^{-1}(\partial^+ \mathcal{B}) + \mathbb{R}^I_{\leq \mathbf{0}}$.

 $^{(\}mathbf{k} \otimes \boldsymbol{\beta}(\overline{\mathbf{a}})) = \sum_{i=1}^{I} 1 \cdot (\nabla_i w(\mathbf{b}^*)\beta_i(\mathbf{a})) \leq 1$, meaning $\boldsymbol{\pi}$ satisfies the optimality conditions (a) over the set $\mathbf{k} \otimes \boldsymbol{\beta}(\boldsymbol{A})$. At the same time, the posterior probability of state *i* conditional on implementing action \mathbf{a} can be written as $\pi_i P_i^*(\mathbf{a})/p^*(\mathbf{a})$ by Bayes rule. By (FONC), this is equal to $\beta_i(\mathbf{a})\pi_i/(\sum_{\mathbf{a}'\in\mathcal{A}} p^*(\mathbf{a}')\beta_i(\mathbf{a}')) = \beta_i(\mathbf{a})\nabla_i w(\mathbf{b}^*) = k_i\beta_i(\mathbf{a})$ for any chosen action. In other words, the posterior is equal to the location of action \mathbf{a} 's attention vector after scaling for any action within the consideration set.

- (b) $\mathcal{A} \subseteq \overline{\mathcal{A}}$, and the two menus share the same IE $\alpha^{\pi} \in \overline{\mathcal{A}}$ under any prior π .
- (c) \overline{A} is closed, bounded above and convex.
- (d) $\bar{\mathcal{A}}$ is equal to the intersection of halfspaces $\bigcap_{\pi\gg 0} \left\{ \boldsymbol{a} \in \mathbb{R}^{I} \mid \boldsymbol{\pi} \cdot (\boldsymbol{\alpha}^{\pi} \boldsymbol{a}) \geq 0 \right\}$.
- (e) Under the set-inclusion order, \overline{A} is weakly larger when new actions are added to A or when the information cost parameter λ decreases.

Proof. See Appendix A.

The learning-proof menu owes its name to property (b), which says that the agent gains nothing from his costly learning opportunities, since \overline{A} always contains its own IE. This stems from the fact that adding an IE or a statewise dominated payoff vector to the menu does not affect the upper boundary $\partial^+ \mathcal{B}$, and so the solutions to (GAP) remain unchanged under all priors. As a consequence, the learning-proof menu is an expected utility representation of the (RI) problem since IE α of \mathcal{A} is also the solution to

$$\max_{\boldsymbol{a}\in\bar{\mathcal{A}}}\boldsymbol{\pi}\cdot\boldsymbol{a}.$$
 (EU)

By Corollary 1 and Theorem 1, this expected utility maximization problem thus allows us to reconstruct the solution to (GAP) and (RI).

The (EU) representation also allow us to link changes in the prior with changes in the IE. For example, if an agent sees state i as more likely, at the expense of state j, then the IE payoff in state i increases while that of state j decreases.

Corollary 5. For any two priors $\boldsymbol{\pi}, \boldsymbol{\rho} \gg \mathbf{0}$, the corresponding IEs $\boldsymbol{\alpha}^{\boldsymbol{\pi}}$ and $\boldsymbol{\alpha}^{\boldsymbol{\rho}}$ satisfy $(\boldsymbol{\alpha}^{\boldsymbol{\pi}} - \boldsymbol{\alpha}^{\boldsymbol{\rho}}) \cdot (\boldsymbol{\pi} - \boldsymbol{\rho}) \geq 0$.

Proof. Optimality of α^{σ} in (EU) implies $\sigma \cdot \alpha^{\sigma} \geq \sigma \cdot a$ for $\sigma \in \{\pi, \rho\}$ and $a \in \{\alpha^{\pi}, \alpha^{\rho}\} \subseteq \overline{A}$. Combining the inequalities yields the desired expression.

The learning-proof menu may arise naturally in games with strategic incentives or robustness consideration. Returning to Example 1, imagine the asset manager does not know Abigail's prior, or is seeking investment in a wider population with heterogeneous beliefs. Is it still possible to ensure participation and capture all the information rents? Yes, it is. The learning-proof menu is the answer to all the manager's problems. By offering a rich menu of funds, it is optimal for Abigail or

any other agent to self-select her own IE without any learning. Note that same is not true if the asset manager were unsure of Abigail's information cost, since then the learning proof menu strictly increases under the set inclusion order by Corollary 4(e).

3.5 Menu expansion

RI and random utility models can have markedly different behavioral implications of menu expansion. In a multinomial logit model, for instance, the implementation probability of an action can be made arbitrarily small by merely adding payoff-equivalent duplicates of other actions in the menu [Debreu, 1960]. Matějka and McKay [2015] show that this is not the case under (RI). Intuitively, duplicate actions do not affect (GAP) and thus do not affect the weight placed on any non-duplicated action.

Moreover, adding a new action never increases the implementation probability of an existing action in random utility models. Matějka and McKay [2015] show by example that this is not the case in (RI), because the new action may generate learning opportunities that render a previously unchosen action attractive under some posteriors. We expand upon this idea by fully characterizing which actions are implemented with positive probability when added to a menu, either by themselves or in conjunction with other actions.

Both the IE α and the learning-proof menu \overline{A} are relevant to what happens as options are added to a menu A. If a single action a^+ is added to A, the new action is implemented with positive probability in all (RI) solutions if and only if the decision maker strictly prefers $\{\alpha, a^+\}$ to $\{\alpha\}$. If other actions are added at the same time, the learning-proof menu \overline{A} helps us determine whether a^+ will be payoff enhancing or not. If a^+ is outside the learning-proof menu \overline{A} , then a^+ will be implemented with positive probability after some menu expansion. Conversely, if a^+ is in the interior of \overline{A} , it will never be implemented no matter what further actions are added to the original menu.

Corollary 6 (Menu Expansion). For an (RI) problem $(\mathcal{A}, \pi, \lambda)$, let α denote the IE and $\overline{\mathcal{A}}$ the learning-proof menu. For any $\alpha^+ \in \mathbb{R}^I$, the following hold:

- (a) $W(\mathcal{A} \cup \{a^+\}) > W(\mathcal{A})$ if and only if $W(\{\alpha, a^+\}) > W(\{\alpha\})$.
- (b) $W(\mathcal{A}' \cup \{a^+\}) > W(\mathcal{A}')$ for some menu $\mathcal{A}' \supseteq \mathcal{A}$ if and only if $a^+ \notin \overline{\mathcal{A}}$.

Proof. See Appendix A.

Part (a) draws upon the optimality conditions in Theorem 2(a), which relate to the menu only through its IE α . It implies that the IE in some way succinctly summarizes the entire menu \mathcal{A} : Not only do the two menus achieve the same (RI) payoff, the addition of a new action a^+ either enhances the payoff from both menus or from neither. Caplin et al. [2018b] refer to the underlying mathematical expression as a 'market entry test' for a new action; we state it in terms of the 'sufficiency' of the IE for evaluating new actions.

However, the full menu has a larger potential for learning than its singleton summary $\{\alpha\}$. Whenever an additional action a^+ is payoff enhancing, the IE offers only a lower bound on the value function, $W(\mathcal{A} \cup \{a^+\}) \geq W(\{\alpha, a^+\})$. It is also possible that even if action a^+ is not payoff enhancing in menu \mathcal{A} , it may be payoff enhancing once further actions are added, that is, in a larger menu $\mathcal{A}' \subseteq \mathcal{A}$. The IE simply does not contain sufficient information to judge whether such a menu \mathcal{A}' exists; but the learning-proof menu does.

Part (b) says that action a^+ is payoff enhancing for at least one menu containing \mathcal{A} if and only if a^+ is not contained in the learning-proof menu $\overline{\mathcal{A}}$. We show this by explicitly constructing an additional action a' such that $W(\mathcal{A} \cup \{a^+, a'\}) > W(\mathcal{A} \cup \{a'\})$, which implies that a^+ is implemented with positive probability in any (RI) solution. Thus, although we initially constructed the learning-proof menu by considering changes in prior, the concept thus carries relevant information in situations where the prior is fixed.

Figure 1 illustrates the distinction between the two situations. If action a^+ lies above the dotted curved line, it satisfies criterion (a): The (RI) agent has an 'actual' interest in implementing a^+ in some contingencies. If a^+ lies between the curved line and the shaded region, it satisfies only (b): The agent has a 'potential' interest in a^+ if further actions are added to the menu. If a^+ lies inside the shaded region, it satisfies neither condition: The agent has no interest in a^+ as long as the actions in \mathcal{A} are available.

4 Practical Implications

In this section, we lay out some practical implications from the theory developments previously exposed. We highlight how our findings prove to be useful for numerical methods, and include some discussion on algorithm design in the last subsection.

4.1 Partial cover

Outside of a handful of cases that admit closed-form solutions, solving RI models requires computation methods which, naturally, are subject to numerical noise. We show here how to use noisy estimates of the optimum to characterize the consideration set of the true optimal solution.

We start by defining the notion of a *partial cover* as a generalization of a consideration set.

Definition 3. A set $A \subseteq \mathcal{A}$ is a *q*-cover of the (RI) problem $(\mathcal{A}, \pi, \lambda)$ if

$$\sum_{i=1}^{I} \pi_i \sum_{\boldsymbol{a} \in A} P_i^*(\boldsymbol{a}) \ge q \in [0, 1]$$

for all optimal P^* .

The consideration set is always a 1-cover, as are any of its supersets. Ideally, we hope to identify q-covers with high probability q and with the smallest cardinality |A| possible. The linear optimality conditions in Theorem 2(b) allow us to generate a q-cover for the (unknown) optimum from any suboptimal choice **b**. Formally, we define the ψ -score $z : \mathcal{A} \times \mathbb{R}^I_+ \to \mathbb{R}$ as $z(\mathbf{a}|\psi) = \psi \cdot \boldsymbol{\beta}(\mathbf{a}) - 1$, and use it to construct a q-cover for any choice of $q \in (0, 1)$.

Corollary 7. For any $b \in \mathcal{B}$ and any $q \in (0, 1)$, the set

$$A = \left\{ \boldsymbol{a} \in \mathcal{A} \mid z(\boldsymbol{a} | \boldsymbol{\nabla} w(\boldsymbol{b})) \ge -\frac{q}{1-q} \left(\max_{\boldsymbol{a}' \in \mathcal{A}} z(\boldsymbol{a}' | \boldsymbol{\nabla} w(\boldsymbol{b})) \right) \right\} \subseteq \mathcal{A}$$

is a q-cover.

Proof. Let $\bar{z} := \max_{a' \in \mathcal{A}} z(a' | \nabla w(b))$ denote the highest $\nabla w(b)$ -score across all actions. By Theorem 2(b), the optimal choice has a nonnegative expected $\nabla w(b)$ -score, $\sum_{a \in \mathcal{A}} p^*(a) z(a | \nabla w(b)) \ge 0$. Trivially, if $\bar{z} = 0$, only actions with $z(a | \nabla w(b)) = 0$ are implemented in (RI), and these are all contained in A. Otherwise, $\bar{z} > 0$ and we proceed by bounding $\nabla w(b)$ -scores above,

$$0 \leq \sum_{\boldsymbol{a} \in \mathcal{A}} p^*(\boldsymbol{a}) z(\boldsymbol{a} | \boldsymbol{\nabla} w(\boldsymbol{b})) \leq \sum_{\boldsymbol{a} \in \mathcal{A} \setminus \mathcal{A}} p^*(\boldsymbol{a}) \left(-\frac{q}{1-q} \bar{z} \right) + \sum_{\boldsymbol{a} \in \mathcal{A}} p^*(\boldsymbol{a}) \bar{z}$$
$$= -\frac{q}{1-q} \bar{z} (1-p^*(\mathcal{A})) + p^*(\mathcal{A}) \bar{z}.$$

Rearranging terms yields $p^*(A) \ge q$.

Corollary 7 can be readily deployed to assess computation accuracy. Let b be the researcher's numerical estimate of the optimum and let A be a partial cover with high q, say 95%, obtained from Corollary 7. The researcher may find that A is pretty large, perhaps close to the full menu A. This may indicate that computational error is substantial, particularly if a narrow consideration set was expected due to high information costs. Alternatively, the researcher may find that A has very few actions or that, from the perspective of the specific application, actions in A are clustered around only a handful of relevant values. In this case, the researcher has effectively identified the key features of the consideration set under the true optimum.

Corollary 7 can also be useful while searching for the right parameters to replicate a salient fact, say, a particular action \boldsymbol{a} being observed with frequency higher than 10%. The researcher does not need a very precise estimate \boldsymbol{b} of the optimum for each parameter value: As soon as the 90%-cover excludes the aforementioned action \boldsymbol{a} , the parameter value can be rejected.

In practice (see Section 5), we find that accurate algorithms yield estimates \boldsymbol{b} that are very close to the optimal (GAP) solution. The resulting q-covers typically have small cardinality even when q is close to one, making this approach very attractive for empirical research.

4.2 Dominated Actions

In many applications, it is even possible to rule out some dominated actions altogether — effectively finding a 1-cover that is significantly smaller than the menu. Sometimes, this is trivial: If an action delivers less payoff in each state than a blind lottery over other actions, it would never be chosen even under full information. By formulating the RI model as an expected utility problem, Corollary 6(b) extends this logic: Only actions that are on the boundary of the learning-proof menu are chosen with positive probability.

We go even one step further by combining numerical estimates and the optimality conditions in Theorem 2. Jointly, the optimality conditions (a) and (b) imply that any action with positive support has a $\nabla w(b^*)$ -score zero. By restricting the optimal gradient using numerical estimates, we can bound the feasible scores for some actions below zero, and thus rule them out for good.

Practically, the set of all bounding hyperplanes to $\mathcal{B} + \mathbb{R}^n_{\leq 0}$ is

$$ar{V}_1 = \left\{ oldsymbol{v} \in \mathbb{R}^I_{\geq 0} \mid oldsymbol{v} \cdot oldsymbol{eta}(oldsymbol{a}) \leq 1 \; orall oldsymbol{a} \in \mathcal{A}
ight\}$$

Note that any action that is interior to $\bar{\mathcal{A}}$ will satisfy $\boldsymbol{v} \cdot \boldsymbol{\beta}(\boldsymbol{a}) < 1$ for all $v \in \bar{V}_1$. By Theorem 2(a), $\nabla w(\boldsymbol{b}^*) \in \bar{V}_1$. Next, we take a numerical estimate \boldsymbol{b}^0 , perturb it slightly along each dimension, and consider the largest feasible attention vector along the perturbed ray.¹⁵ Together, this yields a finite set of near-optimal solutions $\widehat{B} = \{\boldsymbol{b}^0, ..., \boldsymbol{b}^I\} \subset \mathcal{B}$. Theorem 2(b) restricts the optimal attention vector \boldsymbol{b}^* to $\widehat{B} := \mathcal{B} \cap \{\boldsymbol{b} \in \mathbb{R}^I_+ \mid \nabla w(\boldsymbol{b}^k) \cdot \boldsymbol{b} \geq 1 \forall k\}$. Since $\nabla_i w(\boldsymbol{b}) = \pi_i / b_i$ is strictly decreasing in b_i , this also restricts $\nabla w(\boldsymbol{b}^*)$ to the hypercube

$$\bar{V}_2(\widehat{B}) := \bigotimes_{i=1}^{I} \left[\frac{\pi_i}{\max_{\widehat{B}} b_i}, \frac{\pi_i}{\min_{\widehat{B}} b_i} \right].$$

By combining the two feasibility constraints, we know that the optimal gradient $\nabla w(\mathbf{b}^*)$ is contained in $\bar{V}_1 \cap \bar{V}_2(\hat{B})$. We use this to rule out dominated actions.

Corollary 8. For any nonempty subset $\widehat{B} \subseteq \mathcal{B}$, the set

$$A = \left\{ \boldsymbol{a} \in \mathcal{A} \mid \max_{\boldsymbol{v} \in \bar{V}_1 \cap \bar{V}_2(\widehat{B})} z(\boldsymbol{a} | \boldsymbol{v}) \ge 0 \right\} \subseteq \mathcal{A}$$

is a 1-cover.

Proof. See text.

Computationally, finding dominated actions is significantly slower than finding a partial cover. Corollary 8 requires solving $2I + |\mathcal{A}|$ linear optimization problems (one for each bound in $\bar{V}_2(\hat{B})$, and one for each action to maximize $z(\boldsymbol{a}|\boldsymbol{v})$) while Corollary 7 relies only on the explicit score computations. Still, the dominated actions approach is useful in situations where accuracy is paramount.¹⁶

Often the researcher may be solving a finite approximation of a more complex (RI) problem, possibly with a continuous space for actions and states. A caveat is

¹⁵Formally, we solve $\max_{k \in \mathbb{R}, b \in \mathcal{B}} k$ subject to $b \ge k(b^0 + \varepsilon e^i)$, where $\varepsilon > 0$ is a small perturbation scalar and e^i denotes the unit vector in dimension *i*.

¹⁶For an even more conservative approach, the threshold for inclusion in A can be chosen slightly negative to offset the optimality tolerance of the linear program. This increases the cardinality of A, but fully hedges against numerical imprecision.

in order in such cases: Any cover obtained with Corollaries 7 and 8 will be specific to the action grid that is considered, and will rely on a correct characterization of the state distribution. If \mathcal{A} is merely a finite approximation to an infinite choice, it is possible that unmodeled actions generate learning opportunities that increase the attractiveness of actions in \mathcal{A} .¹⁷ Regarding approximations to the state space, our theory is simply not well suited the assess the accuracy of these estimates. That said, it appears in practice that the (GAP) approach also provides sensible numerical estimates when applied to a fine discretization of a continuous (RI) problem.

4.3 Precision metric

Ultimately, the goal of RI models is to rationalize observed patterns of behavior. As such, the primary object of interest in the optimization problem (RI) is the optimal conditional choice P^* . Be it to compare model predictions to empirical data, or to write a stopping criterion for numerical methods, the researcher eventually needs to decide when two conditional choices are 'similar'. As a basis for that call, we here present a notion of distance based on the IE that is both parsimonious and overcomes shortcomings of alternative metrics.

Definition 4. The *IE distance* between choices $P, P' \in \mathcal{C}(\mathcal{A})^I$ is defined as

$$d_{\mathsf{IE}}(\boldsymbol{P},\boldsymbol{P}') := \sqrt{\sum_{i=1}^{I} \pi_i \left(\boldsymbol{\alpha}^{\boldsymbol{P}} - \boldsymbol{\alpha}^{\boldsymbol{P}'}\right)^2},$$

where $\boldsymbol{\alpha}^{\boldsymbol{Q}}$ denotes the implied IE under choice \boldsymbol{Q} , with $\alpha_i^{\boldsymbol{Q}} := \beta_i^{-1}(\mathsf{E}[\beta_i(\boldsymbol{a})|\boldsymbol{a} \sim Q_i]).$

The IE distance is a standard Euclidean distance between the payoff vectors α^{P} and $\alpha^{P'}$, weighted by the prior probability for each state. The weights ensure that the distance is unaffected by a payoff-irrelevant splitting of states. Nonnegativity, symmetry and the triangle inequality are directly inherited from the standard Euclidean distance.

¹⁷If there are actions that are excluded in the model, perhaps as an approximation, one can rule out only actions that are interior to the learning-proof menu by dropping $\bar{V}_2(\hat{B})$ from Corollary 8. By Corollary 6(b), these interior actions are not in the consideration set of any finite menu $\mathcal{A}' \supset \mathcal{A}$. And since the optimal choice over a finite state space always has finite support [Jung et al., 2019], this extends to arbitrary menus. Practically speaking, we expect this approach to be effective in situations where the considered menu \mathcal{A} is large relative to I.

However, the IE distance fails to distinguish between choices that imply the same IE. From a computational perspective, this may be a satisfactory compromise in the interest of parsimony. Indeed, we now show that distance d_{IE} provides a suitable convergence criterion whenever (**RI**) admits a unique solution. From any sequence of choices that converge to the solution under d_{IE} , we can then construct choices that converge both in terms of marginal and conditional probabilities.

Lemma 1. Suppose (*RI*) admits a unique solution P^* . For any sequence of choices $\{P^n\} \subset C(\mathcal{A})^I$, let $\{Q^n\}$ be defined from (FONC) as

$$Q_i^n(\boldsymbol{a}) := \frac{\beta_i(\boldsymbol{a}) \sum_{j=1}^I \pi_j P_j(\boldsymbol{a})}{\sum_{\boldsymbol{a}' \in \mathcal{A}} \beta_i(\boldsymbol{a}') \sum_{j=1}^I \pi_j P_j(\boldsymbol{a}')}.$$

If $\mathbf{P}^n \to \mathbf{P}^*$ according to d_{IE} , then $Q^n \to P^*$.

Proof. See Appendix B.

The IE distance d_{IE} is ideally suited to serve as stopping criterion for numerical solution methods, as it penalizes numerical noise whenever it leads to substantial payoff differences while ensuring that conditional choices converge to the actual optimum. In this sense, the IE distance strikes a balance between other common stopping criteria: Methods that rely on objective values alone can lead to noisy estimates of the model's behavioral implications, since several conditional choices may – and often do [Jung et al., 2019] – share very similar objective values. At the other end, a straightforward comparison between the probability vectors \boldsymbol{P} and \boldsymbol{P}' or their associated marginals p and p', as in $d_{\mathsf{Pr}}(p,p') := ||(p(\boldsymbol{a}) - p'(\boldsymbol{a}))_{\boldsymbol{a} \in \mathcal{A}}||$, treats all actions as equally distinct. However, numerical (RI) estimates over large menus err both along the extensive — the consideration set — and intensive margin — the probability distribution. The vector comparison d_{Pr} disproportionately penalizes errors on the extensive margin, while d_{IE} recognizes when consideration sets contain actions with similar payoff vectors.

4.4 Algorithm design

The geometry of (GAP) lends itself to numerical methods for general finite RI problems or for discrete approximations to continuous RI problems. Standard algorithmic techniques for convex problems perform well, and the reduced dimensionality of (GAP) brings obvious gains in performance. We provide an algorithm based on sequential quadratic programming and active set methods (see, e.g. Judd [1998]), and using d_{IE} as a stopping criterion. The codes are available at https://github.com/ mmulleri/GAP-SQP and a detailed explanation of the code is provided in Appendix B.

For discrete approximations to continuous problems, large state and action spaces are needed. This can routinely lead to memory issues when storing the payoff matrix.¹⁸ We address this problem by starting with a coarse grid over actions, and increasing the grid precision stepwise. At each step, we compute the optimal attention vector \boldsymbol{b}^k and then include K actions from the finer grid with the highest $\nabla w(\boldsymbol{b}^k)$ score. We increase grid precision once the 99% cover stabilizes. When K is large relative to the optimal consideration set, this approach can approximate large action spaces without running into memory management issues.¹⁹ And while the numerical estimates of the algorithm depend on the path of subgrids, any partial covers computed in the last round accurately describe the optimal choice over the *entire* menu \mathcal{A} .

Although the optimization methods we use are relatively unsophisticated, we find that our algorithm performs favorably when compared to other state-of-the-art techniques that are typically used to estimate (RI) models, both in terms of speed and accuracy. The methods and ideas can also be combined with other solution methods to yield further gains in performance.²⁰

5 Applications

In this section, we illustrate by way of example that both the conceptual framework and the computationally tractable algorithm have the potential to expand the purview of further research. We consider three applications: The first is a monopolist problem with uncertain demand as proposed by Matějka [2016]. We use this wellknown application to benchmark the GAP-SQP algorithm described in Section 4.4

¹⁸In the portfolio optimization (Section 5.2) the associated $300^2 \times 300^2$ payoff matrix would require 64.8Gb of memory.

¹⁹Note however that our results are limited to a specific state space. In particular, the output may not approximate the solution of a continuous state distributions even when the distribution is discretized over a very fine grid.

²⁰For instance, one may use the Blahut-Arimotho algorithm with a very high tolerance error to create a starting guess for the GAP-SQL algorithm. Or one may replace the SQP approach with more advanced convex optimization algorithms.

against existing methods, focusing primarily on speed. The second is a portfolio choice problem with a massive state and action space proposed by Jung et al. [2019]. We primarily use it to highlight the precision of GAP-SQP and showcase the more robust behavioral predictions that we develop in Sections 4.1 and 4.2. The third applications is a task assignment problem that is novel to the RI literature. It illustrates that the ideal scenario for the (GAP) approach – finite state spaces coupled with rich action spaces – arises naturally in economically relevant problems.

5.1 Sticky Prices [Matějka, 2016]

Our first illustration is based on the "rationally inattentive seller" model of Matějka [2016]. A monopolistic seller has a per unit input cost of 1 and sets the price p facing an isoelastic demand function whose elasticity, $\frac{d+1}{d}$, is a random variable uniformly distributed. Profits are given by $\Pi(d, p) = p^{-\frac{d+1}{d}}(p-1)$, where the demand variable d is the ex-ante unknown state and the price p corresponds to the seller's action.²¹ As in Matějka [2016], actions and states are discretized. As a benchmark we use a grid of 200×200 points, and we will improve the grid precision to increase the computational demands of the problem without introducing any further complexity in the model.

For comparison to our base routine described in Section 4.4, we also solve the model using the Blahut-Arimoto (BA) algorithm, a solution method that originated in Rate Distortion theory and has recently gained some usage in RI problems. As with our GAP-SQP algorithm, the BA algorithm is guaranteed to converge to the optimum and operates with a reduced dimensionality, updating the marginal distribution over actions. We implement both algorithms in MATLAB.²²

Figure 3 documents the running times in seconds across a range of information costs λ for our benchmark case with a grid of 200 × 200 points. As shown in panel (a), the GAP-SQP algorithm terminates in less than 0.05 seconds in all runs, with minimal differences across information costs. The BA algorithm, reported in panel

²¹The demand variable *d* is uniformly distributed in $(\frac{1}{9}, \frac{1}{2})$, following Matějka [2016], Section 4.2. Matějka [2016] assumes a channel capacity constraint rather than information being acquired at a cost. To match the channel capacity constraint of half a bit we find that we need to set $\lambda = 0.0053$.

²²We use a desktop computer with 16 GB of RAM and an Intel(R) 3.20 GHz Core i7-8700 processor on a Windows 10 Enterprise 64-bit operating system. For further discussion of the BA algorithm, see Caplin et al. [2018b] and Cover and Thomas [2012]. Matějka [2016] instead used proprietary software AMPL/LOQO to solve for the joint probability distribution. Due to licensing restrictions, we were not able to use AMPL/LOQO to document running times. A comparable, freely available solver (IPOPT) was substantially slower and less precise than both the GAP-SQP and BA algorithms.



Figure 3: Running times across information costs

(b), runs in about one second when information costs are very low—and thus the solution is very close to the full information benchmark—but is substantially slower for higher information costs, up to 20 seconds. Both algorithms achieve very similar objective values, with the GAP-SQP algorithm outperforming by about 10^{-8} .²³

Next we ratchet up the computation burden by increasing each grid precision up to 600 points, for a total grid points of $600^2 = 360,000$. As shown in Figure 4(a), running times scale roughly linearly for the GAP-SQP algorithm. Even at a 600×600 grid running times stay well below half a second. Figure 4(b) shows that the BA algorithm also scales well—though this means that computing times approach two minutes for the largest grids.

We also compute the set of dominated actions as well as the 95% cover using the output from our GAP-SQP algorithm. Figure 5 displays the GAP-SQP numerical solution as solid bars over the full price grid — with insets at two points of the full support of prices for visibility. The 99% cover is indicated with a dark blue background. It is identical to the consideration set of the numerical solution. Thus, even if we had low confidence in the accuracy of the algorithm, we would be able to conclude that prices outside of this set occur with no more than 1% probability. Indeed, the main point in [Matějka, 2016] that optimal pricing behavior is discrete, clustering mass on a comparatively small number of points. This observation can also be made by looking at the set of non-dominated actions (light blue background): It shows that the vast majority of prices are not used in any contingency. The validity

²³Numerical solutions for $\lambda = 0.0053$ are also very similar to those reported in Matějka [2016]. See online Appendix C.1. We thank Filip Matějka for sharing his numerical output.



Figure 4: Running times across grid precision



Figure 5: Partial cover and undominated actions

of the claim does not rely on having found the true optimum of the (RI) problem, which makes it significantly more robust.

5.2 Portfolio Choice [Jung et al., 2019]

Our second application considers the portfolio choice problem of Jung et al. [2019], who illustrate that RI can explain low rates of household portfolio rebalancing. In this problem, an investor with unit wealth designs a portfolio composed of three uncorrelated assets, without restrictions on short sales or overall leverage. The investor has constant absolute risk aversion (CARA) utility $u(x) = -e^{-\alpha x}$ with risk aversion parameter α . Asset zero is a safe asset with constant return 1.03. The returns from risky assets j = 1 and j = 2 are each modeled as the sum of two independent random variables around a slightly higher mean return, $1.04 + Z_j + Y_j$. The random variable $Z_j \stackrel{\text{iid}}{\sim} \mathcal{N}(0, \sigma_z^2)$ reflects factors that are inherently unforeseeable. The random variable Y_j reflects factors that are not known at the outset, but can be learned at a cost. Each portfolio $(\theta_1, \theta_2) \in \mathbb{R}^2$ describes an available action, where θ_j is the position in risky asset j and $\theta_0 := 1 - \theta_1 - \theta_2$ is the position on the safe asset. The expected utility from the portfolio conditional on state $\mathbf{Y} = (Y_1, Y_2)$ is

$$U(\boldsymbol{\theta}, \boldsymbol{Y}) = \mathsf{E}\left[u\left(1.03\theta_0 + \sum_{j=1}^2 (1.04 + Z_j + Y_j)\theta_j\right) \middle| \boldsymbol{Y}\right].$$
 (1)

We follow Jung et al. [2019] and assume that \mathbf{Y} follows a discrete distribution over a 300 × 300 grid that is obtained from a normal distribution $\mathcal{N}(\mathbf{0}, 0.02^2 I)$ truncated at three standard deviations. We report results for parameter values $\alpha = 1$, $\lambda = 0.1$ and $\sigma_z = 0.0173$.²⁴

We approximate the continuous menu $(\theta_1, \theta_2) \in \mathbb{R}^2$ by first (without loss of generality) imposing the upper and lower bounds given by the full-information solution, and then iteratively doubling the grid resolution using 99% covers until we reach $513 \times 513 = (2^9 + 1)^2$ points.²⁵ Jung et al. [2019] instead use a variant of the Blahut-Arimoto algorithm that optimizes the points of support at each step of the algorithm. We refer to this algorithm as JKMS. The algorithms reach a comparable objective value, with GAP-SQP only mildly outperforming JKMS.²⁶ Both algorithms perform significantly better than approximating the objective with a second-order polynomial to obtain a Linear-Quadratic Gaussian (LQG) problem (for details, see Online Appendix C.2),²⁷ an approach that is common in the applied literature.²⁸

Turning to the behavioral implications, Figure 6(a) shows the estimated portfolio

²⁴These parameters correspond to scenario B in Jung et al. [2019]. Although not shown, the GAP-SQP solution has larger support and achieves a higher objective value than the JKMS solution in all four parameter scenarios.

²⁵The iterative approach (see Section 4.4 and Appendix B) allows us to handle a large action grid. However, it does not reduce the memory demands imposed by the large state space. In order to compute the solution at the same state grid resolution as Jung et al. [2019], we opted to run the algorithm on a computational cluster.

²⁶The solution published in Jung et al. [2019] closes 0.607153 of the payoff gap between no and full information, while GAP-SQP closes 0.614877 of the gap. For a comparison of the statewise payoff distribution across algorithms, see Online Appendix C.2.

²⁷The (continuous) LQG solution closes roughly 0.4843 of the payoff gap between no and full information.

²⁸Examples of LQG models include Kacperczyk et al. [2016], Luo et al. [2017], Mondria [2010], Van Nieuwerburgh and Veldkamp [2009, 2010].



Figure 6: Portfolio distributions under GAP-SQP (blue), JMKS (orange) and LQG (black). In panel (a), the circle size of each portfolio $\boldsymbol{\theta}$ is proportional to the probability weight $p(\boldsymbol{\theta})$, and the probability that the LQG solution falls between any two dashed contour lines is equal to 0.2.

choice probabilities across all three algorithms. The LQG solution stands out as the only continuous solution, but even the two discrete solutions are measurably different. Jung et al. [2019] caution that their solution method may miss solutions with a larger support, and this is exactly what we find with GAP-SQP. The main point in Jung et al. [2019], that portfolio rebalancing is relatively rare, remains valid and indeed the consideration set shrinks for higher information costs. Quantitatively, though, GAP-SQP finds that portfolio rebalancing is substantially more common and, occasionally, the investor makes small adjustments.

The partial covers displayed in Figure 6(b) allow more robust statements regarding the true optimum given the state and action grid that we use. Contrary to the JKMS estimates, it appears that the investor actually rarely takes large short positions simultaneously on both risky assets. This may be suggestive that the investor looks for good news rather than bad. Another pattern that arises from Figure 6(b) is that the RI investor only selects portfolios from a circle, hinting that some further analytic results are possible — and may help elucidate the relative role of risk aversion and



Figure 7: Payoff distribution across choices, smoothed with a kernel density estimate.

information processing costs, for instance.²⁹ Overall, the example illustrates the need for more robust estimation techniques that not only deliver a 'better estimate', but allow a valid characterization of the true optimal choice.

Figure 7 plots the statewise payoff distribution $U(\boldsymbol{\theta}, \boldsymbol{Y}) - \lambda \mathcal{I}$, assuming $(\boldsymbol{\theta}, \boldsymbol{Y})$ is distributed according to the numeric solution of GAP-SQP (blue) and the information cost \mathcal{I} is borne unconditionally. For comparison, the figure also draws the payoff distribution under no information ($\lambda \to \infty$, grey dotted) and free information ($\lambda \to 0$, grey dashed). The IE (pink) yields the same expected utility as the optimal choice, but – to dissuade learning – it avoids the lowest payoffs in a way that mimics the full-information distribution.

5.3 Task Assignment

Our last application is designed to illustrate how the GAP geometry is particularly helpful in RI problems when the action space is naturally large and discrete. A manager has to assign N workers across three tasks $\{0, 1, 2\}$. Either task one or task two is critical; task zero is never critical and represents dismissal. All but one of the workers are skilled. The unskilled worker is not productive and prevents a skilled worker from contributing (if there are any assigned to the same task). Output is thus determined by the number of skilled workers assigned to the critical task, minus the unskilled worker if he is also assigned to the critical task, n^* . We assume a simple form of decreasing returns to scale, letting output be given by the production function

 $^{^{29}}$ A more thorough investigation of this conjecture is outside of the scope of this research.

 $\Phi(n^*) = \sum_{n=1}^{n^*} \delta^n \text{ for } \delta = 0.9.$

Despite its simple description, the task assignment problem generates a complex optimization problem. There are 2N possible states, indicating which task is critical $c \in \{1, 2\}$ and the identity of the unskilled worker $w \in \{1, ..., N\}$. An action is a task assignment a_w for each worker w, that can be summarized as a vector $\mathbf{a} \in \{0, 1, 2\}^N$. There are 3^N such vectors, and at least 2^N that are optimal under some information structure. We consider a fully symmetric setup with N = 10 workers, resulting in 20 states and 6124 potential assignments.

Figure 8 summarizes expected output, information flow, and optimal assignment strategies for a range of information costs $\lambda \in [.01, 100]$. As information costs increase, the manager uses four distinct allocation strategies (indicated by letters A to D).

When information costs are low, the manager aims for the full-information solution, dismissing the unskilled worker and assigning everyone else to the critical task (Strategy A). Initially, she acquires nearly all the information and consistently achieves the full-information benchmark output $\Phi(9)$. As information cost goes up, the manager occasionally misidentifies the unskilled worker or, with much lower probability, the critical task.³⁰

When λ reaches a certain threshold, the manager changes tack, sending all workers to the task that she identifies as critical (Strategy B). Because all learning on workers is forgone, we see a discrete drop in information acquisition that compensates the manager for the reduction in expected output due to the unskilled worker.³¹ Because the manager is initially very accurate at identifying the critical task, output once again is near constant at $\Phi(8)$ —slightly lower than in the full-information benchmark output $\Phi(9)$. As information costs increase, so does the likelihood of an extreme zerooutput assignment resulting from sending all workers to the misidentified critical task. Output volatility peaks under this strategy.

Once information costs are high enough, the manager aims to hedge and acquires little information. The first hedging strategy is to send all but one of her workers to the task she believes to be critical, but does so with hardly any information on their skills (Strategy C). This strategy is advantageous because the unskilled worker does no harm when he is by himself, while the presence of a single skilled worker is useful

³⁰Although not visible in Figure 8(b), there is a small but positive probability of output $\Phi(0)$ or $\Phi(1)$ resulting from the critical task been misidentified.

³¹Strategy A's expected output decreases with the probability of misidentifying tasks or workers, while output under Strategy B is affected only by task misidentification, by construction.



(a) Mean output (solid black) and information acquisition (dashed blue).



Figure 8: Optimal management strategies under varying information costs.

if the critical task is misidentified.³² As information costs increase even further, the manager assigns an equal number of workers to either tasks (Strategy D), hedging output as much as possible. What little information she still gathers concerns both the worker and the task, but expected output quickly approaches the no-information benchmark.

6 Conclusion

Rational inattention offers a promising research agenda that recognizes that economic agents interact with, and shape, their information environment. As advocated by Sims [2006], researchers need to go beyond the linear-quadratic case if RI models are to be of use for applied work. Notwithstanding recent progress, the sheer size of the information structure that results from RI models can bring substantial challenges, both conceptual and computation.

We hope our contributions here advance our understanding, as well as our ability to solve, RI models. The concept of the ignorance equivalent can effectively summarize the solution to a RI problem, with appealing properties for comparative statics as well as zero-sum games. It also forms the basis for the learning-proof menu, which recasts the RI model in the familiar framework of expected utility maximization. We have also provided an extended toolkit for numerical methods, with a focus to methods that allow a researcher to check the accuracy of the computed solution. More broadly, we hope that the GAP problem proves a fertile ground for expertly-designed algorithms that enable more complex RI models.

There remains important progress to be made in RI models for applied work. A key issue is how to assess RI models empirically. Recent progress by [Caplin et al., 2018a] outlines the kind of ideal dataset that would allow to recover information costs. However, we still do not know how to bring RI models to real-world data, and how to evaluate the resulting fit against, say, rational-agent models with incomplete, but static, information; or behavioral alternatives.

 $^{^{32}\}mathrm{To}$ an outside observer, firm output is most unpredictable over this range, as output entropy peaks under this strategy.

Appendix

A Proofs of Theoretical Results

Proof of Theorem 1:

We establish the equivalence between (RI) and (GAP) by means of a relaxed problem, where the decision maker can separately choose the marginals $p \in C(\mathcal{A})$ and conditionals $\mathbf{P} \in C(\mathcal{A})^{I}$, yielding

$$\max_{p \in \mathcal{C}(\mathcal{A}), \mathbf{P} \in \mathcal{C}(\mathcal{A})^{I}} \sum_{i=1}^{I} \sum_{\boldsymbol{a} \in \mathcal{A}} \pi_{i} P_{i}(\boldsymbol{a}) a_{i} - \lambda \sum_{i=1}^{I} \pi_{i} D_{\mathsf{KL}}(P_{i} \parallel p), \qquad (2)$$

and $D_{\mathsf{KL}}(P_i || p)$ denotes the Kullback-Leibler divergence between P_i and p.

To show the equivalence with (RI), assume pair p, P solves (2). Let $p^{P} = \sum_{i=1}^{I} \pi_{i} P_{i}$. Then

$$\lambda \sum_{i=1}^{I} \pi_i \left(D_{\mathsf{KL}}(P_i \| p) - D_{\mathsf{KL}}(P_i \| p^{\mathbf{P}}) \right) = \lambda D_{\mathsf{KL}}(p^{\mathbf{P}} \| p) \ge 0,$$

with strict inequality whenever $p \neq p^{\mathbf{P}}$. Thus optimality of p, \mathbf{P} requires $p = p^{\mathbf{P}}$, and thus the relaxed problem (2) has the same optimal value and optimal conditional \mathbf{P} than (RI).

For equivalence with (GAP), assume pair p, P solves (2). The necessary first order conditions to (2) imply

$$\pi_i a_i - \lambda \pi_i \ln \left(\frac{P_i(\boldsymbol{a})}{p(\boldsymbol{a})} \right) - \lambda \pi_i = \mu_i,$$

where μ_i denotes the Lagrange multiplier associated with the constraint $\sum_{\boldsymbol{a}\in\mathcal{A}} P_i(\boldsymbol{a}) =$ 1. Re-arranging and solving for μ_i shows that optimality of p, \boldsymbol{P} implies that \boldsymbol{P} satisfies (FONC) given p. Thus without loss of generality we can write (2) exclusively in terms of marginals $p \in \mathcal{C}(\mathcal{A})$,

$$\max_{p \in \mathcal{C}(\mathcal{A})} \lambda \sum_{i=1}^{I} \pi_i \ln \left(\sum_{\boldsymbol{a} \in \mathcal{A}} p(\boldsymbol{a}) \beta_i(\boldsymbol{a}) \right).$$
(3)

For any $p \in \mathcal{C}(\mathcal{A})$ we have $\boldsymbol{b} = \sum_{\boldsymbol{a} \in \mathcal{A}} p(\boldsymbol{a})\beta_i(\boldsymbol{a}) \in \mathcal{B}$; and for any $\boldsymbol{b} \in \mathcal{B}$ there is at least one $p \in \mathcal{C}(\mathcal{A})$. To complete the equivalence with (GAP), note that the objective function in (3) is simply $\lambda w(\boldsymbol{b})$.

Proof of Theorem 2: Since w is strictly concave over a convex domain \mathcal{B} , it admits a unique maximum. We first show that the optimum \mathbf{b}^* necessarily satisfies both conditions. Indeed, consider any $\mathbf{b} \in \mathcal{B} \setminus \{\mathbf{b}^*\}$ and let $\boldsymbol{\eta} : [0,1] \to \mathcal{B}$ be defined as $\boldsymbol{\eta}(t) = t\mathbf{b}^* + (1-t)\mathbf{b}$. The function $w \circ \boldsymbol{\eta}$ is strictly increasing: If t < t', then

$$(w \circ \boldsymbol{\eta})(t') > \frac{t'-t}{1-t}w(\boldsymbol{\eta}(1)) + \frac{1-t'}{1-t}w(\boldsymbol{\eta}(t)) \ge (w \circ \boldsymbol{\eta})(t),$$

where the first inequality follows from strict concavity of w and the second from optimality of \mathbf{b}^* , as $w(\boldsymbol{\eta}(1)) = w(\mathbf{b}^*) \ge w(\boldsymbol{\eta}(t))$. Since the derivative of $(w \circ \boldsymbol{\eta})$ is equal to $\nabla w(\boldsymbol{\eta}(t)) \cdot (\mathbf{b}^* - \mathbf{b})$, and $\nabla w(\mathbf{b}') \cdot \mathbf{b}' = 1$ for all $\mathbf{b}' \in \mathcal{B}$, the nonnegativity of $(w \circ \boldsymbol{\eta})'(1)$ yields condition (a) and the nonnegativity of $(w \circ \boldsymbol{\eta})'(0)$ yields (b).

We show sufficiency through the contrapositive: If b^* is not optimal, then it satisfies neither condition. Let b equal the true optimum and define $\eta(t)$ as above. The function $(w \circ \eta)$ is now strictly decreasing since for any t < t',

$$(w \circ \boldsymbol{\eta})(t) \ge \frac{t'-t}{t'} w(\boldsymbol{\eta}(0)) + \frac{t}{t'} w(\boldsymbol{\eta}(t')) > (w \circ \boldsymbol{\eta})(t'),$$

since $w(\boldsymbol{\eta}(0)) > w(\boldsymbol{\eta}(t'))$ by uniqueness of the optimum. At t = 1, the condition $(w \circ \boldsymbol{\eta})'(t) < 0$ violates (a) and at t = 0 it violates (b).

Lemma 2 (Learning). Assume menus \mathcal{A}^1 and \mathcal{A}^2 have the same IE α and each admit a unique optimal conditional choice \mathbf{P}^1 and \mathbf{P}^2 . Let \mathcal{B}^1 and \mathcal{B}^2 denote the convex hull of their β -images.

If
$$\mathcal{B}^1 \cap \left\{ \boldsymbol{b} \in \mathbb{R}^I_+ \mid \boldsymbol{\nabla} w(\boldsymbol{\beta}(\boldsymbol{\alpha})) \cdot \boldsymbol{b} = 1 \right\} \subseteq \mathcal{B}^2$$
, then $\mathcal{I}(\boldsymbol{P}^1, \boldsymbol{\pi}) \leq \mathcal{I}(\boldsymbol{P}^2, \boldsymbol{\pi})$.

Proof. By Theorem 1, the optimal marginals p^1 and p^2 correspond to the weights in the convex combination that describes $\beta(\alpha)$. Together with the first order condition (FONC), this simplifies mutual information to

$$\mathcal{I}(\boldsymbol{P}^{k},\boldsymbol{\pi}) = \sum_{i=1}^{I} \sum_{\boldsymbol{a} \in \mathcal{A}^{k}} \pi_{i} p^{k}(\boldsymbol{a}) h\left(\frac{\beta_{i}(\boldsymbol{a})}{\beta_{i}(\boldsymbol{\alpha})}\right)$$

for $h(x) = x \ln(x)$. By Theorem 2, actions that are implemented with positive probability map into hyperplane $H_{\alpha} := \{ \boldsymbol{b} \in \mathbb{R}^{I}_{+} \mid \boldsymbol{\nabla} w(\boldsymbol{\beta}(\boldsymbol{\alpha})) \cdot \boldsymbol{b} = 1 \}$. The condition $\mathcal{B}^{1} \cap H_{\alpha} \subseteq \mathcal{B}^{2}$ therefore implies that each $\boldsymbol{\beta}(\boldsymbol{a})$ with $p^{1}(\boldsymbol{a}) > 0$ can be written as a convex combination over points in $\boldsymbol{\beta}(\mathcal{A}^{2})$. We write this as $\boldsymbol{\beta}(\boldsymbol{a}) = \sum_{\tilde{\boldsymbol{a}} \in \mathcal{A}^{2}} w(\boldsymbol{a}, \tilde{\boldsymbol{a}}) \boldsymbol{\beta}(\tilde{\boldsymbol{a}})$. This allows us to express the IE as

$$\boldsymbol{\beta}(\alpha) = \sum_{\boldsymbol{a} \in A^1} p^1(\boldsymbol{a}) \underbrace{\sum_{\tilde{\boldsymbol{a}} \in \mathcal{A}^2} w(\boldsymbol{a}, \tilde{\boldsymbol{a}}) \boldsymbol{\beta}(\tilde{\boldsymbol{a}})}_{\boldsymbol{\beta}(\boldsymbol{a})} = \sum_{\tilde{\boldsymbol{a}} \in \mathcal{A}^2} \underbrace{\sum_{\boldsymbol{a} \in \mathcal{A}^1} p^1(\boldsymbol{a}) w(\boldsymbol{a}, \tilde{\boldsymbol{a}})}_{q(\tilde{\boldsymbol{a}})} \boldsymbol{\beta}(\tilde{\boldsymbol{a}})$$

Since the optimal choice in menu \mathcal{A}^2 is unique, so are the weights q by Theorem 1, implying that $q(\tilde{a}) = p^2(\tilde{a})$ for each $\tilde{a} \in \mathcal{A}^2$.

Since h is strictly convex, Jensen's inequality allows us to conclude that

$$\mathcal{I}(\boldsymbol{P}^{1},\boldsymbol{\pi}) = \sum_{i=1}^{I} \pi_{i} \sum_{\boldsymbol{a} \in A^{1}} p^{1}(\boldsymbol{a}) h\left(\sum_{\tilde{\boldsymbol{a}} \in A^{2}} w(\boldsymbol{a}, \tilde{\boldsymbol{a}}) \frac{\beta_{i}(\tilde{\boldsymbol{a}})}{\beta_{i}(\boldsymbol{\alpha})}\right)$$
$$\leq \sum_{i=1}^{I} \pi_{i} \sum_{\tilde{\boldsymbol{a}} \in A^{2}} \sum_{\boldsymbol{a} \in A^{1}} p^{1}(\boldsymbol{a}) w(\boldsymbol{a}, \tilde{\boldsymbol{a}}) h\left(\frac{\beta_{i}(\tilde{\boldsymbol{a}})}{\beta_{i}(\boldsymbol{\alpha})}\right) = \mathcal{I}(\boldsymbol{P}^{2}, \boldsymbol{\pi}).$$

In other words, P^1 is more informative regarding the state than P^2 .

Proof of Corollary 2: By Theorem 2, the optimal attention vector b^* can be written as a convex combination over points in

$$B = \{ \boldsymbol{\beta}(\boldsymbol{a}) \mid \boldsymbol{a} \in \mathcal{A}, \boldsymbol{\nabla}w(\boldsymbol{b}^*) \cdot \boldsymbol{\beta}(\boldsymbol{a}) = 1 \}.$$

Since *B* is contained in a I - 1-dimensional vector space, Carathéodory's Theorem implies that b^* can be written as a convex combination over at most *I* points. By Theorem 1, these weights correspond to the marginal implementation probabilities in a solution to (RI), and their support therefore describes a consideration set.

The next result establishes that it is without loss of generality to assume that a bounding hyperplane that binds somewhere in

$$\partial^+ \mathcal{B} = \{oldsymbol{b} \in \mathcal{B} \mid \nexists oldsymbol{b}' \in \mathcal{B} : oldsymbol{b}' > oldsymbol{b}\}$$

has only positive coordinates. We use it to restrict attention to priors with full support in Corollary 4 and Corollary 6.

Lemma 3. Let $B = Y + \mathbb{R}^{I}_{\leq 0}$ for a nonempty and convex set $Y \in \mathbb{R}^{I}$.

- (a) Whenever $\boldsymbol{\psi} \in \mathbb{R}^{I}$ satisfies $\boldsymbol{\psi} \cdot \boldsymbol{b} \leq 1 \ \forall \boldsymbol{b} \in B$, then $\boldsymbol{\psi} \geq \boldsymbol{0}$.
- (b) Suppose Y = conv.hull(X) for a nonempty and finite set $X \in \mathbb{R}_+^I$. For any $\mathbf{b}^0 \notin B$, $\exists \boldsymbol{\psi} \gg 0$ such that $\boldsymbol{\psi} \cdot \mathbf{b}^0 > 1$ and $\boldsymbol{\psi} \cdot \mathbf{b} \leq 1 \ \forall \mathbf{b} \in B$. For any $\mathbf{b}^0 \in \partial^+ B$, $\exists \boldsymbol{\psi} \gg 0$ such that $\boldsymbol{\psi} \cdot \mathbf{b}^0 = 1$ and $\boldsymbol{\psi} \cdot \mathbf{b} \leq 1 \ \forall \mathbf{b} \in B$.

Proof. By contradiction, assume $\psi_i < 0$ for some *i*. Then there exists a unit vector e^i and scalar t > 0 large enough such that $\psi^k \cdot (-te^i) = -t\psi_i^k > 1$ for *t* large enough, despite the fact that $-te^i \in B$ for all t > 0. This establishes part (a).

For part (b), the assumption that Y = conv.hull(X) for a nonempty and finite set $X \in \mathbb{R}^{I}_{+}$ implies *B* is a polyhedral set. It can then be written as the intersection of finitely many half spaces (see e.g. Ziegler [2012, Theorem 1.2])

$$B = \bigcap_{k=0}^{K} \left\{ \boldsymbol{y} \in \mathbb{R}^{I} \mid \boldsymbol{\psi}^{k} \cdot \boldsymbol{y} \leq 1 \right\}.$$
(4)

By part (a), $\psi^k \ge 0$ for all k.

Any $\boldsymbol{b}^0 \notin B$ violates $\boldsymbol{\psi}^k \cdot \boldsymbol{b}^0 \leq 1$ for at least one k; without loss of generality let $\boldsymbol{\psi}^0 \cdot \boldsymbol{b}^0 > 1$. By continuity, there exists $\varepsilon > 0$ small enough such that $\boldsymbol{\psi} \cdot \boldsymbol{b}^0 > 1$ for $\boldsymbol{\psi} := (1 - \varepsilon)\boldsymbol{\psi}^0 + \varepsilon \left(\frac{1}{K}\sum_{k=1}^{K}\boldsymbol{\psi}^k\right)$. However, note that $\boldsymbol{\psi} \gg \mathbf{0}$ since B is bounded above, and

$$\boldsymbol{\psi} \cdot \boldsymbol{b} = (1 - \varepsilon) \boldsymbol{\psi}^0 \cdot \boldsymbol{b} + \varepsilon \left(\frac{1}{K} \sum_{k=1}^K \boldsymbol{\psi}^k \cdot \boldsymbol{b} \right) \le (1 - \varepsilon) + \varepsilon = 1$$

for any $\boldsymbol{b} \in B$ by Equation (4).

For any $\mathbf{b}^0 \in \partial^+ B$, let $\bar{K} \subseteq \{\psi^0, ..., \psi^K\}$ denote the nonempty set of half-spaces that are binding at \mathbf{b}^0 . Clearly, $\boldsymbol{\psi} = |\bar{K}|^{-1} \sum_{k \in \bar{K}} \boldsymbol{\psi}^k \geq \mathbf{0}$ satisfies $\boldsymbol{\psi} \cdot \mathbf{b}^0 = 1$ and $\boldsymbol{\psi} \cdot \mathbf{b} \leq 1$ for all $\mathbf{b} \in B$. If $\psi_i = 0$, then there exists $\varepsilon > 0$ small enough such that $\mathbf{b}^0 + \varepsilon \mathbf{e}^i \in B$ by Equation (4), contradicting the initial assumption that $\mathbf{b}^0 \in \partial^+ B$. Hence, $\boldsymbol{\psi} \gg \mathbf{0}$.

Proof of Corollary 4: We prove each claim in turn.

(a) We show first that the upper boundary $\partial^+ \mathcal{B}$ is equal to the set of solutions to (GAP) across all $\pi \gg 0$. By the strict monotonicity of w it is immediate that if b^0 solves (GAP) for some $\pi \gg 0$, then $b^0 \in \partial^+ \mathcal{B}$. Conversely, any point $b^0 \in \partial^+ \mathcal{B}$ solves (GAP) for some $\pi \gg 0$. Indeed, by Lemma 3, there exists a vector $\psi \gg 0$ such that $\psi \cdot b \leq 1$ for all $b \in \mathcal{B}$ and $\psi \cdot b^0 = 1$. The vector $\pi = \psi \otimes b^0$ denotes a valid prior since $\pi \gg 0$ and $\sum_{i=1}^{I} \pi_i = \psi \cdot b^0 = 1$. By construction, $\psi = \nabla w(b^0)$ under π , and thus b^0 solves the corresponding (GAP) problem by Theorem 2(a).

By Corollary 1, it follows that $\{\alpha^{\pi} \mid \pi \gg 0\}$ is equal to $\beta^{-1}(\partial^{+}\beta)$, and the Minkowski sum with $\mathbb{R}^{I}_{\leq 0}$ adds all weakly dominated payoff vectors to complete Definition 2.

(b) If $a \in A$, then there exists $b^0 \in \partial^+ \mathcal{B}$ such that $\mathcal{\beta}(a) \leq b^0 \in \partial^+ \mathcal{B}$, and thus by (a) it follows $\mathcal{A} \subseteq \overline{\mathcal{A}}$.

For the second part, let $\bar{\mathcal{B}} = \beta(\bar{\mathcal{A}})$. By (a) and monotonicity of β , $\partial^+ \bar{\mathcal{B}} = \partial^+ \mathcal{B}$ and trivially the solution to (GAP) is the same for \mathcal{A} and $\bar{\mathcal{A}}$. The result follows from the uniqueness of the IE by Corollary 1.

- (c) By continuity and statewise monotonicity of β , \overline{A} is closed and bounded above. It is strictly convex since β is convex and β^{-1} is statewise strictly concave.
- (d) Let $a \in \overline{A}$. For any $\pi \gg 0$, the agent is indifferent between \overline{A} and $\{\alpha^{\pi}\}$ by part (b). Since blind implementation of a requires zero information cost, the expected utility $\pi \cdot a$ can be no larger than that of α^{π} . Hence a belongs to $A := \bigcap_{\pi \gg 0} \{a \in \mathbb{R}^I \mid \pi \cdot (a^{\pi} a) \ge 0\}.$

Conversely, suppose $a \notin \overline{A}$. Since \overline{A} is closed and convex by part (c), the separating hyperplane theorem implies that there exists $\psi \in \mathbb{R}^{I}$ such that $\psi \cdot a > 1$ and $\psi \cdot a' \leq 1$ for all $a' \in \overline{A}$. By Lemma 3(a), this implies in particular that $\psi \geq 0$. Letting $\rho = \frac{1}{I}\mathbf{1}$ denote the uniform prior, note that by continuity, there exists $\varepsilon > 0$ small enough such that

$$ilde{\psi} = rac{oldsymbol{\psi} + arepsilon oldsymbol{
ho}}{1 + arepsilon (oldsymbol{
ho} \cdot oldsymbol{lpha}^{
ho})} \gg \mathbf{0}$$

also satisfies $\tilde{\psi} \cdot \boldsymbol{a} > 1$. The vector $\boldsymbol{\pi} = \tilde{\psi}/(\tilde{\psi} \cdot \mathbf{1})$ thus forms a valid prior $(\boldsymbol{\pi} \cdot \mathbf{1} = 1)$ with full support $(\boldsymbol{\pi} \gg \mathbf{0})$. Moreover, $\boldsymbol{\rho} \cdot \boldsymbol{a}' \leq \boldsymbol{\rho} \cdot \boldsymbol{\alpha}^{\boldsymbol{\rho}}$ for all $\boldsymbol{a}' \in \bar{\mathcal{A}}$ since $\bar{\mathcal{A}}$ and \mathcal{A} share the IE $\boldsymbol{\alpha}^{\boldsymbol{\rho}}$. Together, this implies that

$$ilde{m{\psi}}\cdotm{a}' = rac{m{\psi}\cdotm{a}'+arepsilon(m{
ho}\cdotm{a}')}{1+arepsilon(m{
ho}\cdotm{lpha}^{m{
ho}})} \leq rac{1+arepsilon(m{
ho}\cdotm{lpha}^{m{
ho}})}{1+arepsilon(m{
ho}\cdotm{lpha}^{m{
ho}})} = 1 \qquad orallm{a}'\inar{m{A}}_{m{P}}$$

and thus $\boldsymbol{\pi} \cdot (\boldsymbol{\alpha}^{\boldsymbol{\pi}} - \boldsymbol{a}) < 0$, or $\boldsymbol{a} \notin A$.

(e) The comparative statics follow from those of the IE in Corollary 1. Indeed, by Corollary 1, α^{π} weakly increases coordinate-wise for each prior π when new actions are added or information costs decrease. Thus, the same is true for the halfspaces defined in part (d), as well as their intersection \overline{A} .

Proof of Corollary 6: Property (a) is a direct consequence of Theorem 1(a). An additional action a^+ alters the solution to (GAP) if and only if

$$\nabla w(\boldsymbol{\beta}(\boldsymbol{\alpha})) \cdot \boldsymbol{\beta}(\boldsymbol{a}^+) > 1,$$

referring to \mathcal{A} only through its IE α .

For property (b), suppose $a^+ \in \overline{\mathcal{A}}$. By Corollary 4(a) and statewise monotonicity of β , there exists $b^0 \in \partial^+ \mathcal{B}$ such that $b^0 \geq \beta(a^+)$. Since \mathcal{B} is weakly increasing in the addition of new actions, b^0 remains a feasible attention vector under any larger menu $\mathcal{A}' \supseteq \mathcal{A}$. Moreover, any feasible attention vector b under menu $\mathcal{A}' \cup \{a^+\}$ can be written as $b = p(a^+)\beta(a^+) + \sum_{a \in \mathcal{A}'} p(a)\beta(a)$ for some $p \in \mathcal{C}(\mathcal{A}' \cup \{a^+\})$. Note that

$$\boldsymbol{b} \leq \boldsymbol{b}^1 := p(\boldsymbol{a}^+)\boldsymbol{b}^0 + \sum_{\boldsymbol{a}\in\mathcal{A}'} p(\boldsymbol{a})\boldsymbol{\beta}(\boldsymbol{a}),$$

where b^1 is feasible under menu \mathcal{A}' . By monotonicity of $w, w(b^1) \ge w(b)$, and so the addition of a^+ does not increase the optimal (GAP) objective value. By Theorem 1, the same is true for (RI).

Conversely, suppose $\mathbf{a}^+ \notin \bar{\mathcal{A}}$. By Lemma 3, there exists $\psi \gg \mathbf{0}$ such that $\psi \cdot \beta(\mathbf{a}^+) > 1$ and $\psi \cdot \beta(\mathbf{a}) \leq 1$ for all $\mathbf{a} \in \bar{\mathcal{A}}$. Let $\mathbf{b}^1 \in \mathbb{R}^I_+$ be defined component-wise as $b_i^1 := (\psi \cdot \beta(\mathbf{a}^+)) \frac{\pi_i}{\psi_i}$. Since $\mathbf{b}^1 \gg \mathbf{0}$, there exists $\varepsilon > 0$ small enough such that $\mathbf{b}^2 := (1 + \varepsilon)\mathbf{b}^1 - \varepsilon\beta(\mathbf{a}^+)) \gg \mathbf{0}$. Labeling the pre-image $\mathbf{a}^2 := \beta^{-1}(\mathbf{b}^2)$ and setting $\mathcal{A}' := \mathcal{A} \cup \{\mathbf{a}^2\}$, we complete the proof by showing that $W(\mathcal{A}' \cup \{\mathbf{a}^+\}) > W(\mathcal{A}')$. By construction, $b^1 \in \mathcal{B}^+ := \text{conv.hull}(\mathcal{A}(\mathcal{A}' \cup \{a^+\}))$ and b^1 satisfies the optimality conditions from Theorem 2(a) over \mathcal{B}^+ since

$$\nabla w(\boldsymbol{b}^{1}) \cdot \boldsymbol{\beta}(\boldsymbol{a}) = \frac{\boldsymbol{\psi} \cdot \boldsymbol{\beta}(\boldsymbol{a})}{\boldsymbol{\psi} \cdot \boldsymbol{\beta}(\boldsymbol{a}^{+})} < 1 \qquad \forall \boldsymbol{a} \in \mathcal{A},$$
(5)

$$\nabla w(\boldsymbol{b}^1) \cdot \boldsymbol{\beta}(\boldsymbol{a}^+) = \frac{\boldsymbol{\psi} \cdot \boldsymbol{\beta}(\boldsymbol{a}^+)}{\boldsymbol{\psi} \cdot \boldsymbol{\beta}(\boldsymbol{a}^+)} = 1$$
(6)

$$\boldsymbol{\nabla}w(\boldsymbol{b}^{1}) \cdot \boldsymbol{\beta}(\boldsymbol{a}^{2}) = (1+\varepsilon) \underbrace{\boldsymbol{\nabla}w(\boldsymbol{b}^{1}) \cdot \boldsymbol{b}^{1}}_{=\sum_{i=1}^{I} \frac{\pi_{i}}{b_{i}^{1}} b_{i}^{1} = 1} -\varepsilon \underbrace{\boldsymbol{\nabla}w(\boldsymbol{b}^{1}) \cdot \boldsymbol{\beta}(\boldsymbol{a}^{+})}_{=1 \text{ by } (6)} = 1.$$
(7)

By Theorem 1, this implies that $W(\mathcal{A}' \cup \{a^+\}) = \lambda w(b^1)$. The strict inequality in (5) implies that $\nabla w(b^1) \cdot b < 1$ for any $b \in \mathcal{B}' \setminus \{b^2\}$, ruling out $b^1 \in \mathcal{B}'$ since $\nabla w(b^1) \cdot b^1 = \sum_{i=1}^{I} \frac{\pi_i}{b_i^1} b_i^1 = 1$ by definition of w. Since the solution to (GAP) is unique, it follows that $w(b) < w(b^1)$ for all $b \in \mathcal{B}'$, and hence $W(\mathcal{A}') < W(\mathcal{A}' \cup \{a^+\})$. \Box

B Derivation of Practical Implications

Proof of Lemma 1: Consider any convergent subsequence $P^{n_k} \to \bar{P}$. Since P^n converges to P^* under d_{IE} , we know that $\beta(\alpha^{\bar{P}}) = \beta(\alpha^{P^*})$. Since the solution to (RI) is unique, this point can be written in a unique way as a convex combination over $\beta(\mathcal{A})$.³³ This implies that the marginals $p^{n_k} = \sum_{i=1}^{I} \pi_i P_i^k$ converge to $p^* = \sum_{i=1}^{I} \pi_i P_i^*$. If all convergent subsequences of the bounded sequence p^n converge to the same limit p^* , the Bolzano-Weierstrass theorem implies that p^n itself converges to p^* . By continuity of Equation (FONC), the convergence translates to the conditional choice $Q^n \to P^*$.

Algorithm Our base routine for small to moderate menus works as follows: We make use of the scaling property (Corollary 3) to avoid floating point imprecision, and store the normalized attention vectors in a *I*-by- $|\mathcal{A}|$ matrix with entries $B_{ia} = \beta_i(\mathbf{a}) / \max_{\tilde{\mathbf{a}} \in \mathcal{A}} \beta_i(\tilde{\mathbf{a}})$. Starting with an initial guess \mathbf{p}^{0} ,³⁴ we iteratively solve a second-

 $^{^{33}}$ See the accompanying discussion on Page 13.

³⁴Practically, we use the full-information marginals by placing weight π_i on $\arg \max_{a \in \mathcal{A}} a_i$.

order Taylor approximation to (GAP), which after dropping constant terms yields

$$\boldsymbol{q}^{k} := \arg \max_{\boldsymbol{p} \in \Delta^{|\mathcal{A}|-1}} \frac{1}{2} \boldsymbol{p}^{\mathsf{T}} B^{\mathsf{T}} H B \boldsymbol{p} - 2 \boldsymbol{\nabla} w (B \boldsymbol{p}^{k})^{\mathsf{T}} B,$$

where *H* refers to the diagonal matrice with entries $H_{ii} = \pi_i / (B\boldsymbol{p}^k)_i^2$.³⁵ If $\nabla w(B\boldsymbol{q}^k) \cdot \boldsymbol{p}^k \leq 1$, we set $\boldsymbol{p}^{k+1} := \boldsymbol{q}^k$ and move to the next iteration. Otherwise, we find the optimal marginals \boldsymbol{p}^{k+1} along the segment $\{t\boldsymbol{p}^k + (1-t)\boldsymbol{q}^k \mid t \in [0,1]\}$ by identifying the root of the monotone function $\nabla w(t\boldsymbol{p}^k + (1-t)\boldsymbol{q}^k) \cdot (\boldsymbol{q}^k - \boldsymbol{p}^k)$.

We repeat this quadratic approximation until the implied IE converges, i.e. until $d_{\mathsf{IE}}(\mathbf{P}^k, \mathbf{P}^{k+1}) < \varepsilon$, where \mathbf{P}^k is defined from \mathbf{p}^k according to (FONC). As default, and in all our applications, we use tolerance parameter $\varepsilon = 10^{-12}$. By construction, our approach ensures that the objective value $w(B\mathbf{p}^k)$ increases with each iteration.

When the action space is rich, the attention matrix B can require a lot of memory. To avoid this limitation, we first apply the base routine to a coarse subgrid of the menu, $A^0 \subset \mathcal{A}$. Upon convergence, we denote the estimated marginals by q^0 and its associated attention vector as $b^0 = B^0 q^0$. At each iteration m, we compute the b^{m-1} -scores over a finer subgrid $A^m \supseteq A^{m-1}$. We add the actions with the highest score until the menu reaches some maximum size K or contains all actions in some p-cover of menu A^m . As long as K is large enough, the p-cover eventually stabilizes, and we move to the next finer subgrid. We continue this process until $A^m = \mathcal{A}$ and the p-cover stabilizes.

References

- Andre Caplin and Mark Dean. Behavioral implications of rational inattention with shannon entropy. *NBER WP*, (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.

³⁵We implement this code using MATLAB's built-in quadprog solver (Version 2019b). Since the solver does not accept an initial guess, we use an equivalent centered problem by solving for $dp = p^{k+1} - p^k$ instead.

- 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. *The Review of Economic Studies*, page rdy037, 2018b. doi: 10.1093/restud/rdy037. URL http://dx.doi.org/10.1093/restud/rdy037.
- Thomas M Cover and Joy A Thomas. *Elements of information theory*. John Wiley & Sons, 2012.
- Kunal Dasgupta and Jordi Mondria. Inattentive importers. Journal of International Economics, 112(C):150-165, 2018. doi: 10.1016/j.jinteco.2018.03. URL https: //ideas.repec.org/a/eee/inecon/v112y2018icp150-165.html.
- Gerard Debreu. Review of rd luce, individual choice behavior: A theoretical analysis. American Economic Review, 50(1):186–188, 1960.
- H. G. Eggleston. Convexity. Cambridge Tracts in Mathematics. Cambridge University Press, 1958. doi: 10.1017/CBO9780511566172.
- Xavier Gabaix. A sparsity-based model of bounded rationality. Quarterly Journal of Economics, 129(4):1661–1710, 2014.
- Wagner Piazza Gaglianone, Raffaella Giacomini, Joao Issler, and Vasiliki Skreta. Incentive-driven inattention. 2019. CEPR Discussion Paper No. DP13619. Available at SSRN: https://ssrn.com/abstract=3363532.
- Matthew Gentzkow and Emir Kamenica. Costly persuasion. The American Economic Review, 104(5):457–462, 2014.
- Lixin Huang and Hong Liu. Rational inattention and portfolio selection. The Journal of Finance, 62(4):1999-2040, 2007. ISSN 00221082, 15406261. URL http://www. jstor.org/stable/4622323.
- Kenneth Judd. Numerical Methods in Economics, volume 1. The MIT Press, 1 edition, 1998. URL https://EconPapers.repec.org/RePEc:mtp:titles:0262100711.

- Junehyuk Jung, Jeong Ho (John) Kim, Filip Matějka, and Christopher A Sims. Discrete Actions in Information-Constrained Decision Problems. *The Review of Economic Studies*, 03 2019. ISSN 0034-6527. doi: 10.1093/restud/rdz011. URL https://doi.org/10.1093/restud/rdz011. rdz011.
- Marcin Kacperczyk, Stijn Van Nieuwerburgh, and Laura Veldkamp. A rational theory of mutual funds' attention allocation. *Econometrica*, 84(2):571–626, 2016.
- Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. The American Economic Review, 101(6):2590–2615, 2011.
- Yulei Luo, Jun Nie, Gaowang Wang, and Eric R. Young. Rational inattention and the dynamics of consumption and wealth in general equilibrium. *Journal of Economic Theory*, 172:55 – 87, 2017. ISSN 0022-0531. doi: https://doi.org/10.1016/ j.jet.2017.08.005. URL http://www.sciencedirect.com/science/article/pii/ S0022053117300832.
- Bartosz Mackowiak, Adam Filip Matejka, and Mirko Wiederholt. Survey: Rational inattention, a disciplined behavioral model. Working paper, CEPR, October 2018.
- Filip Matějka. Rationally inattentive seller: Sales and discrete pricing. *The Review* of *Economic Studies*, 83(3):1125–1155, 2016.
- 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.
- Bartosz Maćkowiak and Mirko Wiederholt. Optimal sticky prices under rational inattention. The American Economic Review, 99(3):769-803, 2009. ISSN 00028282, 19447981. URL http://www.jstor.org/stable/25592482.
- Bartosz Maćkowiak, Filip Matějka, and Mirko Wiederholt. Dynamic rational inattention: Analytical results. Journal of Economic Theory, 176:650 - 692, 2018. ISSN 0022-0531. doi: https://doi.org/10.1016/j.jet.2018.05.001. URL http: //www.sciencedirect.com/science/article/pii/S002205311830139X.
- Jianjun Miao, Jieran Wu, and Eric Young. Multivariate rational inattention. Working paper, Boston University, January 2019.

- Jordi Mondria. Portfolio choice, attention allocation, and price comovement. *Journal* of Economic Theory, 145(5):1837–1864, 2010.
- Lin Peng. Learning with information capacity constraints. *The Journal of Financial* and Quantitative Analysis, 40(2):307–329, 2005. ISSN 00221090, 17566916. URL http://www.jstor.org/stable/27647199.
- Lin Peng and Wei Xiong. Investor attention, overconfidence and category learning. Journal of Financial Economics, 80(3):563 - 602, 2006. ISSN 0304-405X. doi: https://doi.org/10.1016/j.jfineco.2005.05.003. URL http://www.sciencedirect. com/science/article/pii/S0304405X05002138.
- Christopher A Sims. Implications of rational inattention. Journal of monetary Economics, 50(3):665–690, 2003.
- Christopher A Sims. Rational inattention: Beyond the linear-quadratic case. American Economic Review, 96(2):158–163, 2006.
- 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.
- Stijn Van Nieuwerburgh and Laura Veldkamp. Information immobility and the home bias puzzle. The Journal of Finance, 64(3):1187-1215, 2009. doi: 10.1111/j. 1540-6261.2009.01462.x. URL https://onlinelibrary.wiley.com/doi/abs/10. 1111/j.1540-6261.2009.01462.x.
- Stijn Van Nieuwerburgh and Laura Veldkamp. Information acquisition and underdiversification. The Review of Economic Studies, 77(2):779–805, 2010.
- Günter M Ziegler. Lectures on polytopes, volume 152. Springer Science & Business Media, 2012.

C Online Appendix

C.1 Sticky Prices [Matějka, 2016]

Additional Figures. Both the GAP-SQP and BA algorithms replicate the results in Matějka [2016] very closely. Figure 9 shows the marginal distributions over prices for the GAP-SQP algorithm (panel (a)) and BA algorithm (panel(b)), together with the numerical solutions from AMPL provided by Filip Matějka. Solutions are so close that we had to offset the histograms for visibility. We find that increasing grid precision for actions does not meaningfully alter the solution.



Figure 9: Replication of Matějka [2016]

Figure 10 reports the differences in the objective function value, at the computed maximum, between the GAP-SQP and BA algorithms, for the benchmark case. The difference is positively thorough all the information values, indicating that the GAP-SQL algorithm achieves greater precision despite running on a fraction of the time of the BA algorithm. The difference, though, is very small by our choice of stopping values.

C.2 Portfolio Choice [Jung et al., 2019]

Derivation of the LQG solution. Because of the properties of the CARA utility function, it is possible to rewrite Equation (1) as

$$U(\theta, \mathbf{Y}) = -\exp\left(-\alpha \left(1.03 + \sum_{j=1}^{2} (0.01 + Y_j)\theta_j\right) + \frac{\alpha^2}{2} (\theta_1^2 + \theta_2^2)\sigma_z^2\right).$$



Figure 10: Objective function: GAP-SQP minus Blahut-Arimoto algorithm

We now construct a second-order approximation of this objective function around $\boldsymbol{\theta}$ such that

$$\nabla_{\boldsymbol{\theta}} U(\hat{\boldsymbol{\theta}}, \mathbf{0}) = \mathbf{0} \qquad \Longleftrightarrow \qquad \hat{\boldsymbol{\theta}} \approx (33.4124, 33.4124)$$

i.e., those portfolio shares that would be optimal if evaluated at the ex-post realization Y = 0. Note this is not the same as the no-information solution because it does not take into account the risk associated with Y.

Because $\mathsf{E}[\mathbf{Y}] = \mathbf{0}$, the second-order Taylor approximation around $(\hat{\theta}, \mathsf{E}[\mathbf{Y}])$ equals

$$\tilde{U}(\boldsymbol{\theta}, \boldsymbol{Y}) = U(\hat{\boldsymbol{\theta}}, \boldsymbol{0}) + \begin{pmatrix} \boldsymbol{\theta} - \hat{\boldsymbol{\theta}} \\ \boldsymbol{Y} - \boldsymbol{0} \end{pmatrix}^{\mathsf{T}} \boldsymbol{\nabla} U(\hat{\boldsymbol{\theta}}, \boldsymbol{0}) + \frac{1}{2} \begin{pmatrix} \boldsymbol{\theta} - \hat{\boldsymbol{\theta}} \\ \boldsymbol{Y} - \boldsymbol{0} \end{pmatrix}^{\mathsf{T}} \boldsymbol{\nabla}^{2} U(\hat{\boldsymbol{\theta}}, \boldsymbol{0}) \begin{pmatrix} \boldsymbol{\theta} - \hat{\boldsymbol{\theta}} \\ \boldsymbol{Y} - \boldsymbol{0} \end{pmatrix}$$

The LQG approximation seeks to design a random variable, θ , to maximize

$$\max_{\boldsymbol{\theta}} E_{\boldsymbol{\theta},\boldsymbol{Y}} \left[\tilde{U}(\boldsymbol{\theta},\boldsymbol{Y}) \right] - \lambda \mathcal{I}(\boldsymbol{\theta};\boldsymbol{Y})$$

Cover and Thomas [2012] document a well-known solution to this problem: We simply set set $\boldsymbol{\theta}$ to be jointly normal with \boldsymbol{Y} . This follows from the fact that a Gaussian distribution maximizes entropy for a fixed variance. It is not hard to see that, given the choice of approximating point, the optimal mean is just $\hat{\boldsymbol{\theta}}$. Given

this, we need only solve for covariance matrix, Σ , that optimally balances smaller conditional dispersion against information costs.

We can simplify the objective dropping the linear terms, since they do not depend on the covariance matrix. We can then simplify further using a couple of well-known facts: First is the functional form for the mutual information of a pair of multivariate normals, which has a simple closed-form expression; second is for any random variables X_1 and X_2 ,

$$\mathsf{E}[\boldsymbol{X}_1^\mathsf{T} A \boldsymbol{X}_2] = \mathsf{tr}(A\mathsf{Cov}(\boldsymbol{X}_1, \boldsymbol{X}_2)) + \mathsf{E}[\boldsymbol{X}_1]^\mathsf{T} A \mathsf{E}[\boldsymbol{X}_2],$$

Plugging these in implies that solving the RI problem is tantamount to selecting a positive-definite Σ that is consistent with the marginal distribution over Y so as to maximize

$$\frac{1}{2} \mathsf{tr}\left(\left[\boldsymbol{\nabla}^2 U(\hat{\boldsymbol{\theta}}, \mathbf{0})\right] \boldsymbol{\Sigma}\right) - \lambda \frac{1}{2} \log\left(\frac{|\boldsymbol{\Sigma}_{\boldsymbol{\theta}}| \times |\boldsymbol{\Sigma}_{\boldsymbol{Y}}|}{|\boldsymbol{\Sigma}|}\right)$$

where $|\cdot|$ denotes the matrix determinant and Σ_X denotes the marginal covariance of the X.

Plugging in the optimal covariance matrix

$$\Sigma = \begin{bmatrix} \Sigma_{\theta} & \Sigma_{\theta Y} \\ \Sigma_{Y\theta} & \Sigma_{Y} \end{bmatrix} = \begin{bmatrix} 3158.4 & 0 & 0.9453 & 0 \\ 0 & 3158.4 & 0 & 0.9453 \\ 0.9453 & 0 & 0.0004 & 0 \\ 0 & 0.9453 & 0 & 0.0004 \end{bmatrix}$$

yields the distribution found in Figure 6(a). The objective function net of information costs is derived using Monte-Carlo methods. We take 10 million draws from the optimal distribution and compute the sample average utility. We repeat this 100 times and take sample statistics of the estimates. This yields an average payoff, net of information costs, of -.3220 with a 95% confidence band of [-.3221, -.3219].

Additional Figures. For comparison purposes, Figure 11 plots the statewise payoff distribution $U(\boldsymbol{\theta}, \boldsymbol{Y}) - \lambda \mathcal{I}$, assuming $(\boldsymbol{\theta}, \boldsymbol{Y})$ is distributed according to the numeric solution of GAP-SQP (blue) or JKMS (orange), and the information cost \mathcal{I} is borne unconditionally.



Figure 11: Payoff distribution across algorithm estimates, smoothed with a kernel density estimate.