

Information Acquisition, Efficiency, and Non-Fundamental Volatility*

Benjamin Hébert[†]

Jennifer La'O[‡]

August 14, 2022

Abstract

We analyze non-fundamental volatility and efficiency in large games featuring strategic interaction and endogenous information acquisition. We adopt the rational inattention approach to information acquisition but generalize to a large class of information costs. Agents can learn about exogenous states and endogenous aggregate actions. We provide sufficient conditions on information costs that guarantee the existence of efficient or inefficient equilibria, and another set of sufficient conditions that guarantee the existence or non-existence of equilibria with zero non-fundamental volatility. Mutual information, the cost typically assumed in rational inattention models, guarantees the existence of an efficient equilibrium with zero non-fundamental volatility.

JEL codes: C72, D62, D83

Keywords: Information Acquisition, Large Games, Beauty Contests, Rational Inattention

*We thank John Leahy, Andrew Caplin, and Alessandro Pavan for their excellent discussions and feedback on our paper. We also thank Marios Angeletos, Navin Kartik, Elliot Lipnowski, and Mike Woodford for their insightful advice and comments. We would like to thank Harald Uhlig (editor) and four anonymous referees for comments and suggestions that helped improve the paper. All remaining errors are our own.

[†]Stanford University and NBER, bhebert@stanford.edu.

[‡]Columbia University and NBER, jenlao@columbia.edu.

1 Introduction

In many economic environments, agents make choices under incomplete information and have incentives to align their actions with both economic “fundamentals” and the actions of other agents. These games of strategic interaction form the basis underlying many micro-founded macroeconomic environments; examples include firms’ nominal price-setting decisions in New Keynesian models [Woodford, 2003], firms’ real quantity choices in business cycle models [Angeletos and La’O, 2010, 2013], as well as investors’ asset positions in models of financial trade [Grossman and Stiglitz, 1976, 1980].

In these environments, agents’ beliefs over exogenous fundamentals and the endogenous actions of others play a key role in determining equilibrium outcomes. But where do these beliefs come from and how are they formed?

In this paper we investigate the endogenous acquisition of information within games of strategic interaction. We ask two questions. First, what properties of the agents’ information acquisition costs guarantee that an equilibrium is or is not constrained efficient? Second, what properties of the agents’ information acquisition costs guarantee that an equilibrium of the game features or does not feature non-fundamental volatility?

Our framework. We study a class of large games of strategic interaction similar to those analyzed in the beauty contest literature [Morris and Shin, 2002, Angeletos and Pavan, 2007]. A continuum of ex-ante identical agents take actions under incomplete information. Each agent has an incentive to align her action with an exogenous, payoff-relevant state (the “fundamental”) and with the endogenous mean action. We restrict attention to the class of utility functions in which payoff externalities are absent in equilibrium—which we call the “mean-critical” utility functions—to focus on externalities related to information acquisition.

We allow agents to acquire information about the fundamental, the endogenous mean action of other agents, and the realizations of exogenous public signals. These public signals represent news and other forms of publicly-available information that agents can observe and condition their actions upon. We model the exogenous public signals as functions mapping the payoff-relevant states and payoff-irrelevant states, or “noise,” into signal realizations. By allowing for payoff-irrelevant noise in public signals, we introduce a potential source of “non-fundamental” volatility in equilibrium. Agents can also learn by observing the endogenous mean action; this captures agents’ ability to learn from public information that endogenously aggregates the actions of other agents, such as prices or public statistics of aggregate economic activity.

To model information acquisition, we adopt the rational inattention approach to costly information acquisition proposed by Sims [2003]. However, unlike the standard rational inattention framework, we do not assume that information costs are proportional to mutual infor-

mation. We instead consider a more general class of cost functions: those that are “posterior-separable” in the terminology of [Caplin, Dean, and Leahy \[2022\]](#). Posterior-separable cost functions can be described as the expected divergence from the agents’ prior to posterior beliefs. This class nests the standard mutual information cost function as a special case, but also includes many other cost functions, including the log-likelihood ratio (LLR) cost functions of [Pomatto et al. \[2020\]](#), the Tsallis entropy costs of [Caplin, Dean, and Leahy \[2022\]](#), and the neighborhood-based cost functions of [Hébert and Woodford \[2021\]](#).¹

Within this general class of information costs, we ask: what properties of information costs guarantee that equilibria are inefficient? What properties of information costs ensure that equilibria exhibit non-fundamental volatility? And finally, are these properties related? That is, is non-fundamental volatility synonymous with inefficiency?

Partial Monotonicity and Partial Invariance. We find that whether or not equilibria in our context exhibit non-fundamental volatility or inefficiency is related to two key properties of cost functions: “partial monotonicity” and “partial invariance.”

We introduce and define partial monotonicity and partial invariance as properties of the divergences associated with posterior-separable cost functions. Loosely speaking, a divergence can be thought of as a measure of the “distance” between the prior and posterior. Partial monotonicity and partial invariance describe how this divergence responds to different transformations of the prior and posterior.

Suppose an agent is uncertain about a multi-dimensional aggregate state and receives a signal that moves her posterior beliefs “away from” her prior in one dimension of that state. This signal is, in a sense, more informative than another signal that leaves posterior beliefs close to the prior in that dimension. This idea leads to a notion of monotonicity: a divergence is monotonic in some dimension if the divergence decreases as we make the posterior more like the prior in that dimension. But note that a divergence could be monotonic in some dimensions but not in others—our definition of partial monotonicity allows for this flexibility.

Take for example, a two-dimensional state space, $s \in S$ and $r \in R$. We will say that a divergence is *R-monotone* if the divergence decreases when we replace the posterior’s conditional distribution of r given s with the prior’s conditional distribution of r given s . In contrast, we will say that a cost function is *nowhere-R-monotone* when this property fails essentially everywhere.

We define a separate concept, “partial invariance.” Take again a two-dimensional state space $s \in S$ and $r \in R$. We will say that a divergence is *R-invariant* if, for any prior and posterior that share the same conditional distributions of r given s , only their marginal distributions on s matter for the divergence. In contrast, we will say that a cost function is *nowhere-R-invariant* when this property fails essentially everywhere.

¹These alternative cost functions have been proposed in part because they are better able to match observed behavior in experiments (see, e.g., [Dean and Neligh \[2019\]](#)).

The forms of monotonicity and invariance, and nowhere-monotonicity and nowhere-invariance, that we introduce are generalizations of the invariance concept described in the literature on information geometry [Chentsov, 1982, Amari and Nagaoka, 2007] and employed in various economic applications (e.g. Hébert [2018], Caplin, Dean, and Leahy [2022]). The invariance and monotonicity properties described by these authors indicate whether a divergence is invariant and monotone with respect to *all possible* dimensions of the state space. In contrast, our generalization allows for invariance or monotonicity of a divergence with respect to *specific* dimensions of the state space. This flexibility turns out to be important: we find that the equilibrium properties of interest in our game—namely, non-fundamental volatility and constrained efficiency—depend on the partial monotonicity and partial invariance properties with respect to specific dimensions of the state space.

Results. We find that a local form of partial invariance of the cost function with respect to the endogenous mean action is sufficient and necessary to ensure that a solution to the social planner’s problem is an equilibrium.² As a result, when information costs are nowhere-invariant in the aggregate action, all equilibria are (away from certain exceptions) inefficient. On the other hand, when information costs are invariant in the aggregate action, an equilibrium exists that is constrained efficient. If, in addition, actions are strategic substitutes, then all equilibria are efficient.

When information costs are nowhere-invariant in the aggregate action, the sensitivity of the aggregate action to the underlying state affects the ease by which agents acquire information. Consider, as an example, a scenario in which agents learn in part by paying attention to the aggregate action, and more extreme aggregate actions are easier to observe than less extreme ones. If the aggregate action becomes more sensitive to the state, it becomes less costly for agents to learn by observing the aggregate action, which reduces their information acquisition costs. In this context, inefficiency arises due to an informational externality—agents do not internalize how their own strategies affect the sensitivity of the aggregate action to the state. The planner would instead encourage more extreme actions in order to increase the sensitivity of the aggregate action to the state, and in so doing, reduce information costs.

On the other hand, when costs are invariant in the aggregate action, the ease with which agents acquire information is unaffected by how the aggregate action reacts to the underlying state. In this case, an equilibrium exists in which there are no informational externalities: the incentives for agents to acquire and use information coincide with that of the planner.

We find that partial monotonicity of the cost function in noisy public signals is related to whether or not equilibria exhibit non-fundamental volatility. When information costs are nowhere-monotone in noisy public signals, all equilibria in which information is acquired exhibit non-

²This necessity result applies only when beliefs lie in the interior of the simplex.

fundamental volatility. In contrast, when information costs are monotone in noisy public signals, an equilibrium exists that features zero non-fundamental volatility. We also provide additional conditions under which all equilibria exhibit zero non-fundamental volatility.

When information costs are nowhere-monotone in noisy public signals, agents find it relatively cheaper (with respect to information costs) to condition their actions on the noisy public signals than to observe the payoff-relevant states directly. This information cost-saving incentive is what induces agents to condition their actions on public signals, introducing correlated noise in actions. As a result, all equilibria feature non-fundamental volatility. When instead information costs are monotone in noisy public signals, it is always weakly less costly for agents to condition their actions on the payoff-relevant states and aggregate actions alone. As a result, an equilibrium exists in which all agents ignore such signals and there is zero non-fundamental volatility.

Our results clarify that separate properties of information costs determine whether or not equilibria exhibit non-fundamental volatility and whether or not equilibria are efficient. By precisely defining these properties—namely, partial monotonicity and partial invariance—we characterize the relationship between information costs and these equilibrium properties.

Throughout the paper, we consider three example cost functions: mutual information (the standard rational inattention cost function) and two forms of the neighborhood-based cost functions proposed by [Hébert and Woodford \[2021\]](#). The divergence associated with one of the neighborhood-based cost functions we consider is nowhere-monotone in the noisy public signals, and the divergence associated with the other is nowhere-invariant in the endogenous aggregate action. As a result, we show that these cost functions can lead, respectively, to non-fundamental volatility and constrained inefficiency. In contrast, the divergence associated with mutual information (the Kullback-Leibler divergence) is monotone and invariant in all dimensions; consequently, with this cost function, there is always a constrained-efficient equilibrium with zero non-fundamental volatility.

The bulk of our analysis assumes a finite set of exogenous states. We conclude our analysis by presenting examples in the canonical linear-quadratic-Gaussian setting; these examples illustrate the general lessons presented in the rest of this paper using first-order conditions in a familiar environment. We extend some of our results to a continuous state space in the technical appendix, Section [C](#).

Related Literature. Building on [Morris and Shin \[2002\]](#) and [Angeletos and Pavan \[2007\]](#), a large literature has studied the positive and normative implications of large games of strategic interaction and incomplete information, and applied these insights to questions in macro, finance, and industrial organization (see [Angeletos and Lian \[2016\]](#) for a recent survey). Much of this literature assumes linear-quadratic payoffs, Gaussian priors over states, and exogenously specified Gaussian signals. Noisy public signals, in particular, are standard components of ex-

ogenous information sets in these games [Bergemann and Morris, 2013] and play a major role in determining equilibrium outcomes. The “noise” or “errors” in these signals are orthogonal to fundamentals and manifest as “non-fundamental” volatility in equilibrium.

Several authors (e.g. Hellwig and Veldkamp [2009], Myatt and Wallace [2012], Colombo, Femminis, and Pavan [2014], Pavan [2016]) endogenize information acquisition in the linear-quadratic setting, allowing agents to choose the precision with which they observe an exogenously specified set of Gaussian signals about exogenous states. In these papers, the presence or absence of non-fundamental volatility depends on the assumed correlation structure of the exogenously given signals (and, with precision choice across multiple signals, agents’ incentives to coordinate).

Other authors (e.g. Mackowiak and Wiederholt [2009], Paciello and Wiederholt [2014], Afrouzi [2020]) also endogenize information acquisition in this setting, but follow the rational inattention approach of Sims [2003]. These models do not assume a particular set of available signals; instead, agents can choose any signal structure, subject to a cost described by mutual information. With quadratic payoffs, Gaussian priors, and mutual information costs, the agent’s optimal signal is a Gaussian signal about economic fundamentals. As a result, equilibria exhibit zero non-fundamental volatility.

Our paper also follows the rational inattention approach, but generalizes away from the mutual information cost function and incorporates exogenous public signals. As a result, our model accommodates non-fundamental volatility, building a bridge between these two seemingly distinct approaches. We focus on the question of whether agents condition their actions on noisy public signals, leading to non-fundamental volatility in equilibrium. This kind of non-fundamental volatility can occur when the equilibrium is unique. In the presence of multiple equilibria induced by strategic complementarity,³ the realizations of noisy public signals can also act as coordination devices. Our analysis focuses on whether a noisy public signal is a cost-effective means of acquiring information. When the game features strategic substitutability, this property determines whether non-fundamental volatility occurs in equilibrium. When the game instead features strategic complementarity, this property determines whether an equilibrium without non-fundamental volatility exists, but is not sufficient to determine whether additional equilibria with non-fundamental volatility also exist.

Our study of efficiency builds on the work of Angeletos and Pavan [2007] and Colombo, Femminis, and Pavan [2014]. Angeletos and Pavan [2007] study the question of constrained efficiency in the class of linear-quadratic games with exogenous information structures. We extend their results to a more general class of payoff functions, and provide necessary and sufficient conditions for the existence of a constrained efficient equilibrium under endogenous information acquisition. However, we shut down a key channel present in Angeletos and Pavan

³For an example of such a model with endogenous information acquisition, see Yang [2015].

[2007] and Colombo, Femminis, and Pavan [2014] by assuming that only the cross-sectional mean of actions, but not the cross-sectional variance, enters payoffs.⁴ We emphasize instead a different externality: the externality that arises if agents' actions affect other agents' information costs. The inefficiency we highlight is closely related to the informational externality that arises when agents observe exogenous signals about endogenous objects such as prices, as in Laffont [1985], Angeletos and Pavan [2009], Amador and Weill [2010], Vives [2017], Angeletos, Iovino, and La'O [2020].

In studying an environment in which rationally-inattentive agents can acquire information about the endogenous mean action, our work complements Denti [Forthcoming] and Angeletos and Sastry [2021]. Angeletos and Sastry [2021] similarly study the question of efficiency with rationally inattentive agents but in a distinct setting. We focus on large games with strategic interaction, while Angeletos and Sastry [2021] consider a Walrasian general equilibrium environment with complete markets over states and signal realizations in which rationally inattentive agents learn from prices. They find that invariance of the information cost is sufficient to ensure that a planner cannot improve allocations by sending a message that reduces information costs.⁵

Our focus on games with agents who can acquire information about the endogenous actions of other agents builds on Denti [Forthcoming]. Relatedly, Ravid [2020] studies bargaining between a seller and a buyer who is inattentive to both the exogenous state and the endogenous offer of the seller. We adapt the approach of Denti [Forthcoming] to static games with a continuum of identical players, and emphasize agents' beliefs about the joint distribution of the exogenous state and actions of the other agents, as in Ravid [2020]. Relative to Ravid [2020] and Denti [Forthcoming], the “largeness” feature of our class of games permits a simpler definition of equilibrium, which is essentially a Bayesian Nash equilibrium in a static, simultaneous-move game. Our definition of equilibrium can also be thought of as the limit of the dynamic process of strategic information acquisition Denti [Forthcoming] introduces. Our exercise is similar in spirit to Morris and Yang [Forthcoming], in that we relate the properties of information costs to the properties of equilibria in strategic settings.

⁴Colombo, Femminis, and Pavan [2014] study the efficiency of information acquisition within the Angeletos and Pavan [2007] linear-quadratic-Gaussian setting. They show that the absence of payoff externalities involving the mean action does not guarantee efficiency in the acquisition of information because the dispersion of actions enters payoffs—an externality not internalized by agents. This alternative form of payoff externality is absent in our game and distinct from the informational externalities that we emphasize.

⁵The two papers also differ in a number of other respects. For example, we employ the constrained efficiency concept in Angeletos and Pavan [2007] and Colombo, Femminis, and Pavan [2014] for abstract games in which the planner may only control the action functions and information choices of the players. Angeletos and Sastry [2021] consider instead a planner who can send messages to the agents, who can in turn learn about the content of the message in addition to observing exogenous states directly. They ask a different question: whether in markets the price function is an efficient conveyor of information in the sense of Hayek [1945].

2 The Game

We study a large game of strategic interaction with rational inattention.

2.1 Agents, Actions, and Payoffs

There is a continuum of agents with unit mass, indexed by $i \in [0, 1]$. Agent i chooses an action, $a^i \in A \subset \mathbb{R}$. Let $\bar{a} \in \bar{A} \subset \mathbb{R}$ denote the aggregate action, defined as

$$\bar{a} = \int_0^1 a^i di.$$

We assume that A is a convex and compact subset of the real line. Consequently, the sets A and \bar{A} are identical; we use distinct notation only to help differentiate between individual and aggregate actions.

Nature draws a stochastic, payoff-relevant state, s from a finite set $S \subset \mathbb{R}^n$.⁶ We refer to these states as the “fundamentals.” Agents have a payoff function $v : A \times \bar{A} \times S \rightarrow \mathbb{R}$; that is, an agent who takes action $a \in A$ in state $s \in S$ when the aggregate action is $\bar{a} \in \bar{A}$ receives payoff $v(a, \bar{a}, s)$. Individual agents—each of whom is infinitesimal—do not take into account how their own action affects the aggregate action. This is a defining feature of “large games.”

We focus on a specific class of utility functions that we call “mean-critical,” which have the following functional form:

$$v(a, \bar{a}, s) = g(a, s) + G(\bar{a}, s) + (a - \bar{a}) \frac{\partial}{\partial \bar{a}} G(\bar{a}, s), \quad (1)$$

where the functions $G : \bar{A} \times S \rightarrow \mathbb{R}$ and $g : A \times S \rightarrow \mathbb{R}$ are, respectively, twice differentiable in \bar{A} and A , for each value of $s \in S$. Mean-critical utility functions have the particular property that, for any conditional distribution of a given s and any $s \in S$, the mean value $E[a|s]$ is a critical point of $h(\bar{a}) = E[v(a, \bar{a}, s)|s]$.⁷ This property leads to what [Angeletos and Pavan \[2007\]](#) call “efficiency in the use of information,” which is to say that it rules out certain kinds of payoff externalities. We discuss the implications of this property below and in more detail in [Section 5](#). Some of our results will apply beyond this class of utility functions; we will indicate this where applicable.

A leading example of a mean-critical payoff function is a linear-quadratic beauty-contest game, similar to the ones studied in [Morris and Shin \[2002\]](#) and [Angeletos and Pavan \[2007\]](#), in

⁶Using a finite set S simplifies our exposition. We discuss the continuous state case in the technical appendix, [Section C](#). Our approach of using a finite set of states to avoid technical complications follows, e.g., [Kamenica and Gentzkow \[2011\]](#).

⁷In [Appendix Section A](#), we show that (subject to some regularity conditions) these utility functions are the only utility functions with this property; we also provide a multi-dimensional generalization.

which payoffs take the following form:

$$v(a, \bar{a}, s) = -(1 - \beta)(a - s)^2 - \beta(a - \bar{a})^2, \quad (2)$$

where $\beta < 1$ is a constant.⁸ The first component of (2) is a quadratic loss in the distance between the agent's action a and the fundamental s ; the second component is a quadratic loss in the distance between the agent's action a and the aggregate action \bar{a} . For this payoff function, the scalar β governs the extent of strategic complementarity ($\beta > 0$) or substitutability ($\beta < 0$) in this game. For the more general class of mean-critical utility functions in (1), the convexity or concavity of G with respect to \bar{a} determines the extent of strategic complementarity or substitutability, respectively.

2.2 Information Acquisition

Agents are rationally inattentive: they acquire information endogenously, subject to a cost. We allow agents in our framework to flexibly acquire information through multiple channels: by paying attention to the exogenous fundamental state itself, by paying attention to exogenous public signals, and by paying attention to the endogenous aggregate action, $\bar{a} \in \bar{A}$.

States and Priors. In addition to the payoff-relevant states, $s \in S$, there exists a set of exogenous noisy public signals, whose realizations r are drawn from a finite set $R \subset \mathbb{R}^n$. Noisy public signals create the potential for agents to coordinate their actions in a way that does not depend on the payoff-relevant state.

For the purposes of interpretation, it is helpful to think of nature as drawing a set of states $e \in \mathcal{E}$ that are independent of the fundamental states $s \in S$. The public signal r is then generated through some mapping $\mathcal{E} \times S \rightarrow R$. The dependence of the public signal $r \in R$ on the shocks $e \in \mathcal{E}$ introduces a potential source of “non-fundamental” volatility in equilibrium. We treat the set of possible public signal realizations R as exogenous, but allow it to be arbitrarily rich.

Let $\mathcal{U}_0 \equiv \Delta(S \times R)$ denote the space of probability measures on $S \times R$.⁹ Agents share a common, full-support prior probability measure $\mu_0 \in \mathcal{U}_0$.

In addition to the exogenous states s and r , agents can pay attention to the endogenous aggregate action \bar{a} . Let $\bar{\alpha} : S \times R \rightarrow \bar{A}$ be a function mapping exogenous states to an aggregate action; this function will be determined in equilibrium by aggregating over the individual agents' strategies.¹⁰ Let $\bar{\mathcal{A}}$ be the space of all such aggregate action functions.

⁸This is a special case of (1) with $G(\bar{a}, s) = \beta\bar{a}^2$ and $g(a, s) = -(1 - \beta)(a - s)^2 - \beta a^2$. The assumption that $\beta < 1$ ensures that a unique symmetric pure-strategy Nash equilibrium exists under complete information given by $a^i = \bar{a} = s$.

⁹We use the notation $\Delta(X)$, for any topological space X , to denote the set of probability measures defined on the measurable space $(X, \mathcal{B}(X))$, where $\mathcal{B}(X)$ is the Borel sigma-algebra of X .

¹⁰Note that we assume $\bar{\alpha}$ is a deterministic function of (s, r) ; we will subsequently define an equilibrium in a

Consider the space of probability measures $\Delta(S \times R \times \bar{A})$, endowed with the weak* topology. Using the $\bar{\alpha}$ function, we can define an agent's prior $\mu \in \Delta(S \times R \times \bar{A})$ on this space. Let $\phi : \mathcal{U}_0 \times \bar{A} \rightarrow \Delta(S \times R \times \bar{A})$ denote a mapping from any pair $(\mu_0, \bar{\alpha})$ to its induced probability measure, defined by, for all $(s, r) \in S \times R$ and open subsets $\bar{A}_0 \subseteq \bar{A}$,

$$\phi[\mu_0, \bar{\alpha}](s, r, \bar{A}_0) = \begin{cases} \mu_0(s, r) & \text{if } \bar{\alpha}(s, r) \in \bar{A}_0, \\ 0 & \text{otherwise.} \end{cases} \quad (3)$$

The agent's prior is given by $\mu = \phi[\mu_0, \bar{\alpha}]$.

The mapping ϕ induces a measure on $S \times R \times \bar{A}$ by combining the exogenous prior μ_0 over $S \times R$ and the equilibrium conditional distribution of \bar{a} given (s, r) . Because the aggregate action is a deterministic function of s and r , the conditional distribution of \bar{a} given (s, r) is degenerate: it is a point mass at $\bar{\alpha}(s, r)$. This degeneracy arises from the largeness of the game, which ensures that despite the randomness in each individual agent's actions, the aggregate action is a deterministic function of the exogenous states.¹¹

Define the set of all measures in $\Delta(S \times R \times \bar{A})$ that may be generated by some pair $(\mu'_0, \bar{\alpha})$ as

$$\mathcal{U} = \{\mu' \in \Delta(S \times R \times \bar{A}) : \exists(\mu'_0, \bar{\alpha}) \in \mathcal{U}_0 \times \bar{A} \text{ s.t. } \mu' = \phi[\mu'_0, \bar{\alpha}]\}.$$

In what follows, we use the subscript zero to denote objects defined on $S \times R$ as opposed to $S \times R \times \bar{A}$ (as with \mathcal{U}_0 and \mathcal{U}).

Strategies. In rational inattention problems, we typically think of the agent as choosing a signal structure and an action conditional on any signal realization. A signal structure is a conditional distribution over signal realizations; the probability of observing a particular signal realization depends on the underlying state. Observing a particular signal realization leads the agent to form posterior beliefs over the states and, based on that posterior, to choose an action (or distribution over actions) optimally.

We forgo these intermediate steps and write the agent's strategy directly as a joint measure over actions and posteriors, $\pi \in \Delta(A \times \mathcal{U})$.¹² We use the notation $E^\pi[f(a, \mu')]$ to indicate $\int_{\text{supp}(\pi)} f(a, \mu') d\pi(a, \mu')$; in particular, note that the expectation is always taken over the posterior probability measures μ' , not the prior measure μ .

We require strategies to be “Bayes-consistent” with the prior: the expectation of the posterior

manner consistent with this assumption.

¹¹Note that this determinacy does not preclude non-fundamental volatility, as the aggregate action can condition on the public signal realization r .

¹²This joint distribution over actions and posteriors can be thought of as summarizing the agent's state-dependent stochastic choice (Caplin and Dean [2015]). We discuss the relationship between signals, actions, and distributions over posteriors and actions in detail in Appendix Section B, and see also e.g. Proposition 1 of Kamenica and Gentzkow [2011]. These joint distributions can be viewed as a kind of distributional strategy (Milgrom and Weber [1985]), with the posterior playing the role of the type, and with an endogenous distribution of types.

under π must be equal to the prior, $\mu \in \mathcal{U}$. Formally, a strategy $\pi \in \Delta(A \times \mathcal{U})$ is “Bayes-consistent” with the prior $\mu \in \mathcal{U}$ if, for all $(a, \mu') \in \text{supp}(\pi)$, μ' is absolutely continuous with respect to μ (denoted $\mu' \ll \mu$) and if $E^\pi[\mu'] = \mu$.¹³ Because we have assumed S and R are finite and that μ_0 has full support, absolute continuity requires only that the posteriors μ' be consistent with the aggregate action strategy \bar{a} that generated μ from μ_0 .

Let $\Pi(\mu) \subset \Delta(A \times \mathcal{U})$ denote the set of joint measures over actions and posteriors that are Bayes-consistent with the prior $\mu \in \mathcal{U}$, and endow it with the weak* topology. We call $\Pi(\mu)$ the set of feasible strategies.

Information Costs. While all strategies in $\Pi(\mu)$ are feasible, some strategies are more costly to the rationally inattentive agent than others.

We define the agent’s information acquisition costs directly over posteriors. Following [Caplin, Dean, and Leahy \[2022\]](#), we assume costs take the following “posterior-separable” form.

Definition 1. A **posterior-separable cost function** is a function $C : \{(\pi, \mu) \in \Delta(A \times \mathcal{U}) \times \mathcal{U} : \pi \in \Pi(\mu)\} \rightarrow \mathbb{R}_+$ that can be written as

$$C(\pi, \mu) = E^\pi[D(\mu' || \mu)], \quad (4)$$

where $D : \{(\mu', \mu) \in \mathcal{U} \times \mathcal{U} : \mu' \ll \mu\} \rightarrow \mathbb{R}_+$ is a divergence, convex in its first argument.

That is, given a prior $\mu \in \mathcal{U}$, an agent who plays a Bayes-consistent strategy $\pi \in \Pi(\mu)$ incurs information costs $C(\pi, \mu)$. Note that we have written the cost C as a function of the strategy π , which is convenient for notational purposes, even though the cost depends only on the distribution of the posterior beliefs μ' and not the conditional distribution of actions a given each posterior belief.

The cost function in (4) has an intuitive interpretation. A divergence is defined as a function of two measures that is zero if and only if its arguments are equal, and is otherwise strictly positive. It could be interpreted as a measure of how “close” or “far” two measures are from one another, but note that it is not technically a distance: it need not be symmetric nor satisfy the triangle inequality.

In our context, the divergence is a function of the prior and a posterior. As agents gather information, their posterior beliefs move away from their prior as measured by this divergence. How this manifests in information costs is reflected in the cost function (4). Note that the cost of a strategy that gathers no information (all posteriors μ' are equal to the prior μ) is zero.

Finally, to ensure that the agent’s problem is well-behaved, we assume the divergence is continuous and differentiable.

¹³Note here that the equality defining Bayes-consistency should be understood in the weak* sense (see e.g. [Aliprantis and Border \[2006\]](#) section 15.1), as in for example [Jung et al. \[2019\]](#) or the online appendix of [Kamenica and Gentzkow \[2011\]](#).

Assumption 1. D is jointly continuous on $\{(\mu', \mu) \in \mathcal{U} \times \mathcal{U} : \mu' \ll \mu\}$. $D(\phi\{\mu'_0, \bar{\alpha}\} | \phi\{\mu_0, \bar{\alpha}\})$ is continuously differentiable on $(\mu'_0, \bar{\alpha}) \in \text{int}(\mathcal{U}_0 \times \bar{\mathcal{A}})$.¹⁴

2.3 The Agent's Problem

Given the prior μ , the problem of the agent is to choose a strategy $\pi \in \Pi(\mu)$ that maximizes her expected utility less information costs. The separability of expected utility and posterior-separability of information costs together imply that the agent's problem can be written as

$$\max_{\pi \in \Pi(\mu)} E^\pi [V(a, \mu') - D(\mu' | \mu)],$$

where $V(a, \mu') \equiv \int_{\text{supp}(\mu')} v(a, \bar{a}, s) d\mu'(s, r, \bar{a})$ denotes the agent's expected utility from taking action $a \in A$ under posterior measure $\mu' \in \mathcal{U}$. [Caplin, Dean, and Leahy \[2022\]](#) call $V(a, \mu') - D(\mu' | \mu)$ the net utility function.

2.4 Equilibrium Definition

Let $\xi \equiv (\pi, \bar{\alpha})$ denote a symmetric strategy profile, consisting of a strategy $\pi \in \Delta(A \times \mathcal{U})$ for all agents and an aggregate action function $\bar{\alpha} \in \bar{\mathcal{A}}$. In a symmetric equilibrium in which all agents play π , we will require that the aggregate action function $\bar{\alpha}$ be consistent with the mean action generated by the strategy π . Specifically, we assume that, conditional on (s, r, \bar{a}) , the realizations of actions across agents are independent. That is, it is only the measures over actions, not the realizations, that are identical across agents. Independence of realizations allows us to apply the law of large numbers and require that the expected individual action be consistent with the mean action \bar{a} in the population [[Uhlig, 1996](#)]. We impose this as follows.

Definition 2. A symmetric strategy profile $\xi = (\pi, \bar{\alpha})$ is **mean-consistent** if, for all $(s, r) \in S \times R$,

$$\bar{\alpha}(s, r) = E^\pi[a | s, r], \tag{5}$$

where $E^\pi[\cdot | s, r]$ denotes the conditional expectation given $(s, r) \in S \times R$ under π .

Because we have assumed S and R are finite, defining this conditional expectation is straightforward.¹⁵ Note that the conditional independence of actions does not rule out coordination across agents. Agents can attempt to coordinate their actions with other agents by conditioning on r or \bar{a} , or choose to act independently of other agents by conditioning only on s .¹⁶ We model

¹⁴Because $S \times R$ is finite, we can regard μ_0 and $\bar{\alpha}$ as vectors in $\mathbb{R}^{|S| \times |R|}$ and define differentiability in the usual way.

¹⁵With S and R finite, for any $\mu' = \phi[\mu'_0, \bar{\alpha}]$ with $\mu' \ll \mu$, there is a unique density $f_0[\mu', \mu]$ of μ' with respect to μ that does not depend on \bar{a} . We can define the conditional expectation as $E^\pi[a | s, r] = \int_{\text{supp}(\pi)} a f_0[\mu', \mu](s, r) d\pi(a, \mu')$.

¹⁶In the context of posterior beliefs, an agent conditions on r or \bar{a} if the conditional distributions of (r, \bar{a}) given s under her prior and posterior beliefs differ.

information acquisition in this way to capture the idea that agents can observe both public signals (r) and the actions of other agents (\bar{a}) in addition to acquiring private signals.

We define an equilibrium in our game as a symmetric strategy profile in which the strategies π are individually optimal, the prior μ is consistent with the aggregate action function $\bar{a} \in \bar{\mathcal{A}}$, and \bar{a} is consistent with the mean action generated by π .

Definition 3. Given a common prior $\mu_0 \in \mathcal{U}_0$, an **equilibrium** of the game is a mean-consistent, symmetric strategy profile $\xi = (\pi, \bar{a})$ such that the strategy $\pi \in \Pi(\mu)$ is a best response to the prior $\mu = \phi[\mu_0, \bar{a}]$,

$$\pi \in \arg \max_{\pi' \in \Pi(\mu)} E^{\pi'} [V(a, \mu') - D(\mu' || \mu)]. \quad (6)$$

Our equilibrium definition is a hybrid of a Bayesian Nash equilibrium and a Rational Expectations equilibrium. It is a Bayesian Nash equilibrium in the sense that agents play best responses under incomplete information. It is a Rational Expectations equilibrium [Grossman, 1976, Grossman and Stiglitz, 1980, 1976] in the sense that agents can learn from endogenous aggregate actions while simultaneously choosing their strategy. Therefore, beliefs must be consistent with endogenous actions, while actions are best responses to endogenous beliefs. An equilibrium in our game, as in any Rational Expectations equilibrium, is a fixed point of this mapping.¹⁷ Our equilibrium concept builds on Denti [Forthcoming], who shows that our static equilibrium concept can be justified as the steady-state limit of a dynamic game in which rationally inattentive agents acquire information about the past actions of other agents.

2.5 Alternative Formulation

The agent's problem and the definitions of mean consistency and equilibrium can be formulated over the space of posteriors on $S \times R$ instead of $S \times R \times \bar{\mathcal{A}}$. This equivalence is a consequence of the fact that for any $\mu, \mu' \in \mathcal{U}$ with $\mu' \ll \mu$, there is a $\mu'_0 \in \mathcal{U}_0$ and $\bar{a} \in \bar{\mathcal{A}}$ such that $\mu' = \phi[\mu'_0, \bar{a}]$ and $\mu = \phi[\mu_0, \bar{a}]$. In the analysis that follows, we will switch between these two equivalent definitions of equilibrium as convenient. Introducing both definitions avoids the need to use more cumbersome notation in some of our results.

Let $\Pi_0(\mu_0) \subset \Delta(A \times \mathcal{U}_0)$ be the set of strategies satisfying Bayes-consistency with respect to μ_0 , endowed with the weak* topology.

Definition 4. A strategy profile (π_0, \bar{a}) is mean-consistent if $\bar{a}(s, r) = E^{\pi_0} [a | s, r]$. A strategy profile (π_0, \bar{a}) is an equilibrium if it is mean-consistent and

$$\pi_0 \in \arg \max_{\pi'_0 \in \Pi_0(\mu_0)} E^{\pi'_0} [V(a, \phi[\mu'_0, \bar{a}]) - D(\phi[\mu'_0, \bar{a}] || \phi[\mu_0, \bar{a}])].$$

¹⁷Our equilibrium concept can be thought of as a generalization of a Walrasian equilibrium, interpreting the aggregate action as prices. As in Walrasian equilibria, agents' actions aggregate to determine prices, which in turn determine agents' optimal actions. As with Walrasian equilibria, it is convenient to study an equilibrium without explicitly specifying the process by which the economy arrives at that equilibrium.

Note that, given some $(\pi_0, \bar{\alpha}) \in \Pi_0(\mu_0) \times \bar{\mathcal{A}}$, we can always construct the corresponding $\pi \in \Pi(\phi[\mu_0, \bar{\alpha}])$ as the measure induced from π_0 by the mapping $(a, \mu'_0) \mapsto (a, \phi[\mu'_0, \bar{\alpha}])$.¹⁸ Similarly, given some $\pi \in \Pi(\phi[\mu_0, \bar{\alpha}])$, we can construct the corresponding π_0 .¹⁹

2.6 Remarks on the Model

This concludes our description of the model. We have made a few modeling choices that depart from the standard rational inattention paradigm; we discuss these choices below.

Mean-Critical Utility. We restrict attention to the class of “mean-critical” utility functions. Imposing this restriction shuts down payoff externalities, which we discuss in more detail in Section 5. In appendix section A, we show that this class of utility functions extends the [Angelos and Pavan \[2007\]](#) notion of “efficiency in the use of information” beyond the class of linear-quadratic preferences. That is, with this (and, subject to regularity conditions, only this) class of utility functions, when information is exogenous, an equilibrium of the game will coincide with the solution to a planner’s problem. We impose this restriction in order to focus our attention on externalities related to information acquisition rather than payoffs.

Posterior-Separable Costs. The cost function assumed in much of the rational inattention literature, following [Sims \[2003\]](#), is mutual information. Relative to this benchmark, we consider a more general class of cost functionals: those that are “posterior-separable” [[Caplin, Dean, and Leahy, 2022](#)]. Considering a general class of cost functions is necessary to address the question of the relationship between properties of the cost function and properties of equilibria.

Studying posterior-separable cost functions allows us to consider the properties of cost functions in terms of properties of their associated divergences. That said, there are a number of other reasons why posterior-separable cost functions are appealing. First, the posterior-separable class nests a number of cost functions proposed in the literature on information choice. Among these include mutual information [[Sims, 2003](#)], Tsallis entropy costs [[Caplin, Dean, and Leahy, 2022](#)], the Log-Likelihood Ratio (LLR) cost function [[Pomatto, Strack, and Tamuz, 2020](#)], and Fisher information and neighborhood-based cost functions [[Hébert and Woodford, 2021](#)]. Second, while experimental/laboratory data have rejected mutual information (see, e.g., [[Dean and Neligh, 2019](#)]), we are not aware of any results rejecting posterior-separability. Third, posterior-separable cost functions can be studied via Lagrangian methods [[Caplin, Dean, and Leahy, 2022](#)], which facilitates certain proofs.

¹⁸Here and throughout the paper, we say a measure (in this case π) is induced from another measure (in this case π_0) by some mapping when the former is the push-forward measure of the latter under that mapping.

¹⁹The relevant mapping is $(a, \mu') \mapsto (a, \gamma_{-\bar{\alpha}}[\mu'])$, where $\gamma_{-\bar{\alpha}}$ computes the marginal distribution on $S \times R$ and is defined in Section 3 below.

That said, many of our results could be phrased in terms of the cost functions C instead of the divergences D , and thereby extended beyond the class of posterior-separable cost functions. See Section 3.3 below.

Noisy Public Signals. We assume the existence of an arbitrarily rich set of exogenous public signal realizations, $r \in R$. While exogenous public signals are standard in games with incomplete but exogenous information (see, e.g. Morris and Shin [2002], Angeletos and Pavan [2007], Bergemann and Morris [2013]), public signals are less standard in the rational inattention framework.

We do not see the existence of exogenous public signals as antithetical to the spirit of rational inattention; rather, it allows us to build a bridge between these two literatures. As in games with exogenous, incomplete information, the public signal in our framework can serve as a coordinating device for agents. If agents choose actions that correlate with r , non-fundamental volatility can arise in equilibrium (in the form of non-fundamental noise, $e \in \mathcal{E}$).

However, the mere existence of public signals is not sufficient to generate non-fundamental volatility. Rationally inattentive agents optimally choose what to pay attention to; as a result, whether or not non-fundamental noise manifests in equilibrium will depend on whether or not agents choose to correlate their actions with $r \in R$. This, as we will subsequently show, depends on the nature of information costs.²⁰

Learning from the Aggregate Action. In many economic contexts, agents track not only exogenous fundamentals, but also endogenous outcomes. In Walrasian markets, consumers learn from market-clearing prices.²¹ In financial markets, traders track asset prices and order books. In economic recessions and expansions, firms and households monitor aggregate indicators such as GDP, unemployment, and inflation.

In many of these contexts, observing endogenous statistics might in fact be a more efficient way of acquiring information than doing one's own research on economic fundamentals. This motivates our choice to incorporate learning of this type by allowing agents to pay attention to the endogenous aggregate action. In particular, we define the cost function in (4), and its associated divergence, on priors and posteriors over $S \times R \times \bar{A}$.²²

²⁰Note that in exogenous-information games with public signals, agents are often described as separately observing a private signal and the public signal, and then forming (degenerate) posteriors on the joint distribution of the fundamental and noisy public signal realization. In our game, agents can be thought of as observing a private signal about both the fundamental state and the realization of the public signal, and then forming posteriors. The only difference between these two approaches is that in our setting the agents may not perfectly observe the public signal realization.

²¹See e.g. the related analysis of rational inattention with a Walrasian market in Angeletos and Sastry [2021].

²²Rational expectations are maintained by requiring that the prior over $S \times R \times \bar{A}$ is consistent with the aggregate action function (3) and that the aggregate action function is consistent with the mean (5). As in Ravid [2020], defining information costs on this space allows the distribution of the endogenous variables (as summarized by $\bar{\alpha}$ in our model) to potentially affect information costs.

Circumvention of Signal Structures. In defining information costs directly over posteriors, we have said nothing about what kind of signal structure generates these posteriors. In fact, in our setting, there are multiple ways of arriving at the same posterior. We view the cost function C as “reduced-form” in the sense that it represents the cost of reaching the given posteriors by the least costly signals available.

Suppose, by way of example, that in equilibrium the aggregate action depends only on s and is linear: $\bar{\alpha}(s, r) = \bar{\alpha}_s s$ for some non-zero constant $\bar{\alpha}_s$. In this case, the agent could acquire information about s and \bar{a} either by paying attention to s (“doing her own fundamental research”), by paying attention to \bar{a} (“learning from the actions of others”), or some combination thereof. Rather than model these signal structures, we let C summarize the *result* of the agent’s optimization over how best to acquire information, i.e. how best to reach a given set of posteriors. We then ask: given the reduced-form cost function, what strategy π is desirable?

This reduced-form perspective furthermore explains why the divergence can in general depend on the aggregate action function $\bar{\alpha}$. In the example above, it might be optimal for the agent to acquire information by paying attention to the aggregate action \bar{a} , rather than doing her own fundamental research. If this is the case, then changes to the aggregate action function $\bar{\alpha}$ might change the value of the divergence holding the posteriors on $S \times R$ constant. For example, interpreting \bar{a} as a price, we might expect that larger price movements are less costly to observe. See Section 7.4 for a concrete example and Appendix Section B for more on this point. Section 7.4 also illustrates another point: the reduced form cost of information might depend on $\bar{\alpha}$ even if the underlying cost of observing a signal about s or \bar{a} does not depend on the agent’s subjective prior.²³

3 Partial Monotonicity and Partial Invariance

The focus of our investigation is the relationship between properties of the information cost function, in particular its associated divergence, and properties of equilibria. In this section we introduce the concepts of “partial monotonicity” and “partial invariance.” We focus on two particular versions of these properties, “monotonicity in R ” and “invariance in \bar{A} ,” and provide examples of cost functions that do and do not exhibit these properties.

3.1 Monotonicity and Invariance in R and in \bar{A}

The properties we introduce in this section describe how a divergence responds to certain transformations of its inputs, the prior and posterior. We begin by defining these transformations.

²³Section 7.4 makes this point in the context of Gaussian signals; the same point could be made using a fixed-prior (following Denti et al. [Forthcoming]) version of the cost $D_{S\bar{A}}$ presented in the next section.

Let $\gamma_{-R} : \Delta(S \times R \times \bar{A}) \rightarrow \Delta(S \times \bar{A})$ be the function that generates the marginal distribution on $S \times \bar{A}$ from a given measure $\mu \in \Delta(S \times R \times \bar{A})$; we use the subscript $-R$ to indicate the dimension being integrated out.²⁴

Next, define an operator $\eta_{-R|R} : \Delta(S \times R \times \bar{A}) \times \Delta(S \times R \times \bar{A}) \rightarrow \Delta(S \times R \times \bar{A})$. This operator takes a pair of probability measures, $\mu_1, \mu_2 \in \Delta(S \times R \times \bar{A})$, with $\gamma_{-R}[\mu_1] \ll \gamma_{-R}[\mu_2]$, and transforms them into a new probability measure $\eta_{-R|R}[\mu_1|\mu_2] \in \Delta(S \times R \times \bar{A})$, with $\eta_{-R|R}[\mu_1|\mu_2] \ll \mu_2$, defined by

$$\frac{d\eta_{-R|R}[\mu_1|\mu_2]}{d\mu_2} = \frac{d\gamma_{-R}[\mu_1]}{d\gamma_{-R}[\mu_2]}, \quad (7)$$

where $\frac{d\eta_{-R|R}[\mu_1|\mu_2]}{d\mu_2}$ and $\frac{d\gamma_{-R}[\mu_1]}{d\gamma_{-R}[\mu_2]}$ denote Radon-Nikodym derivatives. This expression can be rewritten, using the finiteness of R , as stating that for all $(s, r) \in S \times R$ and all open sets $\bar{A}_0 \subseteq \bar{A}$,

$$\eta_{-R|R}[\mu_1|\mu_2](s, r, A_0) = \begin{cases} \mu_2(s, r, A_0) \frac{\sum_{r' \in R} \mu_1(s, r', A_0)}{\sum_{r' \in R} \mu_2(s, r', A_0)} & \text{if } \sum_{r' \in R} \mu_2(s, r', \bar{A}_0) > 0, \\ 0 & \text{otherwise.} \end{cases}$$

The probability measure $\eta_{-R|R}[\mu_1|\mu_2]$ shares a marginal distribution on “not- R ” ($S \times \bar{A}$ in our context) with μ_1 , and shares a conditional distribution of R given $(s, \bar{a}) \in S \times \bar{A}$ with μ_2 . That is, $\eta_{-R|R}$ replaces the conditional distribution of r given (s, \bar{a}) from μ_1 with that of μ_2 , while preserving the marginal distribution on $S \times \bar{A}$ of μ_1 .²⁵ In this sense, it is closer to μ_2 than μ_1 was originally.

Example. In this example we illustrate the mechanics of the γ_{-R} and $\eta_{-R|R}$ operators. For simplicity, we assume $\bar{\alpha}(s, r) = \bar{a}_0$ is constant and that the sets S and R consist of only two states: $S = \{s_1, s_2\}$ and $R = \{r_1, r_2\}$. Tables 1 and 2 provide examples of two measures, $\mu_1, \mu_2 \in \mathcal{U}$.

$\bar{a} = \bar{a}_0$	r_1	r_2
s_1	0.45	0.15
s_2	0.30	0.10

Table 1. $\mu_1 \in \mathcal{U}$

	$\bar{a} = \bar{a}_0$
s_1	3
s_2	1/2

Table 3. $\frac{d\gamma_{-R}[\mu_1]}{d\gamma_{-R}[\mu_2]}$

$\bar{a} = \bar{a}_0$	r_1	r_2
s_1	0.10	0.10
s_2	0.20	0.60

Table 2. $\mu_2 \in \mathcal{U}$

$\bar{a} = \bar{a}_0$	r_1	r_2
s_1	0.30	0.30
s_2	0.10	0.30

Table 4. $\eta_{-R|R}[\mu_1|\mu_2] \in \mathcal{U}$

²⁴Formally, define the transformation $T_R : S \times R \times \bar{A} \rightarrow S \times \bar{A}$ by $T_R(s, r, \bar{a}) = (s, \bar{a})$, and define $\gamma_{-R}[\mu]$ as the push-forward measure induced from $\mu \in \Delta(S \times R \times \bar{A})$ by T_R .

²⁵Equivalently, the operator replaces the marginal distribution on $S \times \bar{A}$ of μ_2 with that of μ_1 .

First, we apply γ_{-R} to both measures to obtain their marginal distributions on $S \times \bar{A}$. Summing the rows of Tables 1 and 2 yields, respectively, (0.6, 0.4) and (0.2, 0.8). Consequently, $\frac{d\gamma_{-R}[\mu_1]}{d\gamma_{-R}[\mu_2]}$, displayed in Table 3, is equal to 3 if $(s, \bar{a}) = (s_1, \bar{a}_0)$ and 1/2 if $(s, \bar{a}) = (s_2, \bar{a}_0)$.

Multiplying the first row of Table 2 by 3 and the second row of Table 2 by 1/2 results in the measure $\eta_{-R|R}[\mu_1|\mu_2] \in \mathcal{U}$ displayed in Table 4. This measure is arguably “more like” μ_2 than μ_1 was originally in that they share a common conditional distribution of r given (s, \bar{a}) .

Monotonicity in R. We now apply the transformation $\eta_{-R|R}$ defined in (7) to a posterior and prior probability measure, μ' and μ . We will call a divergence “monotone in R ” if making the posterior “more like” the prior in the sense described by $\eta_{-R|R}$ reduces the divergence from the prior to the posterior.

Definition 5. A divergence D is **monotone in R given** $\mu \in \mathcal{U}$ if for all $\mu' \in \mathcal{U}$ with $\mu' \ll \mu$,

$$D(\mu'|\mu) \geq D(\eta_{-R|R}[\mu'|\mu]|\mu). \quad (8)$$

A divergence is **monotone in R** if it is monotone in R given any $\mu \in \mathcal{U}$.

This property compares the divergence from the prior μ to the posterior μ' before and after replacing the conditional distribution on r given (s, \bar{a}) of μ' with that of μ . If replacing the posterior’s conditional distribution of r given (s, \bar{a}) with that of the prior weakly reduces their divergence, the divergence is monotone in R .

Invariance in R. We use the same transformation $\eta_{-R|R}$ to define another property of a divergence, which we will call “invariance in R .”

Definition 6. A divergence D is **invariant in R given** $\mu \in \mathcal{U}$ if for all $\mu', \mu'' \in \mathcal{U}$ with $\mu' \ll \mu$ and $\gamma_{-R}[\mu] \ll \gamma_{-R}[\mu'']$,

$$D(\eta_{-R|R}[\mu'|\mu]|\mu) = D(\eta_{-R|R}[\mu'|\mu'']|\eta_{-R|R}[\mu|\mu'']). \quad (9)$$

A divergence is **invariant in R** if it is invariant in R given any $\mu \in \mathcal{U}$.

This property compares the divergence from the prior μ to the posterior μ' after replacing *both* of their conditional distributions on r given (s, \bar{a}) with those of a third distribution, μ'' . Invariance in R requires that, for any given prior and posterior, (9) holds for all possible μ'' . Invariance in R thereby captures the idea that if the prior and the posterior were to share a common conditional distribution of r given (s, \bar{a}) , the particulars of this conditional distribution would be immaterial for their divergence.

Invariance in R and monotonicity in R are different properties of divergences, and one property does not imply the other. Invariance and monotonicity together require the divergence to

shrink to the same value for *all possible* conditional distributions:

$$D(\mu' || \mu) \geq D(\eta_{-R|R}[\mu' | \mu] || \mu) = D(\eta_{-R|R}[\mu' | \mu''] || \eta_{-R|R}[\mu | \mu''])$$

for all $\mu, \mu', \mu'' \in \mathcal{U}$ with $\mu' \ll \mu$ and $\gamma_{-R}[\mu] \ll \gamma_{-R}[\mu'']$.

Invariance in R will play no role in the analysis that follows; we define it only to introduce the concept of partial invariance and clarify the difference between partial invariance and partial monotonicity. Instead, invariance in \bar{A} (along with monotonicity in R) will be the focus of our analysis. We define invariance in \bar{A} next.

Invariance in \bar{A} . We proceed exactly as above, with \bar{A} in the place of R . Let $\gamma_{-\bar{A}} : \Delta(S \times R \times \bar{A}) \rightarrow \Delta(S \times R)$ be the function that generates the marginal distribution on $S \times R$ from a given measure $\mu \in \Delta(S \times R \times \bar{A})$. We define an operator $\eta_{-\bar{A}|\bar{A}} : \Delta(S \times R \times \bar{A}) \times \Delta(S \times R \times \bar{A}) \rightarrow \Delta(S \times R \times \bar{A})$ that takes a pair of measures, $\mu_1, \mu_2 \in \Delta(S \times R \times \bar{A})$, with $\gamma_{-\bar{A}}[\mu_1] \ll \gamma_{-\bar{A}}[\mu_2]$, and transforms them into a measure $\eta_{-\bar{A}|\bar{A}}[\mu_1 | \mu_2] \in \Delta(S \times R \times \bar{A})$ defined by

$$\frac{d\eta_{-\bar{A}|\bar{A}}[\mu_1 | \mu_2]}{d\mu_2} = \frac{d\gamma_{-\bar{A}}[\mu_1]}{d\gamma_{-\bar{A}}[\mu_2]}. \quad (10)$$

With this operator, we define our concept of invariance in \bar{A} in a manner analogous to our previous definition of invariance in R .

Definition 7. A divergence D is **invariant in \bar{A} given $\mu \in \mathcal{U}$** , if for all $\mu', \mu'' \in \mathcal{U}$ with $\mu' \ll \mu$ and $\gamma_{-\bar{A}}[\mu] \ll \gamma_{-\bar{A}}[\mu'']$,

$$D(\eta_{-\bar{A}|\bar{A}}[\mu', \mu] || \mu) = D(\eta_{-\bar{A}|\bar{A}}[\mu', \mu''] || \eta_{-\bar{A}|\bar{A}}[\mu, \mu'']). \quad (11)$$

A divergence is **invariant in \bar{A}** if it is invariant in \bar{A} given any $\mu \in \mathcal{U}$.

Invariance in \bar{A} captures the idea that if the prior and the posterior were to share a common conditional distribution over $\bar{a} \in \bar{A}$, the exact values of this conditional distribution would be immaterial for the divergence. In our context, the prior and the posterior always share a common conditional distribution over \bar{a} , determined by the \bar{a} function. As a result invariance in \bar{A} determines whether or not information costs depend on the aggregate action function \bar{a} .

3.2 Cost Function Examples

We next provide several examples of posterior-separable cost functions and their associated divergences, and discuss whether these divergences are monotone in R and/or invariant in \bar{A} . All of the example cost functions featured here satisfy our continuity and differentiability assumptions (Assumption 1); see Appendix Section E.1.

Mutual Information. Mutual information is a posterior-separable cost function: its associated divergence is the Kullback-Leibler divergence. In our context, the Kullback-Leibler (KL) divergence from the prior $\mu = \phi[\mu_0, \bar{\alpha}]$ to the posterior $\mu' = \phi[\mu'_0, \bar{\alpha}]$ can be written, for μ'_0 interior, as

$$D_{KL}(\mu' || \mu) = E^{\mu'} \left[\ln \left(\frac{d\mu'(s, r, \bar{a})}{d\mu(s, r, \bar{a})} \right) \right] = \sum_{s \in S} \sum_{r \in R} \mu'_0(s, r) \ln \left(\frac{\mu'_0(s, r)}{\mu_0(s, r)} \right). \quad (12)$$

We extend D_{KL} to all $\mu, \mu' \in \mathcal{U}$ with $\mu' \ll \mu$ by continuity. Observe that, because μ' and μ are degenerate on \bar{a} given (s, r) , the Kullback-Leibler divergence does not depend on $\bar{\alpha}$.

The KL divergence is both monotone in R and invariant in \bar{A} . Invariance in \bar{A} follows immediately from the lack of dependence on the aggregate action function $\bar{\alpha}$.

Monotonicity in R follows from a Jensen's inequality argument. Consider any $\mu, \mu' \in \mathcal{U}$ and with $\mu' \ll \mu$. Applying Jensen's inequality in the R dimension,

$$\begin{aligned} D_{KL}(\mu' || \mu) &= E^\mu \left[E^{\mu'} \left[\frac{d\mu'(s, r, \bar{a})}{d\mu(s, r, \bar{a})} \ln \left(\frac{d\mu'(s, r, \bar{a})}{d\mu(s, r, \bar{a})} \right) \mid s, \bar{a} \right] \right] \\ &\geq E^\mu \left[\frac{d\gamma_{-R}[\mu'](s, \bar{a})}{d\gamma_{-R}[\mu](s, \bar{a})} \ln \left(\frac{d\gamma_{-R}[\mu'](s, \bar{a})}{d\gamma_{-R}[\mu](s, \bar{a})} \right) \right] = D(\eta_{-R|R}[\mu' || \mu] || \mu), \end{aligned}$$

which is the definition of monotonicity in R .

Cost Functions with Perceptual Distance. Several recent papers have described information costs that, unlike mutual information, incorporate a notion of ‘‘perceptual distance’’ (e.g. Hébert and Woodford [2021], Pomatto et al. [2020], Walker-Jones [2019], Bloedel and Zhong [2020]). As an example we consider a particular kind of neighborhood-based cost function proposed in Hébert and Woodford [2021] that is related to the mutual information cost function described above (and to the costs considered in Walker-Jones [2019]).

Let $S = \{s_1, \dots, s_{|S|}\} \subset \mathbb{R}$ with s_i strictly increasing in i , and $R = \{r_1, \dots, r_{|R|}\} \subset \mathbb{R}$ with r_j strictly increasing in j .²⁶ The following neighborhood-based cost function treats each adjacent pair (s_{i-1}, s_i) and (r_{j-1}, r_j) as a ‘‘neighborhood,’’ and penalizes differences between the posterior and prior within each neighborhood using the KL divergence.

Define the conditional KL divergences within these neighborhoods as

$$D_{KL,i,r_j}(\mu'_0 || \mu_0) = E^{\mu'_0} \left[\ln \left(\frac{\mu'_0(s, r_j) \mu_0(s_i, r) + \mu_0(s_{i-1}, r_j)}{\mu_0(s, r_j) \mu'_0(s_i, r_j) + \mu'_0(s_{i-1}, r_j)} \right) \mid s \in \{s_{i-1}, s_i\} \right]$$

and

$$D_{KL,j,s_i}(\mu'_0 || \mu_0) = E^{\mu'_0} \left[\ln \left(\frac{\mu'_0(s_i, r) \mu_0(s_i, r_j) + \mu_0(s_i, r_{j-1})}{\mu_0(s_i, r) \mu'_0(s_i, r_j) + \mu'_0(s_i, r_{j-1})} \right) \mid r \in \{r_{j-1}, r_j\} \right]$$

²⁶The purpose of assuming an order and a metric on S and R is to define a perceptual distance on these sets. As we discuss below, this perceptual distance should be understood as summarizing how difficult or easy it is to distinguish between these dates.

The general family is defined, for $\mu = \phi[\mu_0, \bar{\alpha}]$ and $\mu' = \phi[\mu'_0, \bar{\alpha}]$ with μ'_0 interior, as

$$D_N(\mu' || \mu) = \sum_{i=2}^{|S|} \sum_{j=1}^{|R|} \frac{\mu'_0(s_i, r_j) + \mu'_0(s_{i-1}, r_j)}{d_{ij}^s(\bar{\alpha})} D_{KL, i, r_j}(\mu'_0 || \mu_0) \\ + \sum_{i=1}^{|S|} \sum_{j=2}^{|R|} \frac{\mu'_0(s_i, r_j) + \mu'_0(s_i, r_{j-1})}{d_{ij}^r(\bar{\alpha})} D_{KL, j, s_i}(\mu'_0 || \mu_0),$$

where $d_{ij}^s(\bar{\alpha}) > 0$ denotes the perceptual distance between states (s_i, r_j) and (s_{i-1}, r_j) , and likewise $d_{ij}^r(\bar{\alpha})$ denotes the perceptual distance between states (s_i, r_j) and (s_i, r_{j-1}) . Again, we extend D_N to all $\mu, \mu' \in \mathcal{U}$ with $\mu' \ll \mu$ by continuity.

The perceptual distance $d_{ij}^s(\bar{\alpha})$ captures the idea that two states s_i and s_{i-1} might be easier or more difficult to distinguish in a way that is not captured by $s_i - s_{i-1}$. This is a useful distinction as it allows one to separate the difference in payoffs between two states from the difficulty of distinguishing between those two states.

The Neighborhood Cost Function $D_{S\bar{A}}$. Suppose that D_N is a reduced form divergence arising from a choice over how best to acquire information. In particular, imagine that the agent learns about s in part by observing \bar{a} . To simplify this discussion, we ignore the R dimension by assuming that $|R| = 1$. In this case, we can define, for $j = 1$,

$$d_{ij}^s(\bar{\alpha}) = (s_i - s_{i-1}) \left[1 + \left(\frac{\bar{\alpha}(s_i, r_j) - \bar{\alpha}(s_{i-1}, r_j)}{s_i - s_{i-1}} \right)^2 \right] \quad (13)$$

as the perceptual distance between two states. We denote the neighborhood divergence D_N with this particular perceptual distance as $D_{S\bar{A}}$, to indicate that the agent can observe S and \bar{A} but not R . The perceptual distance $d_{ij}^s(\bar{\alpha})$ captures the idea that if the aggregate action is very different across (s_i, r_1) and (s_{i-1}, r_1) , that is, if $(\bar{\alpha}(s_i, r_1) - \bar{\alpha}(s_{i-1}, r_1))^2$ is large, then it is less costly for the agent to distinguish between states (s_i, r_1) and (s_{i-1}, r_1) .²⁷

This divergence captures the idea that the agent pays attention to the aggregate action and that extreme actions are easier to observe. As a result, the perceptual distance between two states depends on the aggregate action function \bar{a} . It follows immediately that this divergence is not invariant in \bar{A} .²⁸ Monotonicity in R is trivially satisfied under the assumption that $|R| = 1$.

The Neighborhood Cost Function D_{SR} . We can also define versions of the neighborhood-based cost function for which the perceptual distances do not depend on the aggregate action

²⁷In Appendix Section B we derive the continuous-state version of $D_{S\bar{A}}$ from a problem in which the agent optimally allocates her attention over s and \bar{a} ; see also the example constructed in Section 7.4.

²⁸If \bar{a} is not constant, the perturbation $\bar{\alpha}_\epsilon(s, r) = (1 + \epsilon)\bar{\alpha}(s, r) - \epsilon\bar{a}_0$, for any $\bar{a}_0 \in \bar{A}$ and some $\epsilon > 0$, reduces all perceptual distances and hence reduces all information costs.

function but do depend on r . We interpret these cost functions as describing environments in which the agent either cannot or chooses not to learn from the actions of other agents.

As an example, let us return to the assumption that $|R| > 1$, and assume perceptual distances $d_{ij}^r(\bar{\alpha}) = r_j - r_{j-1}$ and $d_{ij}^s(\bar{\alpha}) = s_i - s_{i-1}$. We denote the neighborhood divergence D_N with this particular perceptual distance as D_{SR} .

The D_{SR} divergence, like the KL divergence, does not depend on $\bar{\alpha}$; it is therefore invariant in \bar{A} by construction.²⁹ However, the D_{SR} divergence differs from the KL divergence in that it is not monotone in R (we prove this formally below). Finally, note that it is also possible, at the expense of additional complexity, to construct divergences in this family that both depend on $\bar{\alpha}$ and allow agents to attend to r . More generally, when defining the informational environment one must specify both the set of potential observables (some combination of s , r , and \bar{a} in our framework) and the cost of information given these observables.

Invariance and Blackwell Monotonicity. The special nature of the KL divergence—that it is both monotone in R and invariant in \bar{A} —hints at the relationship between the forms of partial monotonicity and invariance we have defined here and the stronger form of invariance considered by other authors. The literature has focused on divergences that are simply “invariant,” meaning that they are both monotone and invariant with respect to *all* possible transformations of the state space.³⁰ Invariant divergences have been described in the information geometry literature by Chentsov [1982] and Amari and Nagaoka [2007] and employed in another context by Hébert [2018]. Caplin, Dean, and Leahy [2022] show that the standard invariance property of cost functions is equivalent to their invariance-under-compression behavioral axiom.

Our generalization to partial monotonicity and partial invariance allows us to study divergences that are invariant to some but *not necessarily all* partitions of the state space, as the various versions of neighborhood-based costs defined above demonstrate. However, when the divergence is monotone and invariant in all dimensions, then it will be partially monotone and invariant in all dimensions, as in the example of mutual information.

Note that all of the cost functions we consider are Blackwell-monotone (which follows from the convexity of D). As our examples above illustrate, Blackwell-monotonicity of the cost function does not by itself imply a divergence monotone in R or invariant in \bar{A} .

²⁹This point applies to the entire class of posterior-separable cost functions. We could always define such cost functions on $S \times R$, imposing invariance in \bar{A} , or define them on the larger space and not impose invariance in \bar{A} . Only in certain special cases (such as mutual information) will we reach the same conclusion regardless of our assumption; see the remark on general forms of invariance below.

³⁰Our definition of monotonicity is subtly different from what is usually meant by the monotonicity of the KL divergence. The monotonicity property of the KL divergence used in the literature involves comparing KL divergences defined on two different state spaces, whereas our definition of partial monotonicity involves only the operator η , which preserves the state space. Our approach allows us to define partial monotonicity for divergences defined on a single state space as opposed to families of divergences defined on all possible state spaces. Despite this difference, monotonicity in the sense used in the literature implies partial monotonicity in all dimensions. It also implies invariance, whereas partial monotonicity does not imply partial invariance.

3.3 Properties of Cost Functions

The two properties that we focus on for the remainder of the paper are R -monotonicity and \bar{A} -invariance; we will show that these are the key properties related to properties of equilibria. To facilitate this analysis, we next describe how these properties of divergences relate to properties of information cost functions.

Lemma 1. *The posterior-separable cost function $C(\pi, \mu)$ is associated with a divergence D that is monotone in R given $\mu \in \mathcal{U}$ if and only if, for all strategies $\pi \in \Pi(\mu)$, the strategies $\pi' \in \Pi(\mu)$ induced³¹ from π by $(a, \mu') \mapsto (a, \eta_{-R|R}[\mu'|\mu])$ satisfy*

$$C(\pi', \mu) \leq C(\pi, \mu).$$

The posterior-separable cost function $C(\pi, \mu)$ is associated with a divergence D that is \bar{A} -invariant given $\mu \in \mathcal{U}$ if and only if, for all $\mu'' \in \mathcal{U}$ with $\gamma_{-\bar{A}}[\mu] \ll \gamma_{-\bar{A}}[\mu'']$ and all strategies $\pi \in \Pi(\mu)$, the strategies $\pi'' \in \Pi(\eta_{-\bar{A}|\bar{A}}[\mu|\mu''])$ induced from π by $(a, \mu') \mapsto (a, \eta_{-\bar{A}|\bar{A}}[\mu'|\mu''])$ satisfy

$$C(\pi'', \eta_{-\bar{A}|\bar{A}}[\mu|\mu'']) = C(\pi, \mu).$$

Proof. See the Appendix, Section E.2. □

Lemma 1 translates the R -monotonicity and \bar{A} -invariance properties of divergences to properties of information cost functions. Part (i) of the lemma tells us that monotonicity in R is equivalent to the statement that for any strategy, there is a (weakly) less costly strategy that generates the same posteriors on $\Delta(S \times \bar{A})$ in which all conditioning on r is removed. Part (ii) of the lemma tells us that \bar{A} -invariance of the divergence is equivalent to the statement that information costs do not depend on the aggregate action function $\bar{\alpha}$.

The results presented in Lemma 1 suggest how the definitions of R -monotonicity and \bar{A} -invariance can be extended to all cost functions, including those that are not necessarily posterior-separable.³²

Nowhere-Monotone and Nowhere-Invariant Cost Functions. We next define “nowhere” versions of our key properties, what we call “nowhere- R -monotonicity” and “nowhere- \bar{A} -invariance.” These properties imply that R -monotonicity and \bar{A} -invariance, respectively, are violated essentially everywhere.

³¹See footnote 18 for a formal definition of “induced.”

³²Under suitable assumptions on continuity and differentiability, we believe almost all of our results, with the exception of our necessity results for R -monotonicity (Proposition 6), could be readily extended to the general class of cost functions under these definitions. The proof of Proposition 6 relies on explicitly constructing a utility function that induces a particular strategy π , which is facilitated by posterior-separability.

Note, however, that when agents acquire no information, the conditions for R -monotonicity and \bar{A} -invariance hold automatically. Thus, in order to define these nowhere properties, we must first define zero-information strategies.

Definition 8. A strategy $\pi \in \Pi(\mu)$ is **zero-information** if π places full support on action-posterior pairs (a, μ') with $\mu' = \mu$; a strategy is **non-zero-information** otherwise.

Zero information strategies are ones in which agents acquire no information. Armed with this definition, we now provide notions of “nowhere- R -monotonicity” and “nowhere- \bar{A} -invariance; we define these properties directly on cost functions in order to simplify our exposition.³³

Definition 9. A strategy $\pi \in \Pi(\mu)$ is **non- R -measurable** if, for all $(a, \mu') \in \text{supp}(\pi)$, $\eta_{-R|R}[\mu'|\mu] = \mu'$. The posterior-separable cost function $C(\cdot, \mu)$ is **nowhere- R -monotone** given $\mu \in \mathcal{U}$ if, for all non- R -measurable, non-zero-information strategies $\pi' \in \Pi(\mu)$, there exists a strategy $\pi \in \Pi(\mu)$ such that π' is induced from π by $(a, \mu') \mapsto (a, \eta_{-R|R}[\mu'|\mu])$ and for which

$$C(\pi', \mu) > C(\pi, \mu).$$

Non- R -measurable strategies are strategies that could be implemented by observing signals that do not condition on r .³⁴ Loosely speaking, this definition says that if for any strategy π' that collects information and does not condition on r , there is a less-costly strategy π that does condition on r and generates the same posteriors over $S \times \bar{A}$, then the cost function is nowhere- R -monotone. That is, the characterization of R -monotonicity in Lemma 1 fails everywhere, provided the signals are informative.

We turn next to nowhere- \bar{A} -invariance. Let $\text{int}(\mathcal{U}_0)$ denote the interior of the simplex.

Definition 10. The function

$$C_0(\pi_0, \mu_0, \bar{\alpha}) = E^{\pi_0(a, \mu'_0)}[D(\phi\{\mu'_0, \bar{\alpha}\}|\phi\{\mu_0, \bar{\alpha}\})]$$

is **locally invariant in \bar{A}** at $(\pi_0, \bar{\alpha})$ if $\nabla_{\bar{\alpha}} C_0(\pi_0, \mu_0, \bar{\alpha}) = \mathbf{0}$, where $\nabla_{\bar{\alpha}}$ denotes the derivative with respect to $\bar{\alpha}$. The function $C_0(\pi_0, \mu_0, \bar{\alpha})$ is **nowhere- \bar{A} -invariant** given $\mu_0 \in \mathcal{U}_0$ if $\nabla_{\bar{\alpha}} C_0(\pi_0, \mu_0, \bar{\alpha}) \neq \mathbf{0}$ for all non-zero-information π_0 and non-constant $\bar{\alpha}$ in the relative interior of \bar{A} with $\text{supp}(\pi_0) \subset A \times \text{int}(\mathcal{U}_0)$.

Note that D is invariant in \bar{A} if and only if C_0 is locally invariant everywhere. Nowhere- \bar{A} -invariance instead implies a local violation of \bar{A} -invariance at essentially all points.³⁵

³³Specifically, it is possible for a divergence D to be R -monotone nowhere on its domain, and yet for $C(\pi', \mu) \leq C(\pi, \mu)$ to hold (in the context of Lemma 1) at certain points in the parameter space. This case is knife-edge, and defining nowhere- R -monotone in terms of C instead of D avoids the need to consider it.

³⁴That is, the conditional distribution of r given (s, \bar{a}) for each posterior is equal to that of the prior. Such a set of posteriors could have been generated by observing a signal that conditions only on (s, \bar{a}) , but (due to the degeneracy of the prior) could in general also have been generated via other signal structures.

³⁵In our definition of nowhere- \bar{A} -invariance, we allow for local invariance to hold when $\bar{\alpha}$ is constant to account for

Examples. As discussed above, our example cost function $D_{S\bar{A}}$ is not invariant in \bar{A} , and is in fact nowhere- \bar{A} -invariant (see Footnote 28). Likewise, our example cost function D_{SR} is not only non- R -monotone but nowhere- R -monotone for a certain class of priors.

Lemma 2. *Assume $|S| > 1$ and $|R| > 1$. The divergence D_{SR} is not monotone in R , and there are priors for which it is nowhere- R -monotone.*

Proof. See the appendix, E.3. □

The intuition behind this result is as follows. Suppose that a posterior $\mu'_0 \in \mathcal{U}_0$ satisfies, for some $\epsilon > 0$,

$$\mu'_0(s, r) = \mu_0(s, r) f_1(s) (1 + \epsilon(r - \bar{r}(s))),$$

where f_1 is a strictly increasing function of s and $\bar{r}(s) = E^{\mu_0(s, r)}[r|s]$, and suppose for simplicity that $\bar{\alpha}$ is constant across (s, r) . We show in the proof that under these assumptions, if the prior μ_0 satisfies an MLRP condition, so that higher values of s are relatively more likely as r increases, then D_{SR} will violate monotonicity in R for this class of posteriors.

To understand this result, let us take the perspective of an agent who is choosing ϵ but has no particular concern for how her signal correlates with r . Intuitively, because the agent is attempting to differentiate high values of s from low values of s ($f_1(s)$ strictly increasing), and because the expected value of r is increasing in the fundamental ($\bar{r}(s)$ strictly increasing), paying some attention to the public signal is a feasible way of learning about s . Starting from completely ignoring r , paying a little bit of attention to r has a first-order benefit (it reduces the cost of learning via s) but a second-order cost (the marginal cost of learning via r is locally zero near signals that ignore r). Consequently, it is optimal for this agent to attend to r . Nowhere- R -monotonicity follows from essentially the same argument.

4 Equilibrium Existence and Uniqueness

We proceed by establishing equilibrium existence.

Proposition 1. *An equilibrium (Definition 3) exists.*

Proof. See the online appendix, section E.4. This result applies to all continuous utility functions (for a definition of continuity, see the appendix, Definition 13). □

The proof of this result applies the theorem of the maximum and a fixed point theorem in the usual fashion, relying on continuity of payoffs and of costs (Assumption 1).

the possibility that a constant aggregate action maximizes information costs (by preventing the agents from learning the state by observing others' actions). Our example cost function $D_{S\bar{A}}$ has this property. Note by Assumption 1 that $\nabla_{\bar{\alpha}} C_0(\pi_0, \mu_0, \bar{\alpha})$ exists for all $(\pi_0, \bar{\alpha})$ with $\text{supp}(\pi_0) \subset A \times \text{int}(\mathcal{U}_0)$.

We next consider equilibrium uniqueness. Recall that the convexity or concavity of the function G in the mean-critical utility function controls the degree of strategic interaction: when G is strictly convex in \bar{a} , actions are strategic complements, whereas when G is strictly concave in \bar{a} , actions are strategic substitutes. In the linear-quadratic example, $\beta > 0$ indicates strategic complementarity and $\beta < 0$ strategic substitutability. The assumption that $\beta < 1$ guarantees a unique equilibrium under exogenous information.

With endogenous information acquisition, however, we find that a restriction on preferences alone is not sufficient.³⁶ The following proposition provides two conditions that are sufficient (but by no means necessary) to ensure uniqueness of the aggregate action function.

Proposition 2. *If G is strictly concave in $\bar{a} \in \bar{A}$ for all $s \in S$ and the divergence D is invariant in \bar{A} , there is a unique aggregate action function $\bar{\alpha}$ common to all equilibria.*

Proof. See the online appendix, section E.5. □

Relative to the restriction that $\beta < 1$ in the linear-quadratic example, Proposition 2 calls for two additional conditions. One is a stricter condition on payoffs: that actions be strategic substitutes ($\beta < 0$, or more generally G strictly concave). The second is a condition on information costs: that D be invariant in \bar{A} . Together, these conditions essentially rule out strategic complementarity in both payoffs and in costs.

Note that these conditions do not technically ensure a unique equilibrium: they guarantee only the uniqueness of an equilibrium aggregate action function $\bar{\alpha}$. We have not ruled out the possibility that multiple strategies π can be both best responses to and mean-consistent with the same $\bar{\alpha}$. However, for our purposes, this weaker form of uniqueness will suffice.

We should emphasize that strategic complementarity in our context will not always lead to multiplicity—if the convexity of G is not “too strong,” the equilibrium aggregate action function will be unique. In our proof of Proposition 2, we describe a weaker sufficient condition involving g , G , and D .

Sources of Multiplicity. When D is not invariant in \bar{A} , strategic complementarity in information acquisition can arise, opening the door for multiple equilibria $\bar{\alpha}$. To see this, consider the divergence D_{SA} defined above using the perceptual distance in (13). With this cost function, if agents acquire more information, the aggregate action becomes more sensitive to s , generating larger values for $(\bar{\alpha}(s_i, r_j) - \bar{\alpha}(s_{i-1}, r_j))^2$. This reduces the information costs for all agents, increasing their incentive to gather information. In this way agents can face strategic complementarity with respect to information acquisition.

Even in the absence of this form of strategic complementarity (when D is invariant in \bar{A}), strategic complementarity in payoffs combined with endogenous information acquisition can

³⁶In this section, we provide stricter conditions that are sufficient. In Section 7.3, we provide an example of multiplicity in the linear-quadratic-Gaussian setting with $\beta \in (0, 1)$.

still lead to multiplicity as in, e.g., [Hellwig and Veldkamp \[2009\]](#) or [Yang \[2015\]](#). This can be true even when the strategic complementarity in payoffs is too weak to generate multiplicity under exogenous information ($0 \leq \beta < 1$ with the linear-quadratic utility function). Intuitively, with strategic complementarity in payoffs, information is more valuable to an individual agent when other agents are acquiring information and acting upon it. In our linear-quadratic-Gaussian examples, [Section 7.3](#), we construct an example, following [Myatt and Wallace \[2012\]](#), using the KL divergence (which is invariant in \bar{A}) and $\beta = \frac{2}{3}$ in which an equilibrium with no information gathering and an equilibrium with information gathering both exist. [Proposition 2](#) shows that strategic substitutability in payoffs combined with \bar{A} -invariance is sufficient to generate a form of uniqueness; this example illustrates the difficulty of strengthening the result.

5 Efficiency

In this section we consider whether the equilibria of our game coincide with the solution to a planner's problem. To do so, we adapt the constrained efficiency concept of [Angeletos and Pavan \[2007\]](#) to environments with endogenous information acquisition.³⁷ Constrained efficiency is defined as the solution to a planner's problem in which the planner dictates the strategies of the agents with the objective of maximizing ex-ante utility, taking into account the mean-consistency condition and information acquisition costs. We formulate the planner's problem using the "alternative formulation" of the agent's problem ([Definition 4](#)).

Definition 11. Given a common prior $\mu_0 \in \mathcal{U}_0$, a strategy profile $(\pi_0, \bar{\alpha}) \in \Pi_0(\mu_0) \times \bar{\mathcal{A}}$ is **constrained efficient** if it solves

$$\max_{\bar{\alpha} \in \bar{\mathcal{A}}, \pi_0 \in \Pi_0(\mu_0)} E^{\pi_0} [V(a, \phi[\mu'_0, \bar{\alpha}]) - D(\phi[\mu'_0, \bar{\alpha}] || \phi[\mu_0, \bar{\alpha}])]$$

subject to mean-consistency, $\bar{\alpha}(s, r) = E^{\pi_0}[a|s, r]$, $\forall (s, r) \in S \times R$. A strategy profile $(\pi, \bar{\alpha}) \in \Pi(\phi[\mu_0, \bar{\alpha}]) \times \bar{\mathcal{A}}$ is constrained efficient if it is induced from a constrained efficient $(\pi_0, \bar{\alpha})$ by $(a, \mu'_0) \mapsto (a, \phi[\mu'_0, \bar{\alpha}])$.

A constrained efficient strategy profile maximizes welfare subject to information acquisition costs. The key distinction between the planner's problem and an equilibrium is that the planner takes into account how the individual agents' actions affect the aggregate action function $\bar{\alpha}$, and consequently welfare, whereas individual agents do not.

The aggregate action function $\bar{\alpha}$ can in theory affect welfare in two distinct ways. The first is through payoffs, $V(a, \phi[\mu'_0, \bar{\alpha}])$; the second is through information costs, $D(\phi[\mu'_0, \bar{\alpha}] || \phi[\mu_0, \bar{\alpha}])$.

³⁷See also [Radner \[1962\]](#) and [Vives \[1988\]](#). [Angeletos and Pavan \[2007\]](#) define constrained efficiency in an environment with exogenous, incomplete information; the planner in that setting dictates how agents should use their information. In our setting with endogenous information acquisition, the planner dictates what information to acquire as well as what action to take given that information.

We discuss both channels below.

Payoff Externalities. Observe by the definition of V that

$$E^{\pi_0}[V(a, \phi[\mu'_0, \bar{\alpha}])] = E^{\mu_0}[E^{\pi_0}[v(a, \bar{\alpha}(s, r), s)|s, r]].$$

A payoff externality appears to exist in our setting: the aggregate action function $\bar{\alpha}$ enters the utility function $v(\cdot)$ via the latter's dependence on \bar{a} . Each agent, when choosing her own action, does not internalize how her action affects the aggregate action and hence the payoffs of others.

However, for the mean-critical class of payoff functions assumed in (1), if $(\pi_0, \bar{\alpha})$ is a mean-consistent strategy profile, then

$$\frac{\partial}{\partial \bar{a}} E^{\pi_0}[v(a, \bar{a}, s)|s, r]|_{\bar{a}=\bar{\alpha}(s, r)} = E^{\pi_0}[(a - \bar{\alpha}(s, r))|s, r] \frac{\partial^2}{\partial \bar{a}^2} G(\bar{a}, s)|_{\bar{a}=\bar{\alpha}(s, r)} = 0.$$

That is, in equilibrium, small changes to the aggregate action function $\bar{\alpha}$ have no direct effect on the expected payoff of agents. Mean-criticality thus rules out payoff externalities.

Angeletos and Pavan [2007] develop this result in the context of exogenous signals in linear-quadratic games and call this property “efficiency in the use of information.” Our generalization of the Angeletos and Pavan [2007] condition to mean-critical payoffs is potentially of interest on its own. In Appendix Section A, we describe the class of utility functions with this property; we furthermore provide an interpretation of (1) in terms of production economies with complete markets, connecting our efficiency results to the results of Angeletos and Sastry [2021]. Mean-criticality in our context allows us to focus the following analysis on externalities arising in endogenous information acquisition.

Informational Externalities. Information costs can depend on $\bar{\alpha}$, as with the example cost function $D_{S\bar{A}}$. Such a channel is natural if we imagine that agents learn in part by observing other agents' actions, and that some actions are easier to observe than other actions. When this is the case, agents do not internalize how their own strategies affect the information acquisition costs of others and, as a result, an externality arises. The planner, who takes this informational externality into account, can achieve a higher level of welfare by dictating a strategy that puts relatively more weight on easily observable actions, effectively lowering information acquisition costs.³⁸ We construct an example illustrating this point in the linear-quadratic-Gaussian setting (Section 7.1). In this example, more extreme actions are more easily observed, and as a result

³⁸A related phenomena occurs in games with incomplete, exogenous information, when there exists a signal about an endogenous, equilibrium object such as a price; see, e.g. Laffont [1985], Angeletos and Pavan [2009], Amador and Weill [2010], Vives [2017], Angeletos, Iovino, and La'O [2020]. When agents observe signals about endogenous objects, an information-aggregation externality can arise: agents do not take into account how their own use of information affects the information content of these signals. In these environments, the planner may wish for agents to use their information in a way that differs from what is privately optimal, in order to improve the aggregation of information.

the planner would prefer that each agent choose a more extreme action given her beliefs than the agent would choose for herself.

We find that the key to generating these informational externalities is a local violation of invariance in \bar{A} .

Proposition 3. *A constrained efficient strategy profile $(\pi_0, \bar{\alpha})$ with $\text{supp}(\pi_0) \subset A \times \text{int}(\mathcal{U}_0)$ is an equilibrium if and only if C_0 is locally invariant in \bar{A} at $(\pi_0, \bar{\alpha})$.*

Proof. See the appendix, E.7. □

If the gradient of the cost function with respect to $\bar{\alpha}$ is not equal to zero at the planner's optimum, then the planner and the agents' incentives do not coincide. In this case, the solution to the planner's problem is not an equilibrium.

If a divergence guarantees efficiency for all mean-critical utility functions, it must be locally invariant at all interior constrained efficient strategy profiles. This does not require local invariance at all strategy profiles; for example, strategy profiles that are not mean-consistent are not constrained efficient. Likewise, a divergence can generate inefficiency for all mean-critical utility functions without violating local invariance at all possible strategy profiles. However, subject to a few caveats, it follows from this result that nowhere- \bar{A} -invariance and \bar{A} -invariance are sufficient to guarantee, respectively, inefficiency and constrained efficiency. We prove these results below.

Constrained Inefficiency: Sufficient Conditions. When local \bar{A} -invariance is violated at essentially all points, as in our example $D_{S\bar{A}}$, all equilibria are inefficient, with two possible exceptions. The first exception is the special case in which there is a solution to the social planner's problem that involves a constant $\bar{\alpha}$ function (i.e. the aggregate action function does not vary with the states). This includes equilibria in which no information is acquired; it is possible, though not guaranteed, that such equilibria are constrained efficient.³⁹ The second possible exception is if the agents choose posteriors that lie on the boundary of the simplex. In this case, it is not possible for them to acquire additional information (at least in certain directions), and as a result the equilibrium might be efficient despite a violation of local invariance.

Proposition 4. *If the cost function C_0 is nowhere- \bar{A} -invariant, then there is no constrained efficient equilibrium $(\pi_0, \bar{\alpha})$ in which $\bar{\alpha}$ is not constant and $\text{supp}(\pi_0) \subset A \times \text{int}(\mathcal{U}_0)$.*

Proof. See the appendix, E.8. The proof shows that under these conditions $\bar{\alpha}$ lies in the relative interior of \bar{A} , and the result follows from the definition of nowhere- \bar{A} -invariance (Definition 10) and Proposition 3. □

³⁹This could also occur if agents condition the higher moments, but not the mean, of their actions on the state.

Under certain conditions, these exceptions can be ruled out. Specifically, if the scale of the utility function relative to the information costs is such that agents choose to acquire some information without entirely ruling out any possible states (so that all posteriors lie within the interior), and it is socially optimal for the aggregate action to vary with the fundamental state $s \in S$, then only inefficient equilibria will exist.

Constrained Efficiency: Sufficient Conditions. The informational externality, however, is ruled out if we assume that agents can only pay attention to exogenous states, as with our example D_{SR} cost function. It is also ruled out if agents can learn by observing endogenous actions, but learning via the aggregate action and arriving at the same information without observing the aggregate action are equally costly, as with D_{KL} . Both of these scenarios implicitly assume a cost function that is invariant in \bar{A} and therefore rule out externalities in information acquisition.⁴⁰ In this case, the first-order conditions of the planner and the agents coincide.

Proposition 5. *If the divergence D is invariant in \bar{A} , then a constrained efficient equilibrium exists. If in addition G is weakly concave in \bar{a} for all $s \in S$, then all equilibria are constrained efficient.*

Proof. See the appendix, E.9. □

To prove this result, we show that a solution to the planner’s problem exists. Because the first-order conditions of the planner and agents coincide, this solution is also an equilibrium. Thus, \bar{A} -invariance, by eliminating informational externalities, is sufficient for the existence of a constrained efficient equilibrium. In Section 7.1, we provide an example illustrating this result using the KL divergence.

The argument above in fact shows something stronger: if D is invariant in \bar{A} , every equilibrium satisfies the first-order conditions of the planner’s problem. This does not mean every equilibrium is constrained efficient—some equilibria might be minima, saddle points, or local but not global maxima of the planner’s problem.

However, if the planner’s problem is concave, then every such point is a global maximum, and as a result every equilibrium is efficient. A sufficient (but not necessary) condition for concavity of the planner’s problem is concavity of G . By Proposition 2, if G is strictly concave, the aggregate action function is the same in all equilibria, and as a result all equilibria are equivalent from the planner’s perspective. When G is only weakly concave, we cannot rule out equilibria with different aggregate action functions, but all such equilibria generate the same expected utility.

⁴⁰This has been shown in some specific contexts, e.g. the online appendix of Angeletos and La’O [2020], the analysis in Angeletos and Sastry [2021] of fundamental equilibria, or the analysis of Colombo, Femminis, and Pavan [2014] when there is no effect of action dispersion on payoffs.

6 Non-Fundamental Volatility

We next study the conditions under which an equilibrium varies with the public signal $r \in R$. We first summarize our results for this section before providing formal statements below.

We have emphasized the interpretation of $r \in R$ as the realization of a public signal. This interpretation highlights the dual role r plays in our model: it is both informative about s and has a common realization across agents.

Because it is costly for agents to condition their actions on r , the assumption that it is informative about s is not sufficient on its own to induce agents to condition on it. In the absence of strategic considerations (G linear), agents need only consider whether conditioning on r is sufficiently cheap given its informational content. Nowhere- R -monotonicity can be interpreted as saying that conditioning on r is always a cheaper way of obtaining information about (s, \bar{a}) than conditioning on (s, \bar{a}) alone. Monotonicity in R is the opposite property: it is always weakly more expensive to acquire information by conditioning on r than by directly conditioning on (s, \bar{a}) . As a result, absent strategic considerations, nowhere- R -monotonicity generates equilibria that feature non-fundamental volatility, whereas R -monotonicity generates at least one equilibrium that does not exhibit non-fundamental volatility.

In the presence of strategic considerations, the fact that r is common across agents allows it to serve as a coordination device. In the case of strategic substitutability, the public nature of the signal deters agents from relying on its information; in the case of strategic complementarity, it has the opposite effect. However, neither of these effects are operative if other agents are not conditioning on r . As a result, nowhere- R -monotonicity still guarantees that all equilibria exhibit non-fundamental volatility, and R -monotonicity still guarantees the existence of an equilibrium with zero non-fundamental volatility.

We find that R -monotonicity is not on its own sufficient to ensure that all equilibria exhibit zero non-fundamental volatility. It does, however, imply that if there is a unique aggregate action function across equilibria, then that aggregate action function must be s -measurable. As discussed in Section 4, this kind of uniqueness will hold if D is invariant in \bar{A} and the convexity of G is not “too strong.” However, even if uniqueness fails due to strategic complementarity, it is still possible that all equilibria are devoid of non-fundamental volatility. In our linear-quadratic-Gaussian examples (Section 7), this occurs with the KL divergence cost function whenever $\beta \in (1/2, 1)$, highlighting that multiplicity does not on its own imply non-fundamental volatility in the presence of costly coordination devices.⁴¹

⁴¹In the presence of free to observe and completely uninformative coordination devices (i.e. sunspots), multiplicity implies non-fundamental volatility by the usual arguments. Sunspots are an edge case in our framework: if r is slightly costly to observe and uninformative, R -monotonicity holds, whereas if r is free and slightly informative, R -monotonicity fails.

Measurable Equilibria. We first define a notion of measurability. We say that an equilibrium is s -measurable if the strategies and the aggregate action function do not vary with $r \in R$.

Definition 12. An aggregate action function $\bar{\alpha}$ is **s -measurable** if, for all $s \in S$ and $r, r' \in R$, $\bar{\alpha}(s, r) = \bar{\alpha}(s, r')$. A prior $\mu \in \mathcal{U}$ is **s -measurable** if $\mu = \phi[\mu_0, \bar{\alpha}]$ for some s -measurable $\bar{\alpha} \in \bar{\mathcal{A}}$. An equilibrium $(\pi, \bar{\alpha})$ is **s -measurable** if π is non- R -measurable and $\bar{\alpha}$ is s -measurable.

When the aggregate action varies with $r \in R$ conditional on $s \in S$, it varies with the payoff-irrelevant states $e \in \mathcal{E}$; it thereby exhibits non-fundamental volatility. When instead $\bar{\alpha}$ is s -measurable, it exhibits zero non-fundamental volatility.

S -measurability of an equilibrium is a somewhat stronger statement in that it requires not only that the aggregate action be s -measurable but also that equilibrium strategies π be non- R -measurable.⁴² In what follows, we will consider under what conditions equilibria are or are not s -measurable in this stronger sense.

Non- s -measurable Equilibria. We next provide a condition under which all equilibria are non- s -measurable. We show that when R -monotonicity fails essentially everywhere, as with the example cost function D_{SR} , then all equilibria either feature zero information acquisition or are non- s -measurable.

Proposition 6. *If C is nowhere- R -monotone for all $\mu \in \mathcal{U}$ such that $\gamma_{-\bar{A}}[\mu] = \mu_0$, then all equilibria are either not s -measurable or have zero information acquisition.*

Proof. See the online appendix, section E.11. □

Consider an agent with a cost function that is nowhere-monotone in R . Even if all other agents play strategies that generate an s -measurable $\bar{\alpha}$, the best response of this agent is a strategy that conditions on $r \in R$. This is not because the agent cares about r per se, but rather because the agent finds it informationally less-costly to condition actions on r (Lemma 1). It follows that when all agents have nowhere- R -monotone cost functions, an s -measurable equilibrium in which agents collect information cannot exist.

In this case, the exogenous states $r \in R$ resemble the noisy public signal realizations featured in the exogenous information literature. Agents find it cheaper to “pay attention to” or condition their actions on r , introducing correlated noise in actions and generating non-fundamental volatility in equilibrium. We provide an example of such non-fundamental volatility in the linear-quadratic-Gaussian setting (Section 7.2).

⁴²It is possible, in special cases, for an equilibrium to have an s -measurable $\bar{\alpha}$ and yet for the strategy π to condition on r , despite the requirement of mean-consistency. For example, the second and higher moments of the agent’s action could vary in r conditional on s , while the mean conditional on s does not.

Existence of s -measurable Equilibria. We next provide a necessary and sufficient condition for the existence of an s -measurable equilibrium. Recall that the “ s -measurable priors” are the endogenous priors that are induced from the exogenous prior μ_0 by an s -measurable $\bar{\alpha}$.

Proposition 7. *An s -measurable equilibrium exists for all mean-critical utility functions if and only if D is R -monotone on all s -measurable priors.*

Proof. See the appendix, E.12. This result applies to all continuous utility functions. □

If the divergence is D is R -monotone on all s -measurable priors, then an s -measurable equilibrium exists. Consider a single agent with a cost function that is monotone in R . Recall from Lemma 1 that if the divergence is monotone in R , then for any strategy, there is a non- R -measurable strategy that generates the same posteriors on $\Delta(S \times \bar{A})$ and is less costly. This implies that, because the agent does not care about r per se, only how it affects \bar{a} , if all other agents play strategies that generate an s -measurable $\bar{\alpha}$, the individual agent’s set of best responses always includes a strategy that is mean-consistent with an s -measurable aggregate action function. By the same fixed-point argument employed earlier to prove equilibrium existence, it can be shown that an s -measurable equilibrium exists.

Note that this result does not require R -monotonicity given *all* priors; it is sufficient for the divergence to be monotone in R on all s -measurable priors. If instead R -monotonicity fails on some s -measurable prior, then there is a utility function for which all equilibria are non- s -measurable.

Combining Propositions 2 and 7, in the case of strategic substitutability (G strictly concave) and invariance in \bar{A} , the aggregate action function $\bar{\alpha}$ is unique and thus must be s -measurable.

Interpretation. In exogenous information environments, noisy public signals or, more generally, correlated errors in beliefs, are typical components of information structures [Bergemann and Morris, 2013]. Agents rely on public signals as a source of information as well as a method of coordinating their actions with other agents. Non-fundamental volatility arises in equilibrium as a by-product of public signal noise.

In environments with endogenous information acquisition, however, what appears to be an appealing property to impose on cost structures—monotonicity in R —leads to a stark result. If agents have no particular reason to pay attention to $r \in R$, and paying attention to r only increases their information costs, then an equilibrium exists in which all agents ignore r . This equilibrium features zero non-fundamental volatility.⁴³

Consider the case of strategic complementarity, and assume invariance in \bar{A} . We loosely interpret our analysis as describing three (non-exhaustive) possibilities. If observing r is a cheap

⁴³In a model with a large but finite number of agents, idiosyncratic errors across agents do not exactly wash out in the aggregate. By focusing on the continuum limit, we implicitly assume that this kind of non-fundamental volatility is negligible.

way to observe s and allows agents to coordinate, all equilibria must exhibit non-fundamental volatility. If observing r is a costly way to observe s , and the costs outweigh the possible coordination benefits, then no equilibria will exhibit non-fundamental volatility. If observing r is a costly way to observe s , but the coordination benefits are substantial relative to the costs of observing r , there can be equilibria with and without non-fundamental volatility.

We have emphasized the interpretation of $r \in R$ as the realization of a public signal, and for this reason view it as natural to assume that s and r are dependent under the prior μ_0 . However, it is equally valid to interpret r as pure noise (i.e. like the $e \in \mathcal{E}$ introduced above), in which case s and r are independent under the prior. Under this interpretation, D will be R -monotone if incorporating in the noise of r weakly increases information costs, and will not be R -monotone if incorporating in the noise of r reduces information costs. That is, non- R -monotonicity can again be interpreted as equivalent to assuming that agents can reduce costs by paying attention to public signals. The key driver of non-fundamental volatility is R -monotonicity or the lack thereof, and this is independent of whether one interprets r as informative about s or as noise.

Non-fundamental volatility and efficiency. Building on our earlier results, if a divergence is invariant in \bar{A} and but not monotone in R on s -measurable priors, there exists a utility function such that all equilibria are efficient, share a common $\bar{\alpha}$, and are not s -measurable. Our example cost function D_{SR} satisfies these conditions, demonstrating that efficiency and non-fundamental volatility are driven by distinct properties of information costs.

Corollary 1. *If D is invariant in \bar{A} but not monotone in R given some $\mu = \phi[\mu_0, \bar{\alpha}]$ with $\bar{\alpha}$ s -measurable, there exists a mean-critical utility function u such that equilibria exist and all equilibria are constrained efficient, non- s -measurable, and share a common aggregate action function $\bar{\alpha}^*$.*

Proof. See the appendix, E.13. □

7 The Linear-Quadratic Gaussian Setting

In this section we give examples illustrating our general results adapted to the linear-quadratic-Gaussian (LQG) setting. These examples can be understood using first order conditions in a familiar setting; they are not intended as a second set of formal results.

We use the linear-quadratic utility function of (2), but now assume that A , S , and R are the real lines.⁴⁴ We assume a Gaussian prior μ_0 under which $(s, r) \sim \mathcal{N}((0, 0), \Sigma_0)$, and focus on linear equilibria in which $\bar{\alpha}(s, r) = \bar{\alpha}_s s + \bar{\alpha}_r r$ for some constants $(\bar{\alpha}_s, \bar{\alpha}_r) \in \mathbb{R}^2$.

⁴⁴Assuming finite $S \times R$ and compact A avoids technicalities when deriving our general results; restricting attention to the Gaussian setting can be seen as an alternative means of maintaining tractability.

For information costs, we consider the continuous-state analogs of D_{KL} , $D_{S\bar{A}}$, and D_{SR} introduced in Section 3, described below. With these particular cost functions, linear-quadratic payoffs, Gaussian priors, and a linear aggregate action function, rationally inattentive agents optimally form Gaussian posteriors; see results in Sims [2003] and Hébert and Woodford [2021].⁴⁵ By this we mean that the agent’s optimal strategy $\pi_0 \in \Pi_0(\mu_0)$ will be “as if” the agent observes a one-dimensional Gaussian signal $\omega^i \in \mathbb{R}$, updates her beliefs, and then takes the optimal action given those beliefs. The signal ω^i can be represented as

$$\omega^i = (s, r) \cdot \lambda_0 + \varepsilon^i, \quad \text{with} \quad \varepsilon^i \sim \mathcal{N}(0, \tau_0^{-1}), \quad (14)$$

where $\lambda_0 \in \mathbb{R}^2$ is a weighting vector on s vs. r and τ_0 is the precision of this as-if signal. The objects (λ_0, τ_0) are the choice variables of the agent.

Given the utility function in (2), the optimal action given any posterior $\mu'_0 \in \mathcal{U}_0$ is

$$a^*(\mu'_0) = E^{\mu'_0}[(1 - \beta)s + \beta(\bar{\alpha}_s s + \bar{\alpha}_r r)] = E^{\mu'_0}[(s, r) \cdot \psi(\bar{\alpha}_s, \bar{\alpha}_r)],$$

where $\psi(\bar{\alpha}_s, \bar{\alpha}_r)$ is a column vector given by:

$$\psi(\bar{\alpha})' \equiv (1 - \beta + \beta\bar{\alpha}_s, \beta\bar{\alpha}_r). \quad (15)$$

We can associate each ω^i with a posterior $\mu'_0(\omega^i)$ defined by Bayesian updating and an optimal action $a(\omega^i) = \alpha(\lambda_0, \tau_0, \bar{\alpha}_s, \bar{\alpha}_r)\omega^i$, where

$$\alpha(\lambda_0, \tau_0, \bar{\alpha}_s, \bar{\alpha}_r) = \frac{\tau_0}{1 + \tau_0\lambda_0'\Sigma_0\lambda_0} \psi(\bar{\alpha}_s, \bar{\alpha}_r)'\Sigma_0\lambda_0.$$

See the technical appendix, Section D, for a more detailed derivation. The agent’s optimal strategy $\pi_0 \in \Pi_0(\mu_0)$ is induced from the distribution of signals by $\omega^i \mapsto (a(\omega^i), \mu'_0(\omega^i))$.

Mean-consistency in this context requires that $(\bar{\alpha}_s, \bar{\alpha}_r)' = \alpha(\lambda_0, \tau_0, \bar{\alpha}_s, \bar{\alpha}_r)\lambda_0$, which is

$$(\bar{\alpha}_s, \bar{\alpha}_r)' = \frac{\tau_0\psi(\bar{\alpha}_s, \bar{\alpha}_r)'\Sigma_0\lambda_0}{1 + \tau_0\lambda_0'\Sigma_0\lambda_0} \lambda_0. \quad (16)$$

Example Cost Functions. We will consider the continuous-state analogs of the example divergences D_{KL} , $D_{S\bar{A}}$, and D_{SR} introduced in Section 3. All three cost functions imply that Gaussian signals of the form in (14) are optimal. As such, each cost function can be written in reduced-form as a cost $C(\lambda_0, \tau_0, \Sigma_0, \bar{\alpha}_s, \bar{\alpha}_r)$. See the technical appendix, Section D, for derivations of the costs below from results in Sims [2003] and Hébert and Woodford [2021].

⁴⁵An alternative approach, common in the literature (see, e.g., Colombo, Femminis, and Pavan [2014]), is to impose Gaussian signals and define information costs on these signals.

Consider first the mutual information cost function with associated KL divergence D_{KL} :

$$D_{KL}(\mu' || \mu) = E^{\mu'} \left[\ln \left(\frac{d\mu'}{d\mu} \right) \right],$$

The cost function in this case can be written, for some constant $\theta > 0$, as

$$C_{MI}(\lambda_0, \tau_0, \Sigma_0, \bar{\alpha}_s, \bar{\alpha}_r) = \theta \ln(1 + \tau_0 \lambda_0' \Sigma_0 \lambda_0).$$

Second, consider the divergence D_{SR} . By results in [Hébert and Woodford \[2021\]](#), the continuous state limit of D_{SR} is a version of the Fisher information, given by:

$$D_{FI-SR}(\mu' || \mu) = E^{\mu'} \left[\left(\frac{\partial}{\partial s} \ln \left(\frac{d\mu'}{d\mu} \right) \right)^2 + \left(\frac{\partial}{\partial r} \ln \left(\frac{d\mu'}{d\mu} \right) \right)^2 \right].$$

The cost function can be written, for some constant $\theta > 0$, as

$$C_{FI-SR}(\lambda_0, \tau_0, \Sigma_0, \bar{\alpha}_s, \bar{\alpha}_r) = \theta \tau_0 \lambda_0' \lambda_0.$$

Third, consider the divergence $D_{S\bar{A}}$. By results in [Hébert and Woodford \[2021\]](#), the continuous state limit of $D_{S\bar{A}}$ is a version of the Fisher information, given by:

$$D_{FI-S\bar{A}}(\mu' || \mu) = E^{\mu'} \left[\frac{1}{1 + \bar{\alpha}_s^2} \left(\frac{\partial}{\partial s} \ln \left(\frac{d\mu'}{d\mu} \right) \right)^2 \right].$$

The cost function can be written, for some constant $\theta > 0$, as

$$C_{FI-S\bar{A}}(\lambda_0, \tau_0, \Sigma_0, \bar{\alpha}_s, \bar{\alpha}_r) = \frac{\theta}{1 + \bar{\alpha}_s^2} \tau_0 \lambda_0' \lambda_0.$$

Both C_{FI-SR} and $C_{FI-S\bar{A}}$ are linear in the precision τ_0 ; such linearity is commonly assumed in the LQG literature, see, e.g, [Myatt and Wallace \[2012\]](#) or [Van Nieuwerburgh and Veldkamp \[2010\]](#). Note that with $C_{FI-S\bar{A}}$, the cost of any non-zero pair (λ_0, τ_0) is decreasing in $|\bar{\alpha}_s|$ provided that $|\bar{\alpha}_s| > 0$.

We should caution that the signal $\omega^i \in \mathbb{R}$ should not be interpreted as describing the *actual* signal received by the agent. In our setting, the agent might learn about s and r either by paying attention to those objects directly or by paying attention to \bar{a} ; in either case, it is “as if” the agent receives a Gaussian signal of the form in (14). The cost functions C should be understood in a reduced form sense: they represent the cost of reaching the posteriors via the least costly method available. We discuss this issue in more detail and provide an example below.

Equilibrium. The agent's expected payoff given the posterior μ'_0 can be written as

$$E^{\mu'_0}[v(a^*(\mu'_0), \bar{\alpha}(s, r), s)] = -Var^{\mu'_0}[(s, r) \cdot \psi(\bar{\alpha}_s, \bar{\alpha}_r)] - \beta(1 - \beta)E^{\mu'_0}[(s - \bar{\alpha}_s s - \bar{\alpha}_r r)^2]$$

where $Var[\cdot]$ denotes the variance. For the Gaussian strategies described above, this posterior variance is identical across realizations of ω ; a little algebra shows that the agent's ex-ante expected payoff is

$$E^{\tau_0}[u(a, \bar{a}, s)] = \frac{\tau_0}{1 + \tau_0 \lambda'_0 \Sigma_0 \lambda_0} (\psi(\bar{\alpha}_s, \bar{\alpha}_r)' \Sigma_0 \lambda_0)^2 + f(\bar{\alpha}_s, \bar{\alpha}_r),$$

where $f(\bar{\alpha}) \equiv -\psi(\bar{\alpha}_s, \bar{\alpha}_r)' \Sigma_0 \psi(\bar{\alpha}_s, \bar{\alpha}_r) - \beta(1 - \beta)E^{\mu_0}[(s - \bar{\alpha}_s s - \bar{\alpha}_r r)^2]$ are payoff terms that are independent of the agent's individual strategy (λ_0, τ_0) . See the technical appendix, Section D, for details. A profile $(\lambda_0^*, \tau_0^*, \bar{\alpha}_s^*, \bar{\alpha}_r^*)$ is a linear equilibrium if mean consistency (16) holds and

$$(\lambda_0^*, \tau_0^*) \in \arg \max_{\lambda_0 \in \mathbb{R}^2, \tau_0 \geq 0} \frac{\tau_0}{1 + \tau_0 \lambda'_0 \Sigma_0 \lambda_0} (\psi(\bar{\alpha}_s^*, \bar{\alpha}_r^*)' \Sigma_0 \lambda_0)^2 - C(\lambda_0, \tau_0, \Sigma_0, \bar{\alpha}_s^*, \bar{\alpha}_r^*).$$

7.1 Efficiency

We first consider the question of constrained efficiency. To simplify the analysis, we abstract from the r dimension by assuming that the variance of r is zero under the prior μ_0 , setting $\Sigma_0 = \begin{bmatrix} \sigma_s^2 & 0 \\ 0 & 0 \end{bmatrix}$. We thereby need only consider the s dimension. We focus on linear aggregate action functions $\bar{\alpha}(s) = \bar{\alpha}_s s$ and adopt the normalization $\lambda'_0 = (1, 0)$.

Consider the planner's problem (restricted to linear aggregate actions $\bar{\alpha}$ and strategies of the form described above) with the mutual information cost function,

$$\max_{\tau_0 \geq 0, \bar{\alpha}_s} (1 - \beta + \beta \bar{\alpha}_s)^2 \sigma_s^2 \left(1 - \frac{1}{1 + \tau_0 \sigma_s^2} \right) + f(\bar{\alpha}) - \theta \ln(1 + \tau_0 \sigma_s^2),$$

where $f(\bar{\alpha}) = -(1 - \beta + \beta \bar{\alpha}_s)^2 \sigma_s^2 - \beta(1 - \beta)(1 - \bar{\alpha}_s)^2 \sigma_s^2$, subject to mean-consistency:

$$\bar{\alpha}_s = (1 - \beta + \beta \bar{\alpha}_s) \left(1 - \frac{1}{1 + \tau_0 \sigma_s^2} \right). \quad (17)$$

Note that the mutual information cost $\theta \ln(1 + \tau_0 \sigma_s^2)$ does not depend on the coefficient $\bar{\alpha}_s$; this is a manifestation of invariance in \bar{A} .

Let us for a moment ignore the mean-consistency constraint. The first-order condition of the planner's objective with respect to $\bar{\alpha}_s$ is

$$2\beta \sigma_s^2 (1 - \beta + \beta \bar{\alpha}_s) \left(1 - \frac{1}{1 + \tau_0 \sigma_s^2} \right) = 2\beta (1 - \beta + \beta \bar{\alpha}_s) \sigma_s^2 - 2\beta (1 - \beta) (1 - \bar{\alpha}_s) \sigma_s^2,$$

which simplifies to (17). That is, the mean-consistency constraint holds with equality at the planner's unrestricted optimum. It follows that the solution to the planner's problem is also an equilibrium (as the FOCs of the unrestricted planner and agent with respect to τ_0 are identical). This seemingly remarkable coincidence is, of course, a consequence of the mean-critical nature of the payoff function and the invariance of mutual information with respect to \bar{A} . This result is the LQG analog of Proposition 5.

Consider now the problem with the cost function $C_{FI-S\bar{A}} = \theta(1 + \bar{\alpha}_s^2)^{-1}\tau_0$. To simplify the argument, we set $\beta = 0$; mean-consistency is $\bar{\alpha}_s^* = 1 - (1 + \tau_0^*\sigma_s^2)^{-1}$. Note that this cost function is nowhere- \bar{A} -invariant, as it has a non-zero derivative with respect to $\bar{\alpha}_s$ wherever $\bar{\alpha}_s \neq 0$. The agent's problem is

$$\max_{\tau_0 \geq 0} \sigma_s^2 \left(1 - \frac{1}{1 + \tau_0 \sigma_s^2} \right) - \frac{\theta}{1 + \bar{\alpha}_s^2} \tau_0.$$

This problem is concave in τ_0 ; the agent's first-order condition with respect to τ_0 is given by:

$$\frac{(\sigma_s^2)^2}{(1 + \tau_0^* \sigma_s^2)^2} \leq \frac{\theta}{1 + (\bar{\alpha}_s^*)^2},$$

with equality if $\tau_0^* > 0$. By mean-consistency, if $\bar{\alpha}_s^* > 0$,

$$(\sigma_s^2)^2 (1 - \bar{\alpha}_s^*)^2 (1 + (\bar{\alpha}_s^*)^2) = \theta. \quad (18)$$

Consider now a planner's problem in this context,

$$\bar{\alpha}_s^{**} = \arg \max_{\bar{\alpha}_s \in [0,1]} \sigma_s^2 \bar{\alpha}_s - \frac{\theta}{1 + \bar{\alpha}_s^2} \frac{1}{\sigma_s^2} \left(\frac{1}{1 - \bar{\alpha}_s} - 1 \right),$$

where we have substituted the mean-consistency constraint into the planner's objective. The first-order condition is given by

$$(\sigma_s^2)^2 (1 + \bar{\alpha}_s^{**})^2 (1 + (\bar{\alpha}_s^{**})^2) \leq \theta \left(1 - \frac{2\bar{\alpha}_s^{**}}{1 + (\bar{\alpha}_s^{**})^2} \right),$$

with equality if $\bar{\alpha}_s^{**} > 0$. This condition coincides with the equilibrium condition (18) only if no information is gathered, in which case $\bar{\alpha}_s^{**} = \bar{\alpha}_s^* = 0$. This occurs when θ is sufficiently high: $\theta \geq (\sigma_s^2)^2$.⁴⁶ When instead $\theta < (\sigma_s^2)^2$, all equilibria are inefficient: $\bar{\alpha}_s^{**} > \bar{\alpha}_s^*$. This result is driven by the cost function's nowhere- \bar{A} -invariance and is the LQG analog of Proposition 4.

Specifically, in this example, the cost of any precision $\tau_0 > 0$ is a decreasing function of $\bar{\alpha}_s$, the sensitivity of the aggregate action to the underlying state s . More extreme aggregate actions are thus easier or "cheaper" for agents to observe than less extreme aggregate actions;

⁴⁶This point is not specific to the LQG setting. When the solution to the planner's problem involves no information gathering, the aggregate action $\bar{\alpha}$ cannot affect information costs. In this case, the solution to the planner's problem is always an equilibrium provided that the utility function is mean-critical at the relevant point.

for this reason, the cost of information depends on $\bar{\alpha}_s$. Individuals agents do not internalize how their actions and information choices affect $\bar{\alpha}_s$, which leads in equilibrium to inefficiency. A benevolent planner, in an effort to reduce information costs for all agents, would dictate more information acquisition for each agent, resulting in a more sensitive aggregate action function.⁴⁷

7.2 Non-Fundamental Volatility

We next consider the question of non-fundamental volatility. To simplify our analysis, we assume $s \sim \mathcal{N}(0, \sigma_s^2)$ and that r is a noisy public signal about s , given by $r = s + e$, where $e \sim \mathcal{N}(0, \sigma_e^2)$ is independent of s and $\sigma_e^2 > 0$. With this specification, $\Sigma_0 = \begin{bmatrix} \sigma_s^2 & \sigma_s^2 \\ \sigma_s^2 & \sigma_s^2 + \sigma_e^2 \end{bmatrix}$.

Note that (λ_0, τ_0) and $(m\lambda_0, m^{-2}\tau_0)$ for some $m > 0$ are signals that convey identical information (one is just a scaled version of the other). For the purposes of this sub-section, it is convenient to adopt the normalization $\lambda_0' \Sigma_0 \lambda_0 = 1$. Consider the Lagrangian version of the agent's problem with the mutual information cost function C_{MI} with this constraint,

$$\max_{\tau_0 \geq 0, \lambda_0 \in \mathbb{R}^2} \min_{\nu \in \mathbb{R}} \frac{\tau_0}{1 + \tau_0} (\psi(\bar{\alpha}_s, \bar{\alpha}_r)' \Sigma_0 \lambda_0)^2 - \theta \ln(1 + \tau_0) - \nu (\lambda_0' \Sigma_0 \lambda_0 - 1)$$

where ν is the Lagrange multiplier on $\lambda_0' \Sigma_0 \lambda_0 = 1$. The first-order condition with respect to λ_0 is

$$\frac{\tau_0^*}{1 + \tau_0^*} (\psi(\bar{\alpha}_s, \bar{\alpha}_r)' \Sigma_0 \lambda_0^*) \psi(\bar{\alpha}_s, \bar{\alpha}_r) = \nu^* \lambda_0^*,$$

where τ_0^* is the optimal signal precision. Consequently, the agent's optimal weights satisfy $\lambda_0^* \propto \psi(\bar{\alpha}_s, \bar{\alpha}_r)$. That is, with mutual information, the agent optimally chooses a λ_0^* that is proportional to the optimal action weights $\psi(\bar{\alpha}_s, \bar{\alpha}_r)$.

Suppose the aggregate action function is s -measurable: $\bar{\alpha}_r = 0$. In this case, $\psi(\bar{\alpha}_s, 0)' = (1 - \beta + \beta \bar{\alpha}_s, 0)$. It follows that the agent's best response is s -measurable: $(\lambda_0^*)' = (\sigma_s^{-1}, 0)$. By the usual fixed-point arguments, there exists a linear equilibrium featuring zero non-fundamental volatility: $\bar{\alpha}_r^* = 0$. This result is driven by the R-monotonicity of the KL divergence and is the LQG analog of Proposition 7.

In fact, all linear equilibria feature zero non-fundamental volatility, even in the presence of strategic complementarity ($\beta \in (0, 1)$). In a hypothetical equilibrium in which other agents pay some attention to r ($\bar{\alpha}_r \neq 0$), each individual agent would also have an incentive to pay some attention to r , due to the strategic complementarity, but would optimally pay less attention to r than the other agents ($\beta < 1$). Consequently, there is no equilibrium in which the agents attend

⁴⁷Note that, to simplify our analysis, we have assumed in this example that the planner must choose for each agent the action that maximizes that agent's utility given the agent's posterior beliefs. In our general framework, the planner can choose both what information the agents acquire and how they act based on that information. Because the externality in this example operates through $\bar{\alpha}_s$, a planner who could control information choices and actions would prefer both that agents acquire additional information and that they respond with more extreme actions to given their posterior beliefs.

to r . This is true even when there are multiple equilibria (which will occur when $\beta > \frac{1}{2}$, as we show in the next sub-section).

Consider now the agent's problem with cost function C_{FI-SR} :

$$\max_{\tau_0 \geq 0, \lambda_0 \in \mathbb{R}^2} \min_{\nu \in \mathbb{R}} \frac{\tau_0}{1 + \tau_0} (\psi(\bar{\alpha}_s, \bar{\alpha}_r)' \Sigma_0 \lambda_0)^2 - \theta \tau_0 \lambda_0' \lambda_0 - \nu (\lambda_0' \Sigma_0 \lambda_0 - 1).$$

The first order condition with respect to λ_0 can be written as

$$\frac{\tau_0^*}{1 + \tau_0^*} (\psi(\bar{\alpha}_s, \bar{\alpha}_r)' \Sigma_0 \lambda_0^*) \psi(\bar{\alpha}_s, \bar{\alpha}_r)' = \theta \tau_0^* \Sigma_0^{-1} \lambda_0^* + \nu^* \lambda_0.$$

If $\bar{\alpha}_r = 0$ and $\tau_0^* > 0$, then the optimal strategy is not s -measurable: $(\lambda_0^*)' \neq (\sigma_s^{-1}, 0)$.⁴⁸ As long as an agent gathers some information ($\tau_0^* > 0$), the agent will place some weight on r .

Consequently, with the Fisher information cost function C_{FI-SR} , in any equilibrium with non-zero information acquisition, $\tau_0^* > 0$, non-fundamental volatility arises: $\bar{\alpha}_r \neq 0$. We show in the technical appendix Section D that a linear equilibrium exists, and moreover if θ is sufficiently small, all linear equilibria feature non-fundamental volatility. This is the LQG analog of Proposition 6, and is driven by the nowhere-R-monotonicity of C_{FI-SR} .⁴⁹

7.3 Multiplicity

With endogenous information acquisition, the assumption that $\beta < 1$ is not sufficient to ensure uniqueness. In this section we show how multiplicity can arise when $\beta \in (0, 1)$, even when the cost function is invariant in \bar{A} . The example we construct is closely related to certain results in Myatt and Wallace [2012]. With mutual information (which is invariant in \bar{A}), the agent solves

$$\max_{\tau_0 \geq 0} (1 - \beta + \beta \bar{\alpha}_s)^2 \sigma_s^2 \left(1 - \frac{1}{1 + \tau_0 \sigma_s^2} \right) + \theta \ln \left(\frac{1}{1 + \tau_0 \sigma_s^2} \right).$$

Note that this problem is strictly concave in $(1 + \tau_0 \sigma_s^2)^{-1}$ and hence the first-order condition is sufficient. The agent's first-order condition with respect to τ_0 can be written as

$$(1 - \beta + \beta \bar{\alpha}_s)^2 \sigma_s^2 = \theta (1 + \tau_0^* \sigma_s^2). \quad (19)$$

An equilibrium $(\tau_0^*, \bar{\alpha}_s^*)$ jointly satisfies (19) and mean-consistency (17). To demonstrate that multiple equilibrium can exist, we construct two equilibria in this setting.

Set $\theta = (1 - \beta)^2 \sigma_s^2$. It is immediate that $\bar{\alpha}_s^* = 0, \tau_0^* = 0$ satisfies (19) and (17); in this equilibrium no information is acquired.

⁴⁸To see this, observe that the matrix Σ_0 is not diagonal (because r is correlated with s). Hébert and Woodford [2021] show that agents choose a λ_0 that maximally covaries with $\psi(\bar{\alpha}_s, \bar{\alpha}_r)$ under the resulting posterior; when r is a public signal about s , this will always involve putting some weight on r .

⁴⁹We don't prove this formally for the Gaussian setting; it follows from the same arguments used in Section 3.

A second solution can exist, provided β is sufficiently large. For example, let $\beta = 2/3$. In this case, $\bar{\alpha}_s^* = \frac{1}{2}$, $\tau_0^* = 3(\sigma_s^2)^{-1}$ is the other solution to these equations; in this equilibrium agents gather a non-zero amount of information. More generally, whenever $\beta \in (1/2, 1)$, there are two solutions to (17) and (19), one in which $\bar{\alpha}_s^* = 0$ and another in which $\bar{\alpha}_s^* > 0$.⁵⁰

Consider the equilibrium in which all agents gather zero information. If the marginal cost of information at $\tau_0 = 0$ is sufficiently high, then it is individually-optimal for an agent to gather zero information.⁵¹ However, if other agents gather information and act upon it, this increases the individual agent's incentive to gather information. If this form of strategic complementarity is sufficiently strong, an equilibrium in which agents gather information can also exist.

When the two equilibria co-exist, the equilibrium with information gathering Pareto-dominates the equilibrium with zero information gathering; see the technical appendix, Section D, for details. The linear-quadratic utility function is mean-critical and mutual information is invariant in \bar{A} . As a result, the solution to the planner's problem is an equilibrium (as in Proposition 5); in this case, the equilibrium with information gathering. When $\beta \in (0, 1)$, it is possible that another Pareto-dominated equilibrium exists; weak strategic substitutability ($\beta \leq 0$) rules out this possibility.

7.4 The Allocation of Attention

We conclude our discussion of the LQG framework by considering the agent's allocation of attention. We have cautioned against the literal interpretation of ω^i as the agent's signal and λ_0 as the agent's allocation of attention. This is because agents can reach the same posterior beliefs in multiple ways: by paying attention to s or r directly, or by paying attention to \bar{a} . In Appendix Section B, we discuss this in more detail in the context of our general model. Here, we provide a simple illustration in the LQG context.

We define a version of the Fisher information cost function by

$$D_{FI}(\mu' || \mu) = E^{\mu'} \left[\left(\frac{\partial}{\partial s} \ln \left(\frac{d\mu'}{d\mu} \right) \right)^2 + \left(\frac{\partial}{\partial \bar{a}} \ln \left(\frac{d\mu'}{d\mu} \right) \right)^2 \right].$$

Let us require that the agent receive a signal of the form $\bar{\omega}^i = (s, \bar{a}) \cdot \lambda + \bar{\varepsilon}^i$, with $\bar{\varepsilon}^i \sim N(0, \tau^{-1})$, for some $\lambda \in \mathbb{R}^2$ and $\tau \geq 0$; we will interpret λ as describing the allocation of the agent's attention

⁵⁰The first order condition and mean-consistency condition combine into a quadratic equation in $\bar{\alpha}_s$ with no constant, and hence there are at most two solutions, one of which is always $\bar{\alpha}_s^* = 0$. When $\beta \in (1/2, 1)$, $(\bar{\alpha}_s^* = \beta^{-1}(2\beta - 1), \tau_0^* \sigma_s^2 = \beta^2(1 - \beta)^{-2} - 1)$ is also a solution. See the technical appendix, Section D, for details. A second equilibrium exists if $\beta > \frac{1}{2}$, but in general the relevant threshold will depend on the cost function.

⁵¹Our results might seem to contradict certain results in Colombo, Femminis, and Pavan [2014]. However, Colombo, Femminis, and Pavan [2014] impose the assumption that the marginal cost of increasing precision at zero is zero. This assumption is not satisfied by our example cost functions (which are based on cost functions commonly used in the rational inattention literature), and as a result an equilibrium in which no information is acquired exists.

across s and \bar{a} . The resulting cost function can be written as

$$C_{FI}(\lambda, \tau) = \theta \tau \lambda' \lambda.$$

Note that C_{FI} depends only on the nature and precision of the agent's signal and not on the agent's prior uncertainty about s or \bar{a} ; it can thus be viewed as a prior-invariant information cost defined on Gaussian signals.⁵² We show below that the reduced form cost will depend on \bar{a} , despite the prior-invariance of C_{FI} .

In a linear equilibrium ($\bar{a}(s, r) = \bar{\alpha}_s s$), the signal $\bar{\omega}$ is equivalent, from an informational perspective, to receiving the signal $\omega = ((1, \bar{\alpha}_s) \cdot \lambda)s + \bar{\epsilon}^i$. Consequently, any two vectors λ_1, λ_2 such that $(1, \bar{\alpha}_s) \cdot \lambda_1 = (1, \bar{\alpha}_s) \cdot \lambda_2$ will generate the same posteriors. The agent will therefore prefer whichever signal is less costly.

Adopting the normalization $(1, \bar{\alpha}_s) \cdot \lambda = 1$, we can define a reduced-form cost function by

$$C(\tau_0, \bar{\alpha}) = \min_{\lambda \in \mathbb{R}^2} \max_{\nu \in \mathbb{R}} C_{FI}(\lambda, \tau_0) + \nu(1 - (1, \bar{\alpha}_s) \cdot \lambda).$$

This cost function reflects the cost of optimally acquiring a signal about s via whichever means (s , \bar{a} , or a combination thereof) is least costly. The associated first-order condition is $2\theta\tau_0(\lambda^*)' = \nu^*(1, \bar{\alpha}_s)$, and therefore, to satisfy the constraint, $\nu^*(1 + \bar{\alpha}_s^2) = 2\theta\tau_0$. It follows that

$$C(\tau_0, \bar{\alpha}) = \theta \frac{\tau_0}{1 + \bar{\alpha}_s^2} = C_{FI-S\bar{A}}((1, 0)', \tau_0, \Sigma_0, \bar{\alpha}_s, 0).$$

Thus the cost function $C_{FI-S\bar{A}}$ can be understood as the reduced-form representation of C_{FI} . It depends on \bar{a} , even though the “true” cost function, C_{FI} , does not. This dependence is due to the fact that under the “true” cost function, the least costly method of reaching a given set of posteriors on s is to pay some attention to \bar{a} , and the degree to which it is optimal to attend to \bar{a} depends on the volatility of \bar{a} relative to s . This example serves as an illustration of why (λ_0, τ_0) (and in our general framework, π_0 and π) should be understood as describing the posteriors of the agent but not the agent's allocation of attention.

8 Conclusion

In this paper we study the relationship between information cost functions and properties of equilibria in large games with strategic interaction and rationally inattentive agents. We demonstrate that there is a close connection between certain properties of information cost functions and whether or not an equilibrium is inefficient and/or exhibits non-fundamental volatility.

⁵²That is, given a Gaussian prior, D_{FI} (which is not prior-invariant) induces a prior-invariant cost on Gaussian signals. As a result, a prior-invariant version of D_{FI} , constructed from a fixed Gaussian prior along the lines of [Denti et al. \[Forthcoming\]](#), would induce the same cost on Gaussian signals.

Even in the absence of payoff externalities, equilibria will be efficient if and only if a local form of invariance in the aggregate action holds. As a result, invariance in the aggregate action ensures the existence of a constrained efficient equilibrium, and nowhere-invariance leads to inefficiency. Likewise, monotonicity with respect to noisy public signals leads to the existence of an equilibrium with zero non-fundamental volatility, and nowhere-monotonicity ensures that all equilibria with information acquisition exhibit non-fundamental volatility.

Our results show that the structure of the information cost function is an important determinant of equilibria properties in this class of games. The standard rational inattention cost function, mutual information, ensures the existence of an equilibrium that is both efficient and exhibits zero non-fundamental volatility. In contrast, alternative cost functions such as the neighborhood-based cost functions proposed by Hébert and Woodford [2021] can ensure non-fundamental volatility and/or inefficiency.

References

- Hassan Afrouzi. Strategic inattention, inflation dynamics, and the non-neutrality of money. *Columbia mimeo*, 2020.
- Charalambos D Aliprantis and Kim C Border. *Infinite Dimensional Analysis: A Hitchhiker's Guide*. Springer Science & Business Media, 2006.
- Manuel Amador and Pierre-Olivier Weill. Learning from prices: Public communication and welfare. *Journal of Political Economy*, 118(5):866–907, 2010.
- Shun-ichi Amari and Hiroshi Nagaoka. *Methods of information geometry*, volume 191. American Mathematical Soc., 2007.
- George-Marios Angeletos and Jennifer La'O. Noisy business cycles. In *NBER Macroeconomics Annual 2009, Volume 24*, pages 319–378. University of Chicago Press, 2010.
- George-Marios Angeletos and Jennifer La'O. Sentiments. *Econometrica*, 81(2):739–779, 2013.
- George-Marios Angeletos and Jennifer La'O. Optimal monetary policy with informational frictions. *Journal of Political Economy*, 128(3):1027–1064, 2020.
- George-Marios Angeletos and Chen Lian. Incomplete information in macroeconomics: Accommodating frictions in coordination. *Handbook of Macroeconomics*, 2:1065–1240, 2016.
- George-Marios Angeletos and Alessandro Pavan. Efficient use of information and social value of information. *Econometrica*, 75(4):1103–1142, 2007.
- George-Marios Angeletos and Alessandro Pavan. Policy with dispersed information. *Journal of the European Economic Association*, 7(1):11–60, 2009.
- George-Marios Angeletos and Karthik Sastry. Inattentive economies: General equilibrium and welfare theorems. *MIT mimeo*, 2021.
- George-Marios Angeletos, Luigi Iovino, and Jennifer La'O. Learning over the business cycle: Policy implications. *Journal of Economic Theory*, 190, 2020.
- Arindam Banerjee, Xin Guo, and Hui Wang. On the optimality of conditional expectation as a bregman predictor. *IEEE Transactions on Information Theory*, 51(7):2664–2669, 2005.
- Dirk Bergemann and Stephen Morris. Robust predictions in games with incomplete information. *Econometrica*, 81(4):1251–1308, 2013.

- Alexander W Bloedel and Weijie Zhong. The cost of optimally-acquired information. Technical report, Technical report, Working paper, Stanford University, 2020.
- Andrew Caplin and Mark Dean. Revealed preference, rational inattention, and costly information acquisition. *American Economic Review*, 105(7):2183–2203, 2015.
- Andrew Caplin, Mark Dean, and John Leahy. Rationally inattentive behavior: Characterizing and generalizing shannon entropy. *Journal of Political Economy*, 130(6):1676–1715, 2022.
- Nikolai Nikolaevich Chentsov. *Statistical decision rules and optimal inference*. Number 53. American Mathematical Soc., 1982.
- Luca Colombo, Gianluca Femminis, and Alessandro Pavan. Information acquisition and welfare. *The Review of Economic Studies*, 81(4):1438–1483, 2014.
- Thomas M Cover and Joy A Thomas. *Elements of information theory*. John Wiley & Sons, 2012.
- Mark Dean and Nathaniel Neligh. Experimental tests of rational inattention. *Columbia mimeo*, 2019.
- Tommaso Denti. Unrestricted information acquisition. *Theoretical Economics*, Forthcoming.
- Tommaso Denti, Massimo Marinacci, Aldo Rustichini, et al. Experimental cost of information. Technical report, Forthcoming.
- Sanford Grossman. On the efficiency of competitive stock markets where trades have diverse information. *The Journal of Finance*, 31(2):573–585, 1976.
- Sanford J. Grossman and Joseph E. Stiglitz. Information and competitive price systems. *The American Economic Review Papers and Proceedings*, 66(2):246–253, 1976.
- Sanford J Grossman and Joseph E Stiglitz. On the impossibility of informationally efficient markets. *The American economic review*, 70(3):393–408, 1980.
- F. A. Hayek. The use of knowledge in society. *The American Economic Review*, 35(4):519–530, 1945.
- Benjamin Hébert. Moral hazard and the optimality of debt. *The Review of Economic Studies*, 85(4):2214–2252, 2018.
- Benjamin Hébert and Michael Woodford. Neighborhood-based information costs. *American Economic Review*, 111(10):3225–55, 2021.
- Christian Hellwig and Laura Veldkamp. Knowing what others know: Coordination motives in information acquisition. *Review of Economic Studies*, 76(1):223–251, 2009.
- Junehyuk Jung, Jeong Ho Kim, Filip Matějka, and Christopher A Sims. Discrete actions in information-constrained decision problems. *The Review of Economic Studies*, 86(6):2643–2667, 2019.
- Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. *American Economic Review*, 101(6):2590–2615, 2011.
- Jean-Jacques M Laffont. On the welfare analysis of rational expectations equilibria with asymmetric information. *Econometrica*, 53(1):1–29, January 1985.
- Bartosz Mackowiak and Mirko Wiederholt. Optimal sticky prices under rational inattention. *American Economic Review*, 99(3):769–803, June 2009.
- Filip Matějka, Alisdair McKay, et al. Rational inattention to discrete choices: A new foundation for the multinomial logit model. *American Economic Review*, 105(1):272–98, 2015.
- Paul R Milgrom and Robert J Weber. Distributional strategies for games with incomplete information. *Mathematics of operations research*, 10(4):619–632, 1985.
- Stephen Morris and Hyun Song Shin. Social value of public information. *American Economic Review*, 92(5):1521–1534, 2002.

- Stephen Morris and Ming Yang. Coordination and continuous stochastic choice. *Review of Economic Studies*, Forthcoming.
- David P Myatt and Chris Wallace. Endogenous information acquisition in coordination games. *The Review of Economic Studies*, 79(1):340–374, 2012.
- Luigi Paciello and Mirko Wiederholt. Exogenous information, endogenous information, and optimal monetary policy. *Review of Economic Studies*, 81(1):356–388, 2014.
- Alessandro Pavan. Attention, coordination, and bounded recall. *Northwestern University mimeo*, 2016.
- Luciano Pomatto, Philipp Strack, and Omer Tamuz. The cost of information. *arXiv preprint arXiv:1812.04211*, 2020.
- R. Radner. Team decision problems. *The Annals of Mathematical Statistics*, 33(3):857–881, 1962.
- Doron Ravid. Ultimatum bargaining with rational inattention. *American Economic Review*, 110(9):2948–63, 2020.
- Christopher A. Sims. Implications of rational inattention. *Journal of Monetary Economics*, 50(3):665 – 690, 2003. ISSN 0304-3932. Swiss National Bank/Study Center Gerzensee Conference on Monetary Policy under Incomplete Information.
- Christopher A Sims. Rational inattention and monetary economics. *Handbook of Monetary Economics*, 3:155–181, 2010.
- Harald Uhlig. A law of large numbers for large economies. *Economic Theory*, 8(1):41–50, 1996.
- Stijn Van Nieuwerburgh and Laura Veldkamp. Information acquisition and under-diversification. *The Review of Economic Studies*, 77(2):779–805, 2010.
- Xavier Vives. Aggregation of information in large cournot markets. *Econometrica*, 56(4):851–76, 1988.
- Xavier Vives. Endogenous public information and welfare in market games. *Review of Economic Studies*, 84(2):935–963, 2017.
- David Walker-Jones. Rational inattention and perceptual distance. *arXiv preprint arXiv:1909.00888*, 2019.
- Michael Woodford. Imperfect common knowledge and the effects of monetary policy. *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*, 2003.
- Ming Yang. Coordination with flexible information acquisition. *Journal of Economic Theory*, 158:721–738, 2015.

Appendix

A Efficiency and Mean-Critical Preferences

In this appendix section we provide additional results highlighting the role of mean-critical preferences and their connection with efficiency. We first establish that, subject to some regularity conditions, preferences with the functional form given in (1) are in fact the only preferences that satisfy the mean-critical property; we also provide a multi-dimensional generalization. We then provide formal results on the connection between efficiency for all invariant divergences and

the mean-critical preferences.

A.1 Mean-Critical Preferences

For the purposes of this subsection, to allow our results to be generalized beyond our particular application, we will assume that A and S are compact subsets of \mathbb{R}^N , and that A is convex (and hence $\bar{A} = A$). We begin by defining the class of continuous utility functions.

Definition 13. A utility function $v : A \times \bar{A} \times S \rightarrow \mathbb{R}$ is **continuous** if it is continuous on $A \times \bar{A}$ for all $s \in S$.

Some of our proofs in the main text (as indicated) apply to this class. Within continuous utility functions, we consider the class of regular utility functions.

Definition 14. A utility function $v : A \times \bar{A} \times S \rightarrow \mathbb{R}$ is **regular** if it is concave on A for all $(\bar{a}, s) \in \bar{A} \times S$, continuously twice-differentiable on $A \times \bar{A}$ for all $s \in S$, and either (i) concave on \bar{A} for all $(a, s) \in A \times S$ or (ii) convex on \bar{A} for all $(a, s) \in A \times S$.

The linear-quadratic function in (2) is regular in this sense. Regularity allows for either concavity or convexity with respect to $\bar{a} \in \bar{A}$, but restricts utility functions to exhibiting either concavity or convexity everywhere (as opposed to exhibiting concavity on some parts of the domain and convexity on other parts).

We next formally define what it means for a utility function to be mean-critical.

Definition 15. A utility function $v : A \times \bar{A} \times S \rightarrow \mathbb{R}$ is **mean-critical given** $\sigma : S \rightarrow \Delta(A)$ if it is differentiable on $A \times \bar{A}$ for all $s \in S$ and if $E^{\sigma(s)}[a]$ is a critical point of $h(\bar{a}) = E^{\sigma(s)}[v(a, \bar{a}, s)]$. A utility function is **mean-critical** if it is mean-critical for all such σ .

Armed with these definitions, we characterize the class of regular mean-critical utility functions.

Lemma 3. *A regular utility function v is mean-critical if and only if there exists a function $G : \bar{A} \times S \rightarrow \mathbb{R}$ and a function $g : A \times S \rightarrow \mathbb{R}$ such that*

$$v(a, \bar{a}, s) = g(a; s) + G(\bar{a}; s) + (a - \bar{a}) \cdot \nabla G(\bar{a}; s), \quad (20)$$

where $\nabla G(\bar{a}; s)$ denotes the gradient of G with respect to its first argument.

Proof. See the appendix, F1. □

The “if” part of the proof is straightforward with a one-dimensional action space (as in the main text). The first-order partial derivative of the payoff function in (20) with respect to \bar{a} is

given by⁵³

$$\frac{\partial v(a, \bar{a}, s)}{\partial \bar{a}} = \frac{\partial}{\partial \bar{a}} G(\bar{a}, s) - \frac{\partial}{\partial \bar{a}} G(\bar{a}, s) + (a - \bar{a}) \frac{\partial^2 G(\bar{a}, s)}{\partial \bar{a}^2} = (a - \bar{a}) \frac{\partial^2 G(\bar{a}, s)}{\partial \bar{a}^2}.$$

Taking expectations demonstrates this utility function is mean-critical. The “only if” part builds on the proofs in [Banerjee et al. \[2005\]](#).

In the context of an abstract game, there is no obvious reason why the utility function should have this particular functional form. However, the functional form in (20) has an interpretation in terms of production economies with full risk-sharing. Consider the following example, which is similar to the “coconuts and money” example of [Angeletos and Sastry \[2021\]](#). Let good one be the numeraire and let $a \in A$ be agent’s consumption of good two. Define $-G(\bar{a}, s)$ as the cost of producing \bar{a} units of good two, and assume that the price of good two is equal to the marginal cost, $p = -\frac{\partial}{\partial \bar{a}} G(\bar{a}, s)$. Each agent receives utility $g(a, s)$ from her consumption of good two and has linear utility in good one. The agent’s consumption of good one is equal to her share of the profits from production, $p \cdot \bar{a} + G(\bar{a}, s)$, less her spending on good two, $p \cdot a$. As a result, the agent’s utility is given by (20).

This example illustrates the connection between an absence of payoff externalities at the margin in games (characterized by the mean-critical property) and the classic welfare theorem results, that is, the absence of pecuniary externalities at the market equilibrium under full risk-sharing.⁵⁴

A.2 Mean-Critical Preferences and Efficiency

We next show that a local version of the mean-critical property is necessary and sufficient to guarantee efficiency for all divergences that are invariant in \bar{A} . To simplify the statement of this result, we return to our assumptions of a one-dimensional action space and a finite set S .

Proposition 8. *Assume the divergence D is invariant in \bar{A} and that the utility function is continuous. Let $(\pi_0, \bar{\alpha})$ be a constrained-efficient strategy profile with $\text{supp}(\pi_0) \subset A \times \text{int}(\mathcal{U}_0)$ and such that $\bar{\alpha}$ lies in the relative interior of \bar{A} , and let σ be the conditional distribution of a given s under π_0 . Then $(\pi_0, \bar{\alpha})$ is an equilibrium if and only if v is mean-critical given σ .*

Proof. See the appendix, [F.2](#). □

Intuitively, absent informational externalities, an equilibrium will be efficient if and only if there are no payoff externalities. When there are both payoff and information externalities, we in general expect that inefficiency will arise (as in [Proposition 4](#)), subject to the caveat that it is possible that the payoff and informational externalities can exactly offset.

⁵³Note by regularity that G is twice-differentiable.

⁵⁴In this example, full risk-sharing (agents equating their marginal values of wealth across all realizations of aggregate and idiosyncratic shocks) is achieved by assuming quasi-linear preferences. In the main part of their analysis, [Angeletos and Sastry \[2021\]](#) instead impose a complete markets assumption to achieve full risk-sharing.

B The Allocation of Attention

In this appendix, we describe the problem of an agent who can acquire signals subject to a posterior-separable cost. We show that the problem of the agent described in the main text can be understood as a reduced form for this more general problem. We discuss the key example of $D_{S\bar{A}}$ at the end of this section, and argue that it can be motivated as a discrete approximation of the reduced form of a problem in which the agents allocate their attention between the fundamentals $s \in S$ and the actions of other agents, $\bar{a} \in \bar{A}$.

In what follows, we assume S and R are compact but not necessarily finite subsets of \mathbb{R}^n , while continuing to assume μ_0 has full support on $S \times R$. We endow \mathcal{U}_0 with the weak* topology and \bar{A} with the sup-norm topology. Note that the definitions of the operations γ and η used in the main text continue to be valid in the continuous state setting.

Signals Let Ω denote a space of possible signal realizations, which we assume is rich (we state a sufficient assumption below). An admissible signal structure consists of a signal alphabet $\bar{\nu} \in \Delta(\Omega)$ and a function $\nu : \Omega \times S \times R \times \bar{A} \rightarrow \mathbb{R}_+$ satisfying, for each $(s, r, \bar{a}) \in S \times R \times \bar{A}$,

$$\int_{\Omega} \nu(\omega|s, r, \bar{a}) d\bar{\nu}(\omega) = 1.$$

and such that, for each $\omega \in \Omega$, $\nu(\omega|\cdot) \in L^1(\mu)$, meaning that

$$\int_{S \times R \times \bar{A}} \nu(\omega|s, r, \bar{a}) d\mu(s, r, \bar{a}) < \infty.$$

Note that the signal can condition on \bar{a} in addition to s and r ; we interpret this as describing a situation in which the agent pays attention to the actions of other agents. That is, unlike the distributions over posteriors $\pi \in \Pi(\mu)$ and $\pi_0 \in \Pi_0(\mu_0)$ considered in the main text, we will think of the signal structure $(\bar{\nu}, \nu) \in \mathcal{V}$ as describing the allocation of the agent's attention.

Let $\mu \in \mathcal{U}$ be the agent's prior on $S \times R \times \bar{A}$. From this prior, each signal structure induces a density of signal realizations with respect to $\bar{\nu}$, $\bar{f}[\nu, \mu] : \Omega \rightarrow \mathbb{R}_+$, defined by

$$\bar{f}[\nu, \mu](\omega) = \int_{S \times R \times \bar{A}} \nu(\omega|s, r, \bar{a}) d\mu(s, r, \bar{a}),$$

and a set of posterior probability densities $f^\omega[\nu, \mu] \in L^1(\mu)$ defined by, for all ω such that $\bar{f}[\nu, \mu](\omega) > 0$,

$$f^\omega[\nu, \mu](s, r, \bar{a}) = \frac{\nu(\omega|s, r, \bar{a})}{\bar{f}[\nu, \mu](\omega)}.$$

Note that these functions are the posterior densities with respect to the measure $\mu \in \mathcal{U}$. We adopt the convention that $f^\omega[\nu, \mu](s, r, \bar{a}) = 1$ for all ω such that $\bar{f}[\nu, \mu](\omega) = 0$. Define the set of

admissible densities with respect to the prior μ as

$$\mathcal{F}(\mu) = \{f \in L^1(\mu) : 1 = \int_{S \times R \times \bar{A}} f(s, r, \bar{a}) d\mu(s, r, \bar{a})\}.$$

Observe by construction that $f^\omega[\nu, \mu] \in \mathcal{F}(\mu)$ for all $\mu \in \mathcal{U}$, $\omega \in \Omega$, and $(\bar{\nu}, \nu) \in \mathcal{V}$. Note also that each density $f \in \mathcal{F}(\mu)$ induces a measure $\hat{\mu}[f, \mu] \in \mathcal{U}$ by $d\hat{\mu}[f, \mu](s, r, \bar{a}) = f(s, r, \bar{a})d\mu(s, r, \bar{a})$.

We next define the class of cost functions on signals.

Definition 16. A **cost function defined on signals** is a cost function $\bar{C} : \mathcal{V} \times \mathcal{U} \rightarrow [0, \infty]$.

Given such a cost function, we can define the agent's problem:

$$\sup_{(\bar{\nu}, \nu) \in \mathcal{V}, \sigma : \Omega \rightarrow \Delta(A)} E^\mu \left[\int_{\text{supp}(\bar{\nu})} E^{\sigma(\omega)}[v(a, \bar{a}, s)] \nu(\omega | s, r, \bar{a}) d\bar{\nu}(\omega) \right] - \bar{C}(\bar{\nu}, \nu, \mu),$$

where $\sigma : \Omega \rightarrow \Delta(A)$ describes the agent's mixed strategy over actions conditional on receiving the signal realization $\omega \in \Omega$.

Posteriors Take as given a prior $\mu \in \mathcal{U}$ and signal structure $(\bar{\nu}, \nu) \in \mathcal{V}$. Together, these induce a measure on signal realizations $\bar{f}[\nu, \mu](\bar{\nu}) \in \Delta(\Omega)$. This measure, combined with the action strategy $\sigma : \Omega \rightarrow \Delta(A)$, defines via the product measure a measure on $A \times \Omega$. From this product measure, the mapping $(a, \omega) \mapsto (a, \hat{\mu}[f^\omega[\nu, \mu], \mu])$ induces a strategy $\hat{\pi}[\bar{\nu}, \nu, \mu] \subset \Delta(A \times \mathcal{U})$. By construction, this strategy satisfies Bayes-consistency: for all $(s, r, \bar{a}) \in S \times R \times \bar{A}$,

$$\begin{aligned} \int_{\text{supp}(\hat{\pi})} d\mu'(s, r, \bar{a}) d\hat{\pi}[\bar{\nu}, \nu, \mu](a, \mu') &= \int_{\text{supp}(\bar{\nu})} f^\omega[\nu, \mu](s, r, \bar{a}) \bar{f}[\nu, \mu](\omega) d\bar{\nu}(\omega) d\mu(s, r, \bar{a}) \\ &= \int_{\text{supp}(\bar{\nu})} \nu(\omega | s, r, \bar{a}) d\bar{\nu}(\omega) d\mu(s, r, \bar{a}) = d\mu(s, r, \bar{a}), \end{aligned}$$

and by construction $\hat{\mu}[f^\omega[\nu, \mu], \mu] \ll \mu$. Let $\Pi(\mu)$ be the set of Bayes-consistent strategies with full support on posteriors absolutely continuous with respect to the prior μ .

Moreover, we can write the agent's expected payoff (not including information costs) under $(\bar{\nu}, \nu)$ as

$$E^\mu \left[\int_{\text{supp}(\bar{\nu})} E^{\sigma(\omega)}[v(a, \bar{a}, s)] \nu(\omega | s, r, \bar{a}) d\bar{\nu}(\omega) \right] = E^{\hat{\pi}[\bar{\nu}, \nu, \mu]}[V(a, \mu')]$$

where $V(a, \mu') = E^{\mu'}[v(a, \bar{a}, s)]$ as in the main text. Note here that if $\hat{\pi}[\bar{\nu}, \nu, \mu] = \pi$ in the weak* topology, then $E^{\hat{\pi}[\bar{\nu}, \nu, \mu]}[V(a, \mu')] = E^\pi[V(a, \mu')]$. That is, the expected payoff is the same for all signal structures $(\bar{\nu}, \nu) \in \mathcal{V}$ that result in the same reduced-form strategy $\pi \in \Pi(\mu)$.

It follows that the agent can divide her problem into two parts: choosing some feasible strategy $\pi \in \Pi(\mu)$, and choosing the signal structure $(\bar{\nu}, \nu) \in \mathcal{V}$ to induce that strategy while minimizing information costs. The following lemma shows that all $\pi \in \Pi(\mu)$ are feasible under the assumption that Ω is rich (we assume $\Omega = A \times \mathcal{U}_0$, which is sufficient).

Lemma 4. Assume $\Omega = A \times \mathcal{U}_0$. For any prior $\mu \in \mathcal{U}$ and any $\pi \in \Pi(\mu)$, there exists a signal structure $(\bar{\nu}, \nu) \in \mathcal{V}$ such that $\pi = \hat{\pi}[\bar{\nu}, \nu, \mu]$ and $\nu(\omega|s, r, \bar{a}) = \nu(\omega|s, r, \bar{a}')$ for all $\omega \in \Omega$, $s \in S$, $r \in R$, and $\bar{a}, \bar{a}' \in \bar{A}$.

Proof. See the technical appendix, H.10. □

As a consequence, the agent's problem can be written as

$$\max_{\pi \in \Pi(\mu)} E^\pi[V(a, \mu')] - C(\pi, \mu),$$

where

$$C(\pi, \mu) = \inf_{(\bar{\nu}, \nu) \in \mathcal{V}: \hat{\pi}[\bar{\nu}, \nu, \mu] = \pi} \bar{C}(\bar{\nu}, \nu, \mu).$$

This is the problem considered in the main text (if $C(\pi, \mu)$ is posterior-separable). The reduced-form cost function $C(\pi, \mu)$ is the result of a cost-minimization problem over signals.

Given the aggregate action strategy $\bar{\alpha} \in \bar{\mathcal{A}}$, there is a one-to-one relationship between $\pi \in \Pi(\mu)$ and $\pi_0 \in \Pi_0(\mu)$. Define $\Phi(\pi_0, \bar{\alpha})$ as the strategy induced from π_0 by $(a, \mu'_0) \mapsto (a, \phi[\mu'_0, \bar{\alpha}])$, and observe that $\pi = \Phi(\pi_0, \bar{\alpha})$. Similarly, the strategy π_0 is induced from π by $(a, \mu') \mapsto (a, \gamma_{\bar{A}}\{\mu'\})$. We can therefore write the agent's problem and reduced form cost function in the alternative formulation,

$$\max_{\pi_0 \in \Pi_0(\mu_0)} E^{\pi_0}[V(a, \phi[\mu'_0, \bar{\alpha}])] - C_0(\pi_0, \mu_0, \bar{\alpha}),$$

where $C_0(\pi_0, \mu_0, \bar{\alpha}) = C(\Phi(\pi_0, \bar{\alpha}), \phi[\mu_0, \bar{\alpha}])$.

Posterior Separability In the main text, we focus on the case in which $C(\pi, \mu)$ is posterior-separable. In the examples below, we construct reduced form posterior-separable cost functions $C(\pi, \mu)$ from posterior-separable cost functions $\bar{C}(\bar{\nu}, \nu, \mu)$ of the form

$$C(\bar{\nu}, \nu, \mu) = \int_{\text{supp}(\bar{\nu})} \bar{D}(f^\omega[\nu, \mu] || \mu) \bar{f}[\nu, \bar{\mu}](\omega) d\bar{\nu}(\omega), \quad (21)$$

where $\bar{D} : \{(f, \mu) \in \mathcal{F} \times \mathcal{U} : f \in \mathcal{F}(\mu)\} \rightarrow [0, \infty]$ is a sort of divergence defined on densities. We assume that if $f(s, r, \bar{a}) = 1$ everywhere, then $\bar{D}(f || \mu) = 0$.

We will define the kind of divergence studied in the main text by

$$D(\mu' || \mu) = \inf_{f \in \mathcal{F}(\mu): \hat{\mu}[f, \mu] = \mu'} \bar{D}(f || \mu).$$

Note that this definition automatically ensures convexity and that $D(\mu || \mu) = 0$.

Our strategy when constructing our example cost functions $C(\pi, \mu)$ is to minimize \bar{D} at each posterior in the support of some Bayes-consistent π , and then verify that the resulting densities can be generated by some signal structure. Such densities are always optimal if feasible. For the KL divergence example, we verify the assumption in this section. For our other examples, we

will verify this property in the context of our discussion of the continuous state case (technical appendix Section C).

In this exercise, it is convenient to consider densities in \mathcal{F}_0 , the set of continuous functions $S \times R \rightarrow \mathbb{R}_+$, instead of in \mathcal{F} . To this end, let us define the operation $\bar{\phi} : \mathcal{F}_0 \times \bar{\mathcal{A}} \rightarrow \mathcal{U}$ by

$$d\bar{\phi}[f_0, \bar{\alpha}](s, r, \bar{a}) = f_0(s, r) d\mu_0(s, r) d\delta_{\bar{\alpha}(s, r)}(\bar{a}), \quad (22)$$

where $d\delta_{\bar{\alpha}(s, r)}$ denotes the point mass on $\bar{\alpha}(s, r)$, which is the analog of (3) for densities. For $f_0 \in \mathcal{F}_0(\mu_0)$ (the set of densities with $\int_{S \times R} f_0(s, r) \mu_0(s, r) = 1$) and $\bar{\alpha} \in \bar{\mathcal{A}}$, we can define a reduced-form divergence on densities,

$$D^{\mathcal{F}}(f_0 || \mu_0; \bar{\alpha}) = D(\bar{\phi}[f_0, \bar{\alpha}] || \phi[\mu_0, \bar{\alpha}]).$$

Note that this definition requires that for any $f_0, f'_0 \in \mathcal{F}_0(\mu_0)$, if $\bar{\phi}[f_0, \bar{\alpha}] = \bar{\phi}[f'_0, \bar{\alpha}]$, then $D^{\mathcal{F}}(f_0 || \mu_0; \bar{\alpha}) = D^{\mathcal{F}}(f'_0 || \mu_0; \bar{\alpha})$. We will verify this property in each of our examples.

Examples We next provide examples of how reduced form divergences like the examples D_{KL} , $D_{S\bar{\mathcal{A}}}$, and D_{SR} considered in the main text can arise from cost functions defined on signal structures.

The KL Divergence Mutual information is a posterior-separable cost function defined on signal structures (i.e. it satisfies (21)). The corresponding divergence, the KL divergence, can be defined on densities as

$$\bar{D}_{KL}(f || \mu) = \int_{S \times R \times \bar{\mathcal{A}}} f(s, r, \bar{a}) \ln(f(s, r, \bar{a})) d\mu(s, r, \bar{a}).$$

By definition, $\mu = \phi[\mu_0, \bar{\alpha}]$ for some $\bar{\alpha} \in \bar{\mathcal{A}}$, and consequently

$$\bar{D}_{KL}(f || \mu) = \int_{S \times R} f(s, r, \bar{\alpha}(s, r)) \ln(f(s, r, \bar{\alpha}(s, r))) d\mu_0(s, r).$$

It follows immediately that all $f' \in \mathcal{F}(\mu)$ satisfying $\hat{\mu}[f', \mu] = \mu'$ for some $\mu' = \phi[\mu'_0, \bar{\alpha}]$ will generate the same value of $\bar{D}_{KL}(f || \mu)$ (since they are all identical μ -a.e.), and hence

$$D^{\mathcal{F}}_{KL}(f_0 || \mu_0; \bar{\alpha}) = \int_{S \times R} f_0(s, r) \ln(f_0(s, r)) d\mu_0(s, r)$$

and

$$D_{KL}(\phi[\mu'_0, \bar{\alpha}] || \phi[\mu_0, \bar{\alpha}]) = D^{\mathcal{F}}_{KL}(f_0 || \mu_0; \bar{\alpha})$$

for any f_0 such that $f_0 = \frac{d\mu'_0}{d\mu_0}$. In the finite state case, D_{KL} is the example cost function described in the text.

Because all signal structures $(\bar{\nu}, \nu)$ generate the same cost, by Lemma 4 there exists a signal

structure that induces any $\pi \in \Pi(\mu)$ and satisfies

$$\bar{C}(\bar{\nu}, \nu, \mu) = E^\pi[D_{KL}(\mu' || \mu)].$$

Fisher Information: D_{SR} Hébert and Woodford [2021] define a family of posterior-separable cost functions using Fisher information. When considering these cost functions, we will assume S and R are convex subsets of \mathbb{R}^N and that μ_0 has full support on $S \times R$. The general version of the divergence associated with the Fisher information cost can be defined for any $\mu \in \mathcal{U}$ and any strictly positive and weakly differentiable density f as

$$\bar{D}_{FI}(f || \mu) = \int_{S \times R \times \bar{A}} \|\nabla \ln [f(s, r, \bar{a})]\|^2 f(s, r, \bar{a}) d\mu(s, r, \bar{a}), \quad (23)$$

where ∇ denotes the gradient with respect to (s, r, \bar{a}) and $\|\cdot\|$ is the Euclidean norm. We define $\bar{D}_{FI}(f || \mu)$ by continuity for any f that is not strictly positive and weakly differentiable.

We can generate a restricted version of this cost function by assigning infinite cost to learning about the actions of others. Define

$$\bar{D}_{FI-SR}(f || \mu) = \begin{cases} \infty & \exists (s, r, \bar{a}, \bar{a}') \in S \times R \times \bar{A} \times \bar{A} \text{ s.t. } f(s, r, \bar{a}) \neq f(s, r, \bar{a}') \\ \bar{D}_{FI}(f || \mu) & \text{otherwise.} \end{cases}$$

This cost function rules out learning about the actions of others. It induces the reduced form divergence

$$D_{FI-SR}^{\mathcal{F}}(f_0 || \mu_0; \bar{\alpha}) = \int_{S \times R} \|\nabla \ln [f_0(s, r)]\|^2 f_0(s, r) d\mu_0(s, r)$$

for any weakly differentiable $f_0(s, r)$. Observe that, for any f_0, f'_0 that are equal μ_0 -a.e., if h_0 is a weak derivative of f_0 , it is also a weak derivative of f'_0 . Consequently, for any μ'_0 such that a weakly differentiable Radon-Nikodym derivative with respect to μ_0 exists,

$$D_{FI-SR}(\phi[\mu'_0, \bar{\alpha}] || \phi[\mu_0, \bar{\alpha}]) = D_{FI-SR}^{\mathcal{F}}(f_0 || \mu_0; \bar{\alpha})$$

for any $f_0 = \frac{d\mu'_0}{d\mu_0}$. We define the divergence by continuity if no weakly differentiable Radon-Nikodym derivative exists. We verify in the technical appendix, Section C that there is a signal structure that achieves

$$\bar{C}(\bar{\nu}, \nu, \mu) = E^\pi[D_{FI-SR}(\mu' || \mu)]$$

for all admissible π (see that section for a definition of admissible strategies).

Hébert and Woodford [2021] show that the cost function D_{SR} described in the text is a discretized version of this cost function, in the sense that there is a sequence of problems with a growing number of states such that D_{SR} converges to D_{FI-SR} as the number of states go to infinity.

Fisher Information: $D_{S\bar{A}}$ Lastly, let us next consider the case in which $R = \{r\}$ is a singleton, and define

$$\bar{D}_{FI-S\bar{A}}(f||\mu) = \int_{S \times R \times \bar{A}} \left[\left(\frac{\partial \ln(f(s, r, a))}{\partial s} \right)^2 + \left(\frac{\partial \ln(f(s, r, \bar{a}))}{\partial \bar{a}} \right)^2 \right] f(s, r, \bar{a}) d\mu(s, r, \bar{a}).$$

Consider the problem of choosing f to minimize costs and satisfy $f(s, r, \bar{\alpha}(s, r)) = f_0(s)$ for some weakly differentiable $f_0 : S \rightarrow (0, \infty)$. This problem can be thought of as choosing the allocation of attention. We assume in what follows that $\bar{\alpha}$ is differentiable with respect to s .⁵⁵

Because f and f_0 are both weakly differentiable, we must have, μ_0 -a.e.,

$$f_{0,s}(s) = f_s(s, r, \bar{\alpha}(s, r)) + f_{\bar{a}}(s, r, \bar{\alpha}(s, r))\bar{\alpha}_s(s, r).$$

Consequently,

$$\begin{aligned} \left[\left(\frac{\partial \ln(f(s, r, a))}{\partial s} \right)^2 + \left(\frac{\partial \ln(f(s, r, \bar{a}))}{\partial \bar{a}} \right)^2 \right] f(s, r, \bar{a})^2 &= f_{0,s}(s, r)^2 + f_a(s, r, \bar{\alpha}(s, r))^2(1 + \bar{\alpha}_s(s, r)^2) \\ &\quad - 2f_a(s, r, \bar{\alpha}(s, r))\bar{\alpha}_s(s, r)f_{0,s}(s, r). \end{aligned}$$

These equations reflect the fact that the agent has a choice about how to gather information. If the agent would like her signal/posterior to be sensitive to s , she can either attend to those variables directly (choosing $|f_a|$ to be small), or learn about them via the actions of other agents (choosing $|f_a|$ to be large). The optimal f_a satisfies

$$f_a(s, r, \bar{\alpha}(s, r)) = \frac{\bar{\alpha}_s(s, r)}{1 + \bar{\alpha}_s(s, r)^2} f_{0,s}(s), \quad (24)$$

which is to say the agent will attend to \bar{a} to the extent that it is informative about s . The resulting reduced form divergence can be written, for any $\mu_0, \mu'_0 \in \mathcal{U}_0$, differentiable $\bar{\alpha} \in \mathcal{A}$, and some weakly differentiable f_0 satisfying $f_0(s)d\mu_0(s, r) = d\mu'_0(s, r)$ for all $(s, r) \in S \times R$, as

$$D_{FI-S\bar{A}}^{\mathcal{F}}(f_0||\mu_0; \bar{\alpha}) = \int_{S \times R} \frac{1}{1 + \bar{\alpha}_s(s, r)^2} \left(\frac{\partial \ln(f_0(s))}{\partial s} \right)^2 f_0(s) d\mu_0(s, r),$$

Again, all weakly differentiable f_0 satisfying the required condition will generate the same reduced form divergence, and hence for any μ'_0 such that a weakly differentiable Radon-Nikodym derivative with respect to μ_0 exists,

$$D_{FI-S\bar{A}}(\phi[\mu'_0, \bar{\alpha}]||\phi[\mu_0, \bar{\alpha}]) = D_{FI-S\bar{A}}^{\mathcal{F}}(f_0||\mu_0; \bar{\alpha}).$$

We again define the divergence by continuity if no Radon-Nikodym derivative exists. We verify

⁵⁵Because this cost function makes non-differentiable densities infinitely costly, mean-consistency implies this property must hold in equilibrium.

in the technical appendix, Section C that there is a signal structure that achieves

$$\bar{C}(\bar{\nu}, \nu, \mu) = E^\pi[D_{FI-S\bar{A}}(\mu' || \mu)]$$

for all admissible π (see that section for a definition of admissible strategies).

As with D_{SR} , the arguments in Hébert and Woodford [2021] show that the cost function $D_{S\bar{A}}$ described in the text is a discretized version of this cost function.

The reduced form cost function $D_{FI-S\bar{A}}$ and its discrete counterpart $D_{S\bar{A}}$ depend on $\bar{\alpha}$ because $\bar{\alpha}$ influences the optimal allocation of the agent's attention. This is the key point of this example: the agent's cost of acquiring information (in the reduced form representation) will depend on $\bar{\alpha}$ if the agent is acquiring that information in part by learning about the actions of the other agents and if some actions are more easily observed than others.

Technical Appendix

C Results for Continuous States

In this section we describe how some of our results can be extended to the case of an continuous state space, in which S and R are compact, convex subsets of \mathbb{R}^n . In particular, we will extend our existence and uniqueness results (Propositions 1 and 2), our sufficient conditions for the presence or absence of non-fundamental volatility (Proposition 6 and the “if” part of Proposition 7), and our sufficient conditions for efficiency (Proposition 5) to the continuous state case. Our other proofs all in various ways rely on the construction of counter-examples. These constructions are simplified by the assumption of finiteness for $S \times R$; we speculate that similar constructions are possible in the infinite state case.

We will also consider a larger action space, $A \subset \mathbb{R}^n$, while continuing to assume A is convex and compact (and hence that $\bar{A} = A$). We do this in part because extending our existence, uniqueness, and sufficiency results to multi-dimensional action spaces is relatively straightforward, and in part because we are motivated by the interpretation of \bar{A} as affecting agents’ utility via prices.

The key challenge in extending our existence, uniqueness, and sufficiency results to the infinite state case arises from the need to consider divergences that assign infinite cost to certain posteriors. For example, consider the KL divergence. In the finite state case, the full-support assumption on $\mu_0 \in \mathcal{U}_0$ was sufficient to ensure that all posteriors have finite cost. In the infinite state case, there are sequences of posteriors absolutely continuous with respect to prior such that the KL divergence approaches infinity. As a consequence, the expected value of the divergence under arbitrary strategies is no longer continuous. Put another way, the set of posteriors with finite cost is not compact. To resolve this issue, we will restrict the set of strategies available to the agent to ensure compactness.

The second challenge concerns the mean-consistency condition. To avoid the difficulties associated with conditional expectations in infinite-dimensional settings, we define strategies as measures on posterior densities, as opposed to measures on measures. Specifically, let \mathcal{F}_0 be the set of continuous functions $S \times R \rightarrow \mathbb{R}$ satisfying, for all $f_0 \in \mathcal{F}_0$,

$$\int_{S \times R} f_0(s, r) d\mu_0(s, r) = 1,$$

and let $\mathcal{F}_0^+ \subset \mathcal{F}_0$ be the set of admissible densities. We describe our assumptions concerning \mathcal{F}_0^+ below.

The strategy space for the agent is $\bar{\Pi}_0 \subset \Delta(A \times \mathcal{F}_0^+)$, the set of measures satisfying Bayes-

consistency: for all $\bar{\pi} \in \bar{\Pi}_0$ and all $(s, r) \in S \times R$,

$$1 = \int_{\text{supp}(\bar{\pi}_0)} f_0(s, r) d\bar{\pi}_0(a, f_0).$$

The agents problem is to solve

$$\bar{\pi}_0^* \in \arg \max_{\bar{\pi}_0 \in \bar{\Pi}_0} E^{\bar{\pi}_0} [V(a, \bar{\phi}[f_0, \bar{\alpha}]) - D(\bar{\phi}[f_0, \bar{\alpha}] || \phi[\mu_0, \bar{\alpha}])]$$

which is a version of the alternative formulation of the agent's problem (Definition (4)) for densities. Recall that $\bar{\phi}$ is defined in Appendix Section B. Mean-consistency requires

$$\bar{\alpha}(s, r) = \int_{\text{supp}(\bar{\pi}_0)} a f_0(s, r) d\bar{\pi}_0(a, f_0).$$

Recall here that $\bar{\alpha}(s, r)$ is a vector in R^n for each $(s, r) \in S \times R$. This relatively simple definition of mean-consistency is one of the chief benefits of considering densities instead of measures. Let $\bar{\mathcal{A}}^+ \subseteq \bar{\mathcal{A}}$ be the set of aggregate action functions that can be generated by some $\bar{\pi}_0 \in \bar{\Pi}_0$.

An equilibrium is a mean-consistent strategy profile $(\bar{\pi}_0, \bar{\alpha})$ such that $\bar{\pi}_0$ is a best reply to $\bar{\alpha}$. Similarly, a strategy profile $(\bar{\pi}_0^*, \bar{\alpha}^*)$ is constrained efficient if it solves

$$(\bar{\pi}_0^*, \bar{\alpha}^*) \in \arg \max_{\bar{\pi}_0 \in \bar{\Pi}_0, \bar{\alpha} \in \bar{\mathcal{A}}} E^{\bar{\pi}_0} [V(a, \bar{\phi}[f_0, \bar{\alpha}]) - D(\bar{\phi}[f_0, \bar{\alpha}] || \phi[\mu_0, \bar{\alpha}])]$$

subject to mean-consistency.

A strategy profile $(\bar{\pi}_0, \bar{\alpha})$ is s -measurable if, for all $s \in S$ and $r, r' \in R$, $\bar{\alpha}(s, r) = \bar{\alpha}(s, r')$, and if $(a, f_0) \in \text{supp}(\bar{\pi}_0)$, there exists a function $f : S \times \bar{\mathcal{A}} \rightarrow \mathbb{R}_+$ such that $f(s, \bar{a}) d\phi[\mu_0, \bar{\alpha}](s, r, \bar{a}) = d\bar{\phi}[f_0, \bar{\alpha}](s, r, \bar{a})$ for all $(s, r, \bar{a}) \in S \times R \times \bar{\mathcal{A}}$.

We will assume in what follows that D is sufficiently continuous (the analog of Assumption 1 in the main text). Define \mathcal{U}^+ as the range of $\bar{\phi}$, which is the relevant domain of the problem:

$$\mathcal{U}^+ = \{\mu \in \mathcal{U} : \exists (f_0, \bar{\alpha}) \in \mathcal{F}_0^+ \times \bar{\mathcal{A}}^+ \text{ s.t. } \bar{\phi}[f_0, \bar{\alpha}] = \mu\}.$$

Note that for the existence and sufficiency results, it is not necessary to assume differentiability. Note also that, by definition, for any $\bar{\alpha} \in \bar{\mathcal{A}}$, $f_0 \in \mathcal{F}_0^+$,

$$D(\bar{\phi}[f_0, \bar{\alpha}] || \phi[\mu_0, \bar{\alpha}]) = D^{\mathcal{F}}(f_0 || \mu_0; \bar{\alpha}).$$

Assumption 2. For the continuous state case, we impose the follow assumptions.

1. The reduced form divergence $D^{\mathcal{F}}(f_0 || \mu_0; \bar{\alpha})$ is jointly continuous on $\mathcal{F}_0^+ \times \mathcal{A}^+$.
2. The sets of admissible densities \mathcal{F}_0^+ and aggregate action strategies $\bar{\mathcal{A}}^+$ are convex, non-empty, and compact in the sup-norm topology. The set \mathcal{F}_0^+ contains the function that is equal to one everywhere.

3. For any $f_0 \in \mathcal{F}_0^+$ and s -measurable $\bar{\alpha} \in \bar{\mathcal{A}}$, there exists an $f'_0 \in \mathcal{F}_0^+$ such that $\bar{\phi}[f'_0, \bar{\alpha}] = \eta_{-R|R}[\bar{\phi}[f_0, \bar{\alpha}] | \phi[\mu_0, \bar{\alpha}]]$.

The second assumption ensures the agent's strategy space is compact, while the third guarantees that for any non- s -measurable strategy, we can construct an s -measurable version using admissible strategies. These two properties were guaranteed by the assumption of a finite state space, but can apply to continuous state spaces as well.

Observe also that the operations $\eta_{-R|R}$ and $\eta_{-\bar{A}|\bar{A}}$ are well-defined in the continuous state case, and consequently the definitions of monotonicity in R and invariance in \bar{A} apply without modification to the infinite state case.

The following proposition generalizes our existence, sufficiency, and uniqueness results to the infinite-dimensional setting.

Proposition 9. *In the continuous state setting, assuming $v : A \times \bar{A} \times S \rightarrow \mathbb{R}$ is a continuous function,*

1. *An equilibrium exists,*
2. *If the divergence D is monotone in R , an s -measurable equilibrium exists,*
3. *If C is nowhere- R -monotone for all $\mu \in \mathcal{U}$ such that $\gamma_{-\bar{A}}[\mu] = \mu_0$, then all equilibria are either not s -measurable or have zero information acquisition,*
4. *If u is mean-critical and the divergence D is invariant in \bar{A} , a constrained efficient equilibrium exists.*
5. *If the divergence D is invariant in \bar{A} and u is mean-critical with an associated G function that is concave, then all equilibria are constrained efficient. If in addition G is strictly concave, then in any pair of equilibria $(\bar{\pi}_{0,1}, \bar{\alpha}_1)$ and $(\bar{\pi}_{0,2}, \bar{\alpha}_2)$, $\bar{\alpha}_1 = \bar{\alpha}_2$.*

Proof. See the technical appendix, section G. □

C.1 Examples

The KL Divergence To satisfy compactness in the sup-norm topology, restricting attention to Lipschitz-continuous functions is sufficient. For the KL divergence, define \mathcal{F}_0^+ as the set of K -Lipschitz-continuous functions for some $K > 0$, and observe that $\bar{\mathcal{A}}^+$ inherits Lipschitz-continuity from \mathcal{F}_0^+ . The second part of Assumption 2 is therefore satisfied by construction. Moreover, the operation $f'_0(s) = E^{\mu_0}[f_0(s, r) | r]$ preserves Lipschitz-continuity, and hence the third part of the assumption is satisfied.

Let us now suppose there exists a sequence $(f_{0,n}, \bar{\alpha}_n) \in \mathcal{F}_0^+ \times \bar{\mathcal{A}}^+$ that converges to some $(f_0, \bar{\alpha}) \in \mathcal{F}_0^+ \times \bar{\mathcal{A}}^+$ in the sup-norm topology. By assumption, the functions in \mathcal{F}_0^+ are uniformly

bounded, and therefore by the dominated convergence theorem $D_{KL}^{\mathcal{F}}$ is continuous,

$$\int_{S \times R} f_{0,n}(s, r) \ln(f_{0,n}(s, r)) d\mu_0(s, r) \rightarrow \int_{S \times R} f_0(s, r) \ln(f_0(s, r)) d\mu_0(s, r).$$

We conjecture, but do not prove, that choosing K sufficiently large enough is equivalent, in terms of the resulting equilibria, to allowing $K = \infty$. That is, for K sufficiently large, we expect that the constraint on Lipschitz-continuity is not in fact binding. We base this conjecture on the logit-type optimal policies generated by the KL divergence (see, e.g., [Matějka et al. \[2015\]](#)).

The Fisher Information Divergences For these cost functions, let \mathcal{F}_0^+ be the set of differentiable functions with a K -Lipschitz-continuous derivative that satisfy $\min_{(s,r) \in S \times R} f_0(s, r) \geq \epsilon$ for some $\epsilon > 0$ and $K > 0$. Such functions have a bounded derivative, and therefore Lipschitz-continuous. Note again that $\bar{\mathcal{A}}^+$ inherits Lipschitz-continuity from \mathcal{F}_0^+ . The second part of Assumption 2 is therefore satisfied by construction. Again, the operation $f_0'(s) = E^{\mu_0}[f_0(s, r)|r]$ preserves these properties, and hence the third part of the assumption is satisfied.

Let us now suppose there exists a sequence $(f_{0,n}, \bar{\alpha}_n) \in \mathcal{F}_0^+ \times \bar{\mathcal{A}}^+$ that converges to some $(f_0, \bar{\alpha}) \in \mathcal{F}_0^+ \times \bar{\mathcal{A}}^+$ in the sup-norm topology. The quantity

$$\|\nabla \ln [f_0(s, r)]\|^2 f_0(s, r)$$

is bounded above by $\epsilon^{-1} K^2$. By the dominated convergence theorem $D_{FI-SR}^{\mathcal{F}}(f_{0,n} || \mu_0; \bar{\alpha}_n)$ converges to $D_{FI-SR}^{\mathcal{F}}(f_0 || \mu_0; \bar{\alpha})$. Consequently, $D_{FI-SR}^{\mathcal{F}}$ satisfies the required continuity property.

Likewise, the quantity

$$\frac{1}{1 + \bar{\alpha}_s(s, r)^2} \left(\frac{\partial \ln(f_0(s))}{\partial s} \right)^2 f_0(s)$$

is bounded above by $\epsilon^{-1} K^2$, and again by the dominated convergence theorem $D_{FI-S\bar{\mathcal{A}}}^{\mathcal{F}}(f_{0,n} || \mu_0; \bar{\alpha}_n)$ converges to $D_{FI-S\bar{\mathcal{A}}}^{\mathcal{F}}(f_0 || \mu_0; \bar{\alpha})$.

We conjecture, based on results in [Hébert and Woodford \[2021\]](#) (who show that with Fisher information optimal policies are twice-differentiable and bounded away from zero in certain examples) that for sufficiently large K and small ϵ these constraints do not bind.

C.2 Fisher Information and the Allocation of Attention

We conclude with the following lemma, which shows that for both of these cost functions (D_{FI-SR} and $D_{FI-S\bar{\mathcal{A}}}$), the optimal posteriors are achievable given a sufficiently rich space of signal realizations. This lemma uses the admissible space of densities \mathcal{F}_0^+ , defined just above for these functions. It demonstrates the existence of signal structures (the $(\bar{\nu}, \nu) \in \mathcal{V}$ referenced in the lemma) as defined in Appendix Section C that achieve the required measure on posteriors. Note that in this lemma we assume a Lipschitz-continuous $\bar{\alpha} \in \bar{\mathcal{A}}$; given our definition of \mathcal{F}_0^+ , all mean-consistent $\bar{\alpha}$ are Lipschitz-continuous.

Lemma 5. Assume $\Omega = A \times \mathcal{F}_0^+$. For any $\bar{\pi}_0 \in \bar{\Pi}_0(\mu_0)$ and Lipschitz-continuous $\bar{\alpha} \in \bar{\mathcal{A}}$, let $\pi \in \Pi(\mu_0)$ be the measure induced from $\bar{\pi}_0$ by $(a, f_0) \mapsto (a, \bar{\phi}[f_0, \bar{\alpha}])$. If \bar{C} is the Fisher information restricted to $S \times R$, there exists a $(\bar{\nu}, \nu) \in \mathcal{V}$ such that

$$\bar{C}(\bar{\nu}, \nu, \mu) = E^\pi[D_{FI-SR}(\mu' || \mu)] = E^{\bar{\pi}_0}[D_{FI-SR}(\bar{\phi}[f_0, \bar{\alpha}] || \phi[\mu_0, \bar{\alpha}])].$$

If \bar{C} is the Fisher information and if $R = \{r\}$ is a singleton, there exists a $(\bar{\nu}, \nu) \in \mathcal{V}$ such that

$$\bar{C}(\bar{\nu}, \nu, \mu) = E^\pi[D_{FI-S\bar{A}}(\mu' || \mu)] = E^{\bar{\pi}_0}[D_{FI-S\bar{A}}(\bar{\phi}[f_0, \bar{\alpha}] || \phi[\mu_0, \bar{\alpha}])].$$

Proof. See the technical appendix, H.11. □

D Details for the LQG Model

D.1 Derivation of the Objective Function

By assumption, given the posterior μ'_0 the agent takes the action

$$a^*(\mu'_0) = E^{\mu'_0}[(s, r) \cdot \psi(\bar{\alpha}_s, \bar{\alpha}_r)].$$

The agent receives a signal $\omega \sim N((s, r) \cdot \lambda_0, \tau_0^{-1})$. By the standard Bayesian updating formula, her posterior precision matrix is

$$\Lambda = (\Sigma_0^{-1} + \tau_0 \lambda_0 \lambda'_0).$$

Her posterior mean concerning $(s, r) \cdot \lambda_0$ is $(\tau_0 + (\lambda'_0 \Sigma_0 \lambda_0)^{-1})^{-1} \tau_0 \omega$, and her posterior mean concerning any z such that $z' \Sigma_0 \lambda_0 = 0$ is unchanged (and hence zero). Consequently, her posterior mean concerning (s, r) is

$$E[(s, r) | \omega] = \frac{\Sigma_0 \lambda_0}{\lambda'_0 \Sigma_0 \lambda_0} \frac{\tau_0 \lambda'_0 \Sigma_0 \lambda_0}{1 + \tau_0 \lambda'_0 \Sigma_0 \lambda_0} \omega.$$

By the Sherman-Morrison lemma,

$$\Lambda^{-1} = \Sigma_0 - \frac{\tau_0 \Sigma_0 \lambda_0 \lambda'_0 \Sigma_0}{1 + \tau_0 \lambda'_0 \Sigma_0 \lambda_0}. \tag{25}$$

which yields

$$\mu'_0(\omega) = N\left(\omega \frac{\Sigma_0 \lambda_0}{\lambda'_0 \Sigma_0 \lambda_0} \frac{\tau_0 \lambda'_0 \Sigma_0 \lambda_0}{1 + \tau_0 \lambda'_0 \Sigma_0 \lambda_0}, \Sigma_0 - \frac{\tau_0 \Sigma_0 \lambda_0 \lambda'_0 \Sigma_0}{1 + \tau_0 \lambda'_0 \Sigma_0 \lambda_0}\right)$$

and therefore

$$\alpha(\lambda_0, \tau_0, \bar{\alpha}_s, \bar{\alpha}_r) = \frac{\psi(\bar{\alpha}_s, \bar{\alpha}_r)' \Sigma_0 \lambda_0}{\lambda'_0 \Sigma_0 \lambda_0} \frac{\tau_0 \lambda'_0 \Sigma_0 \lambda_0}{1 + \tau_0 \lambda'_0 \Sigma_0 \lambda_0}.$$

The linear-quadratic utility function in (2) can be written as

$$\begin{aligned} v(a, \bar{a}, s) &= -(a - \beta s - (1 - \beta)\bar{a})^2 \\ &\quad + \beta^2 s - \beta s^2 + (1 - \beta)^2 \bar{a}^2 - (1 - \beta)\bar{a}^2 \\ &\quad + 2\beta(1 - \beta)s\bar{a}, \end{aligned}$$

which yields

$$E^{\mu'_0}[u(a^*(\mu'_0), \vec{\alpha} \cdot (s, r), s)] = -Var^{\mu'_0}[(s, r) \cdot \psi(\bar{\alpha}_s, \bar{\alpha}_r)] - E^{\mu'_0}[\beta(1 - \beta)(s - \bar{\alpha}_s s - \bar{\alpha}_r r)^2]$$

where $Var[\cdot]$ denotes the variance. We have by (25)

$$Var^{\mu'_0}[(s, r) \cdot \psi(\bar{\alpha}_s, \bar{\alpha}_r)] = \psi(\bar{\alpha}_s, \bar{\alpha}_r)' \Sigma_0 \psi(\bar{\alpha}_s, \bar{\alpha}_r) - \frac{\tau_0 (\psi(\bar{\alpha}_s, \bar{\alpha}_r)' \Sigma_0 \lambda_0)^2}{1 + \tau_0 \lambda_0' \Sigma_0 \lambda_0}$$

The equation

$$\begin{aligned} E^{\pi_0}[E^{\mu'_0}[v(a, \bar{a}, s)]] &= \frac{\tau_0 (\psi(\bar{\alpha}_s, \bar{\alpha}_r)' \Sigma_0 \lambda_0)^2}{1 + \tau_0 \lambda_0' \Sigma_0 \lambda_0} \\ &\quad - \psi(\bar{\alpha}_s, \bar{\alpha}_r)' \Sigma_0 \psi(\bar{\alpha}_s, \bar{\alpha}_r) - \beta(1 - \beta) E^{\mu_0}[(s - \bar{\alpha}_s s - \bar{\alpha}_r r)^2] \end{aligned}$$

follows. Note that this derivation assumes Σ_0 is invertible. However, the resulting expression for expected utility is continuous in the limit as $\Sigma_0 \rightarrow \begin{bmatrix} \sigma_s^2 & 0 \\ 0 & 0 \end{bmatrix}$ and could be derived in analogous fashion for that case.

D.2 Derivations of the Cost Function Examples

Mutual Information The mutual information between the signal ω and state (s, r) can be derived using textbook methods (Cover and Thomas [2012]).

In our context, because the signal ω is one-dimensional, the mutual information is simply the mutual information between ω and $(s, r) \cdot \lambda_0$. As a shortcut, let us start from section 3.1 of Sims [2010], in our notation. We have

$$C_{MI}(\lambda_0, \tau_0, \Sigma_0, \bar{\alpha}_s, \bar{\alpha}_r) = -\theta \ln(1 - \rho(\lambda_0, \tau_0, \Sigma_0)^2),$$

where $\rho(\lambda_0, \tau_0, \Sigma_0)$ is the correlation between ω and $(s, r) \cdot \lambda_0$,

$$\rho(\lambda_0, \tau_0, \Sigma_0) = \frac{\lambda_0' \Sigma_0 \lambda_0}{\sqrt{\tau_0^{-1} + \lambda_0' \Sigma_0 \lambda_0} \sqrt{\lambda_0' \Sigma_0 \lambda_0}},$$

which yields

$$C_{MI}(\lambda_0, \tau_0, \Sigma_0, \bar{\alpha}_s, \bar{\alpha}_r) = -\theta \ln\left(1 - \frac{\lambda_0' \Sigma_0 \lambda_0}{\tau_0^{-1} + \lambda_0' \Sigma_0 \lambda_0}\right),$$

or

$$C_{MI}(\lambda_0, \tau_0, \Sigma_0, \bar{\alpha}_s, \bar{\alpha}_r) = \theta \ln(1 + \tau_0 \lambda_0' \Sigma_0 \lambda_0).$$

Note again that this formula remains valid even if Σ_0 is degenerate.

Fisher Information By corollary 3 of Hébert and Woodford [2021], the Fisher information cost function in the linear-quadratic-Gaussian case is, in our context,

$$C_{FI-SR}(\lambda_0, \tau_0, \Sigma_0, \bar{\alpha}_s, \bar{\alpha}_r) = \theta(\text{tr}[\Lambda] - \text{tr}[\Sigma_0^{-1}]),$$

where Λ is the posterior precision matrix, which is

$$C_{FI-SR}(\lambda_0, \tau_0, \Sigma_0, \bar{\alpha}_s, \bar{\alpha}_r) = \theta \tau_0 \text{tr}[\lambda_0 \lambda_0'] = \theta \tau_0 \lambda_0' \lambda_0.$$

Similarly,

$$C_{FI-S\bar{A}}(\lambda_0, \tau_0, \Sigma_0, \bar{\alpha}_s, \bar{\alpha}_r) = \frac{\theta}{1 + \bar{\alpha}_s^2} \tau_0 \lambda_0' \lambda_0.$$

Again, these formulas remain valid even in the degenerate limit.

D.3 Formal Results

Below we state formally our results on the existence of equilibria in the LQG setting.

Proposition 10. (i) *With the mutual information cost C_{MI} , there exists a linear equilibrium with zero non-fundamental volatility: $\bar{\alpha}_r = 0$, and all linear equilibria feature zero non-fundamental volatility.* (ii) *With the Fisher information cost C_{FI-SR} , there exists a linear equilibrium. Any such equilibrium either features non-fundamental volatility, $\bar{\alpha}_r \neq 0$, or features no information acquisition, $\tau_0^* = 0$.*

Proof. See the technical appendix, F.3. □

We next prove that in the context of Section 7.3, with the mutual information cost, the equilibrium with information acquisition Pareto-dominates the equilibrium without information acquisition.

Proposition 11. *In the setting of Section 7.3, with the mutual information cost C_{MI} , $\beta \in (\frac{1}{2}, 1)$, and $\theta = (1 - \beta)^2 \sigma_s^2$, there are two symmetric equilibrium, one of which features zero information acquisition, and the equilibrium with information acquisition Pareto-dominates the equilibrium without information acquisition.*

Proof. See the technical appendix, F.4. □

E Proofs

E.1 Additional Technical Lemmas

Lemma 6. *The reduced form divergences D_{KL} , $D_{S\bar{A}}$, and D_{SR} satisfy Assumption 1.*

Proof. Continuous differentiability on the interior follows immediately from the fact that all three functions are continuously differentiable in $\mu'_0(s, r)$ and $\bar{\alpha}(s, r)$ for each $(s, r) \in S \times R$. All three functions are defined on the boundary of the simplex by continuity, and therefore by construction are continuous provided they are finite. Finiteness follows from the assumption of full-support for μ_0 , the compactness of $S \times R$, and the boundedness of $\mu'_0(s, r)$, which together demonstrate the boundedness of both the conditional and unconditional KL divergences. \square

Lemma 7. *The set $\Pi_0(\mu_0)$ is non-empty, compact, and convex.*

Proof. See technical appendix, H.1. \square

Lemma 8. *The function $\phi : \mathcal{U}_0 \times \bar{\mathcal{A}} \rightarrow \mathcal{U}$ is continuous.*

Proof. See technical appendix, H.2. \square

Lemma 9. *The expected net utility function $J_0 : \Pi_0(\mu_0) \times \bar{\mathcal{A}} \rightarrow \mathbb{R}$ defined by*

$$J_0(\pi_0, \bar{\alpha}) = E^{\pi_0}[V(a, \phi\{\mu'_0, \bar{\alpha}\}) - D(\phi[\mu'_0, \bar{\alpha}] || \phi[\mu_0, \bar{\alpha}])]$$

is continuous in both its arguments, and concave in π .

Proof. See technical appendix, H.3. \square

Lemma 10. *The function $M : \Pi_0(\mu_0) \rightarrow \bar{\mathcal{A}}$ defined by $M[\pi_0](s, r) = E^{\pi_0}[a|s, r]$ for all $(s, r) \in S \times R$ is continuous.*

Proof. See technical appendix, H.4. \square

Lemma 11. *Define M as Lemma 10 and J_0 as in Lemma 9. A solution to the social planner's problem, $\max_{\pi_0 \in \Pi_0(\mu_0)} J_0(\pi_0, M[\pi_0])$, exists.*

Proof. The set $\Pi_0(\mu_0)$ is compact, $M(\pi_0)$ is continuous, and $J_0(\pi_0, \bar{\alpha})$ is continuous (by Lemmas 7, 9, and 10). \square

Lemma 12. *Define M as Lemma 10 and J_0 as in Lemma 9. For any $\pi_0 \in \Pi_0(\mu_0)$, if D is locally invariant in $\bar{\mathcal{A}}$ at $M[\pi_0]$ (Definition 10), then for all $\pi'_0 \in \Pi_0(\mu_0)$,*

$$\lim_{\epsilon \rightarrow 0^+} \frac{1}{\epsilon} (J_0((1 - \epsilon)\pi_0 + \epsilon\pi'_0, M[(1 - \epsilon)\pi_0 + \epsilon\pi'_0]) - J_0(\pi_0, M[\pi_0])) = J_0(\pi'_0, M[\pi_0]) - J_0(\pi_0, M[\pi_0]).$$

Proof. This result follows the observation that

$$E^{\pi_0}[V(a, \phi[\mu'_0, \bar{\alpha}])]$$

is directionally differentiable in $\bar{\alpha}$ (which follows from the mean-critical property, and note this derivative is zero) and the linearity of conditional expectations,

$$\lim_{\epsilon \rightarrow 0^+} \frac{1}{\epsilon} (M[(1 - \epsilon)\pi_0 + \epsilon\pi'_0] - M[\pi_0]) = M[\pi'_0] - M[\pi_0],$$

along with the local invariance in \bar{A} (which ensures that the expected divergence is continuously differentiable with respect to $\bar{\alpha}$, with a derivative of zero). \square

Lemma 13. *Define M as Lemma 10 and J_0 as in Lemma 9. If $(\pi_0^*, \bar{\alpha}^*)$ is an equilibrium and D is locally invariant in \bar{A} at $M[\pi_0]$, then for all $\pi'_0 \in \Pi_0(\mu_0)$,*

$$\lim_{\epsilon \rightarrow 0^+} \frac{1}{\epsilon} (J_0((1 - \epsilon)\pi_0^* + \epsilon\pi'_0, M[(1 - \epsilon)\pi_0^* + \epsilon\pi'_0]) - J_0(\pi_0^*, M[\pi_0^*])) \leq 0.$$

Proof. By the definition of an equilibrium, π_0^* is a maximum of $J_0(\cdot, \bar{\alpha}^*) : \Pi_0(\mu_0) \rightarrow \mathbb{R}$. By the definition of mean-consistency, $\bar{\alpha}^* = M[\pi_0^*]$. It follows by Lemma 12 that, for any $\pi'_0 \in \Pi_0(\mu_0)$,

$$\begin{aligned} \lim_{\epsilon \rightarrow 0^+} \frac{1}{\epsilon} (J_0((1 - \epsilon)\pi_0^* + \epsilon\pi'_0, M[(1 - \epsilon)\pi_0^* + \epsilon\pi'_0]) - J_0(\pi_0^*, M[\pi_0^*])) = \\ J_0(\pi'_0, M[\pi_0^*]) - J_0(\pi_0^*, M[\pi_0^*]) \leq 0. \end{aligned}$$

\square

Lemma 14. *Define J_0 as in Lemma 9, and assume D is invariant in \bar{A} . If G is concave in \bar{a} for all $s \in S$, then for any mean-consistent pairs $(\pi_{0,1}, \bar{\alpha}_1)$ and $(\pi_{0,2}, \bar{\alpha}_2)$ and any $\lambda \in (0, 1)$, $(\lambda\pi_{0,1} + (1 - \lambda)\pi_{0,2}, \lambda\bar{\alpha}_1 + (1 - \lambda)\bar{\alpha}_2) \in \Pi_0(\mu_0) \times \bar{A}$ and*

$$J_0(\lambda\pi_{0,1} + (1 - \lambda)\pi_{0,2}, \lambda\bar{\alpha}_1 + (1 - \lambda)\bar{\alpha}_2) \geq \lambda J_0(\pi_{0,1}, \bar{\alpha}_1) + (1 - \lambda) J_0(\pi_{0,2}, \bar{\alpha}_2),$$

strictly if $\bar{\alpha}_1 \neq \bar{\alpha}_2$ and G is strictly concave in \bar{a} for all $s \in S$.

Proof. Plugging in the functional form of a mean-critical utility function, (1), and invoking invariance in \bar{A} , for any $\bar{\alpha}' \in \bar{A}$,

$$J_0(\pi_0, \bar{\alpha}) = E^{\pi_0}[g(a; s) - D(\phi[\mu'_0, \bar{\alpha}'] || \phi[\mu_0, \bar{\alpha}'])] + E^{\mu_0}[G(\bar{\alpha}(s, r); s)],$$

and the result follows from the linearity of expectations under the assumption that μ_0 has full support. \square

Lemma 15. (*The Maximum Theorem*) Let $\Pi_0^* : \bar{\mathcal{A}} \Rightarrow \Pi_0(\mu_0)$ be the best reply correspondence in the agent's problem,

$$\Pi_0^*(\bar{\alpha}) = \{\pi_0 \in \arg \max_{\pi'_0 \in \Pi_0(\mu_0)} J(\pi'_0, \bar{\alpha})\},$$

where J is defined as in Lemma 9. Π_0^* is non-empty, compact-valued, convex-valued, and upper hemi-continuous.

Proof. By Lemma 7, the correspondence $\Gamma : \bar{\mathcal{A}} \Rightarrow \Delta(A \times \mathcal{U})$ defined by $\Gamma(\bar{\alpha}) = \Pi_0(\mu_0)$ satisfies the conditions of the theorem of the maximum (continuity with non-empty compact values). In this context, continuity is trivial as $\Gamma(\bar{\alpha})$ does not actually depend on $\bar{\alpha}$. Therefore Π_0^* is non-empty, compact-valued, and upper hemi-continuous by Aliprantis and Border [2006] 17.31 (the maximum theorem) and Lemma 9. Π_0^* is convex-valued by the linearity of J in π_0 and the convexity of $\Pi_0(\mu_0)$ (Lemma 7). \square

E.2 Proof of Lemma 1

R-Monotonicity By the definition of R-monotonicity (Definition 5), it is immediate that if D is R-monotone,

$$E^{\pi'}[D(\mu'|\mu)] \leq E^\pi[D(\mu'|\mu)].$$

It is therefore sufficient to prove that if this condition holds for all $\pi \in \Pi(\mu)$, then D is R-monotone given μ .

Proof by contradiction: suppose there exists some $\mu''' \ll \mu$ such that

$$D(\mu'''|\mu) < D(\eta_{-R|R}[\mu'''|\mu]|\mu).$$

If this inequality holds, by the definition of \mathcal{U} it must hold for some μ'_0, μ_0 such that $\mu'_0 \ll \mu_0$ and, for some $\bar{\alpha} \in \bar{\mathcal{A}}$, $\mu''' = \phi[\mu'_0, \bar{\alpha}]$ and $\mu = \phi[\mu_0, \bar{\alpha}]$.

Let us consider a two-point distribution for $\pi_\epsilon \in \Pi(\mu)$, placing probability $\epsilon > 0$ on (a', μ''') and probability $1 - \epsilon$ on $(a'', \mu'' = \frac{\mu - \epsilon\mu'''}{1 - \epsilon})$ for some arbitrary $a', a'' \in A$. By the finiteness of $S \times R$ and $\mu'_0 \ll \mu_0$, μ'' lies in the relative interior of the simplex for sufficiently small ϵ . By construction, this distribution satisfies Bayes' consistency. We have

$$E^{\pi_\epsilon}[D(\mu'|\mu)] = \epsilon D(\mu'''|\mu) + (1 - \epsilon) D\left(\frac{\mu - \epsilon\mu'''}{1 - \epsilon}|\mu\right),$$

and, defining π'_ϵ as the measure induced from π_ϵ by $(a, \mu') \mapsto (a, \eta_{-R|R}[\mu'|\mu])$,

$$\begin{aligned} E^{\pi'_\epsilon}[D(\mu'|\mu)] &= \epsilon D(\eta_{-R|R}[\mu'''|\mu]|\mu) + (1 - \epsilon) D(\eta_{-R|R}\left[\frac{\mu - \epsilon\mu'''}{1 - \epsilon}|\mu\right]|\mu) \\ &= \epsilon D(\eta_{-R|R}[\mu'''|\mu]|\mu) + (1 - \epsilon) D\left(\frac{\mu - \epsilon\eta_{-R|R}[\mu'''|\mu]}{1 - \epsilon}|\mu\right), \end{aligned}$$

where the last follows from the linearity of $\eta_{-R|R}$ in its first argument (see (7)). By assumption,

$D(\phi[\mu'_0, \bar{\alpha}]||\phi[\mu_0, \bar{\alpha}])$ is differentiable in μ'_0 on the interior of the simplex, and by definition D is minimized when $\mu' = \mu$, implying that the derivative is zero at this point. We have

$$\lim_{\epsilon \rightarrow 0^+} \frac{(1-\epsilon)}{\epsilon} D\left(\frac{\mu - \epsilon\mu'''}{1-\epsilon} || \mu\right) = \frac{D(\mu + \frac{\epsilon}{1-\epsilon}(\mu''' - \mu)||\mu) - D(\mu||\mu)}{\frac{\epsilon}{1-\epsilon}} = 0,$$

and likewise $\lim_{\epsilon \rightarrow 0^+} \frac{1-\epsilon}{\epsilon} D\left(\frac{\mu - \epsilon\eta_{-R|R}[\mu''']}{1-\epsilon} || \mu\right) = 0$. Therefore,

$$\begin{aligned} \lim_{\epsilon \rightarrow 0^+} \epsilon^{-1} (E^{\pi_\epsilon}[D(\mu' || \mu)] - E^{\pi'_\epsilon}[D(\mu' || \mu)]) = \\ D(\mu''' || \mu) - D(\eta_{-R|R}[\mu'''] || \mu) < 0, \end{aligned}$$

contradicting $E^{\pi'_\epsilon}[D(\mu' || \mu)] \leq E^{\pi_\epsilon}[D(\mu' || \mu)]$ for all $\epsilon \geq 0$.

\bar{A} -Invariance Next, consider invariance in \bar{A} . By the definition of invariance in \bar{A} , it is immediate that if D is invariant in \bar{A} ,

$$E^{\pi'}[D(\eta_{-\bar{A}|\bar{A}}[\mu' || \mu''] || \eta_{-\bar{A}|\bar{A}}[\mu || \mu''])] = E^\pi[D(\mu' || \mu)].$$

It is therefore sufficient to prove that if this condition holds for all $\mu'' \in \mathcal{U}$ with $\gamma_{-\bar{A}}[\mu] \ll \gamma_{-\bar{A}}[\mu'']$ and $\pi \in \Pi(\mu)$ (where π' is defined as in the statement of the lemma), then D is invariant in \bar{A} given μ .

Proof by contradiction: suppose the D is not invariant but such a condition holds.

Observe by definition that for $\mu' = \phi[\mu'_0, \bar{\alpha}]$ and $\mu'' = \phi[\mu''_0, \bar{\alpha}']$, $\eta_{-\bar{A}|\bar{A}}[\mu', \mu''] = \phi[\mu'_0, \bar{\alpha}']$. Consequently, if

$$D(\mu' || \mu) \neq D(\eta_{-\bar{A}|\bar{A}}[\mu', \mu''] || \eta_{-\bar{A}|\bar{A}}[\mu, \mu''])$$

for some $\mu, \mu', \mu'' \in \mathcal{U}$ with $\mu' \ll \mu$ and $\gamma_{-\bar{A}}[\mu] \ll \gamma_{-\bar{A}}[\mu'']$, then there exists some $\bar{\alpha}, \bar{\alpha}' \in \bar{\mathcal{A}}$ and $\mu'_0 \in \mathcal{U}_0$ such that

$$D(\phi[\mu'_0, \bar{\alpha}] || \phi[\mu_0, \bar{\alpha}]) \neq D(\phi[\mu'_0, \bar{\alpha}'] || \phi[\mu_0, \bar{\alpha}']).$$

Let us consider a two-point distribution for $\pi_\epsilon \in \Pi(\phi[\mu_0, \bar{\alpha}])$, placing probability $\epsilon > 0$ on $(a', \phi[\mu'_0, \bar{\alpha}])$ and probability $1 - \epsilon$ on $(a'', \mu'' = \phi[\frac{\mu_0 - \epsilon\mu'_0}{1-\epsilon}, \bar{\alpha}])$ for some arbitrary $a', a'' \in A$. By the finiteness of $S \times R$ and $\mu'_0 \ll \mu_0$, μ'' lies in the relative interior of the simplex for sufficiently small ϵ . By construction, this distribution satisfies Bayes' consistency. We have

$$E^{\pi_\epsilon}[D(\mu' || \phi[\mu_0, \bar{\alpha}])] = \epsilon D(\phi[\mu'_0, \bar{\alpha}] || \phi[\mu_0, \bar{\alpha}]) + (1 - \epsilon) D(\phi[\frac{\mu_0 - \epsilon\mu'_0}{1-\epsilon}, \bar{\alpha}] || \phi[\mu_0, \bar{\alpha}])$$

and, defining π'_ϵ as the measure induced from π_ϵ by $(a, \mu') \mapsto (a, \eta_{-\bar{A}|\bar{A}}[\mu'|\mu''])$,

$$\begin{aligned} E^{\pi'_\epsilon}[D(\mu'|\phi\{\mu_0, \bar{\alpha}'\})] &= \epsilon D(\phi[\mu'_0, \bar{\alpha}'|\phi[\mu_0, \bar{\alpha}']]) + (1 - \epsilon) D(\phi[\frac{\mu_0 - \epsilon\mu'_0}{1 - \epsilon}, \bar{\alpha}'|\phi[\mu_0, \bar{\alpha}']]) \\ &= \epsilon D(\phi[\mu'_0, \bar{\alpha}'|\phi[\mu_0, \bar{\alpha}']]) + (1 - \epsilon) D(\frac{\phi[\mu_0, \bar{\alpha}'] - \epsilon\phi[\mu'_0, \bar{\alpha}']}{1 - \epsilon}|\phi[\mu_0, \bar{\alpha}']]), \end{aligned}$$

where the last follows from the linearity of ϕ in its first argument. By assumption, $D(\phi[\mu'_0, \bar{\alpha}]|\phi[\mu_0, \bar{\alpha}])$ is differentiable in μ'_0 on the interior of the simplex, and by definition D is minimized when $\mu' = \mu$, implying that the derivative is zero at this point. Therefore,

$$\begin{aligned} \lim_{\epsilon \rightarrow 0^+} \epsilon^{-1} (E^{\pi_\epsilon}[D(\mu'|\phi[\mu_0, \bar{\alpha}])] - E^{\pi'_\epsilon}[D(\mu'|\phi[\mu_0, \bar{\alpha}'])]) &= \\ D(\phi[\mu'_0, \bar{\alpha}]|\phi[\mu_0, \bar{\alpha}]) - D(\phi[\mu'_0, \bar{\alpha}]|\phi[\mu_0, \bar{\alpha}']) &\neq 0, \end{aligned}$$

contradicting $E^{\pi'_\epsilon}[D(\mu'|\phi[\mu_0, \bar{\alpha}']]) = E^{\pi_\epsilon}[D(\mu'|\phi[\mu_0, \bar{\alpha}])] for all $\epsilon \geq 0$.$

E.3 Proof of Lemma 2

We prove non-monotonicity via a counter-example.

Let s_i denote the i -th smallest element of $S \subset \mathbb{R}$, and let r_j denote the j -th smallest element of $R \subset \mathbb{R}$. Assume an MLRP property: $\frac{\mu_0(s_i, r_j)}{\mu_0(s_{i-1}, r_j)}$ is strictly increasing in j for all $i > 1$. Define $f_1(s)$ as any strictly positive, non-constant function of s such that $\sum_{s \in S, r \in R} \mu_0(s, r) f_1(s) = 1$. Define i^* as the smallest i such that $f(s_i) \neq f(s_{i+1})$, which must exist by the assumption that $f_1(s)$ is not constant.

Consider the example posterior described in the text:

$$\mu'_0(s, r; \epsilon) = \begin{cases} \mu_0(s, r) f_1(s) (1 + \epsilon(r - E^{\mu_0}[r'|s])) & s = s_{i^*}, \\ \mu_0(s, r) f_1(s) & s \in S \setminus \{s_{i^*}\}. \end{cases}$$

Here and in what follows we use the notation $E^{\mu_0}[r'|s]$ to indicate

$$E^{\mu_0}[r'|s] = \frac{\sum_{r' \in R} r' \mu_0(s, r')}{\sum_{r' \in R} \mu_0(s, r')}.$$

Note that this is a probability measure provided that $\epsilon^{-1} \geq -\min_{r \in R} (r - E^{\mu_0}[r'|s = s_i^*]) > 0$, where the latter follows from the full-support property of μ_0 . Suppose without loss of generality that $\bar{\alpha}(s, r) = \bar{a}_0$ for some $\bar{a}_0 \in \bar{A}$ and all $(s, r) \in S \times R$ (recall that D_{SR} is invariant in \bar{A}). Consequently, for $\mu'(\epsilon) = \phi[\mu'_0(\epsilon), \bar{\alpha}]$ and $\mu = \phi[\mu_0, \bar{\alpha}]$,

$$d\mu'(s, r, \bar{a}; \epsilon) = \begin{cases} d\mu(s, r, \bar{a}) f_1(s) (1 + \epsilon(r - E^{\mu_0}[r'|s])) & s = s_{i^*}, \\ d\mu(s, r, \bar{a}) f_1(s) & s \in S \setminus \{s_{i^*}\}. \end{cases}$$

By construction,

$$d\gamma_{-R}[\mu'(\epsilon)](s, \bar{a}) = f_1(s) d\gamma_{-R}[\mu](s, \bar{a}).$$

Therefore,

$$\begin{aligned} d\eta_{-R|R}[\mu'(\epsilon)|\mu] &= f_1(s)d\mu(s, r, \bar{a}) \\ &= d\phi[\mu_0'', \bar{\alpha}], \end{aligned}$$

where $\mu_0''(s, r) = f_1(s)\mu_0(s, r)$. We can write this as $\eta_{-R|R}[\mu'(\epsilon)|\mu] = \mu'(0)$.

Now consider the conditional KL divergences given this μ_0' . We have

$$\begin{aligned} (\mu_0'(s_{i^*}, r_j; \epsilon) + \mu_0'(s_{i^*}, r_{j-1}; \epsilon))D_{KL, j, s_{i^*}}(\mu_0' || \mu_0) &= \\ \sum_{r \in \{r_j, r_{j-1}\}} \mu_0'(s_{i^*}, r; \epsilon) [\ln(1 + \epsilon(r - E^{\mu_0}[r'|s]))] & \\ - \ln(1 + \epsilon \frac{\mu_0(s_{i^*}, r_j)(r_j - E^{\mu_0}[r'|s]) + \mu_0(s_{i^*}, r_{j-1})\epsilon(r_{j-1} - E^{\mu_0}[r'|s])}{\mu_0(s_{i^*}, r_j) + \mu_0(s_{i^*}, r_{j-1})}) &]. \end{aligned}$$

Taking the derivative with respect to ϵ and evaluating at $\epsilon = 0$,

$$\begin{aligned} \frac{\partial}{\partial \epsilon} (\mu_0'(s_{i^*}, r_j; \epsilon) + \mu_0'(s_{i^*}, r_{j-1}; \epsilon))D_{KL, j, s_{i^*}}(\mu_0'(\epsilon) || \mu_0)|_{\epsilon=0} &= \\ f_1(s_{i^*}) \sum_{r \in \{r_j, r_{j-1}\}} \left\{ \frac{\mu_0(s_{i^*}, r)}{\mu_0(s_{i^*}, r_j) + \mu_0(s_{i^*}, r_{j-1})} \times \right. & \\ \left. \left[r - \frac{\mu_0(s_{i^*}, r_j)}{\mu_0(s_{i^*}, r_j) + \mu_0(s_{i^*}, r_{j-1})} r_j - \frac{\mu_0(s_{i^*}, r_{j-1})}{\mu_0(s_{i^*}, r_j) + \mu_0(s_{i^*}, r_{j-1})} r_{j-1} \right] \right\} &= 0. \end{aligned}$$

For any other value of s ($s \in S \setminus \{s_i^*\}$), by construction,

$$\frac{\partial}{\partial \epsilon} (\mu_0'(s, r_j; \epsilon) + \mu_0'(s, r_{j-1}; \epsilon))D_{KL, j, s}(\mu_0'(\epsilon) || \mu_0)|_{\epsilon=0} = 0.$$

We also have

$$\begin{aligned} (\mu_0'(s_{i^*+1}, r; \epsilon) + \mu_0'(s_{i^*}, r; \epsilon))D_{KL, i^*+1, r}(\mu_0' || \mu_0) &= \\ \mu_0(s_{i^*}, r)f_1(s_{i^*})(1 + \epsilon(r - E^{\mu_0}[r'|s = s_{i^*}])) \{ \ln(f_1(s_{i^*})(1 + \epsilon(r - E^{\mu_0}[r'|s = s_{i^*}])))) & \\ - \ln(\frac{\mu_0(s_{i^*+1}, r)f_1(s_{i^*+1}) + \mu_0(s_{i^*}, r)f_1(s_{i^*})(1 + \epsilon(r - E^{\mu_0}[r'|s = s_{i^*}]))}{\mu_0(s_{i^*+1}, r) + \mu_0(s_{i^*}, r)}) & \\ + \mu_0(s_{i^*+1}, r)f_1(s_{i^*+1}) \{ \ln(f_1(s_{i^*+1})) & \\ - \ln(\frac{\mu_0(s_{i^*+1}, r)f_1(s_{i^*+1}) + \mu_0(s_{i^*}, r)f_1(s_{i^*})(1 + \epsilon(r - E^{\mu_0}[r'|s = s_{i^*}]))}{\mu_0(s_{i^*+1}, r) + \mu_0(s_{i^*}, r)}) & \}. \end{aligned}$$

Taking the derivative with respect to ϵ and evaluating at $\epsilon = 0$,

$$\begin{aligned} & \frac{\partial}{\partial \epsilon} (\mu'_0(s_{i^*+1}, r; \epsilon) + \mu'_0(s_{i^*}, r; \epsilon)) D_{KL, i^*+1, r}(\mu'_0 || \mu_0) |_{\epsilon=0} = \\ & \mu_0(s_{i^*}, r) f_1(s_{i^*}) (r - E^{\mu_0}[r' | s = s_{i^*}]) (\ln(f_1(s_{i^*}))) - \ln\left(\frac{\mu_0(s_{i^*+1}, r) f_1(s_{i^*+1}) + \mu_0(s_{i^*}, r) f_1(s_{i^*})}{\mu_0(s_{i^*+1}, r) + \mu_0(s_{i^*}, r)}\right)) + \\ & \mu_0(s_{i^*}^*, r) f_1(s_{i^*}^*) (r - E^{\mu_0}[r' | s' = s_{i^*}^*]) - \\ & \sum_{s \in \{s_{i^*}^*, s_{i^*+1}\}} \mu_0(s, r) f_1(s) \frac{\mu_0(s_{i^*}, r) f_1(s_{i^*}) (r - E^{\mu_0}[r' | s' = s_{i^*}^*])}{\mu_0(s_{i^*+1}, r) f_1(s_{i^*+1}) + \mu_0(s_{i^*}, r) f_1(s_{i^*})}. \end{aligned}$$

The last two lines simplify to zero.

By a similar argument, if $i^* > 1$,

$$\begin{aligned} & \frac{\partial}{\partial \epsilon} (\mu'_0(s_{i^*}, r; \epsilon) + \mu'_0(s_{i^*-1}, r; \epsilon)) D_{KL, i^*, r}(\mu'_0 || \mu_0) |_{\epsilon=0} = \\ & \mu_0(s_{i^*}, r) f_1(s_{i^*}) (r - E^{\mu_0}[r' | s_{i^*}]) (\ln(f_1(s_{i^*}))) - \ln\left(\frac{\mu_0(s_{i^*}, r) f_1(s_{i^*}) + \mu_0(s_{i^*-1}, r) f_1(s_{i^*-1})}{\mu_0(s_{i^*}, r) + \mu_0(s_{i^*-1}, r)}\right)) + \\ & \mu_0(s_{i^*}^*, r) f_1(s_{i^*}^*) (r - E^{\mu_0}[r' | s' = s_{i^*}^*]) - \\ & \sum_{s \in \{s_{i^*-1}, s_{i^*}\}} \mu_0(s, r) f_1(s) \frac{\mu_0(s_{i^*}, r) f_1(s_{i^*}) (r - E^{\mu_0}[r' | s' = s_{i^*}^*])}{\mu_0(s_{i^*}, r) f_1(s_{i^*}) + \mu_0(s_{i^*-1}, r) f_1(s_{i^*-1})} = 0, \end{aligned}$$

which follows from the assumption that $f_1(s_{i^*}) = f_1(s_{i^*-1})$ if $i^* > 1$. By construction, for any $i \notin \{i^*, i^* + 1\}$,

$$\frac{\partial}{\partial \epsilon} (\mu'_0(s_i, r; \epsilon) + \mu'_0(s_{i-1}, r; \epsilon)) D_{KL, i, r}(\mu'_0 || \mu_0) |_{\epsilon=0} = 0.$$

It follows that

$$\begin{aligned} & \frac{\partial}{\partial \epsilon} D_{SR}(\mu'(\epsilon) || \mu) |_{\epsilon=0} = \\ & \sum_{j=1}^{|R|} \mu_0(s_{i^*}, r_j) (r_j - E^{\mu_0}[r' | s_{i^*}]) \frac{f_1(s_{i^*})}{s_{i^*+1} - s_{i^*}} [\ln(f_1(s_{i^*}^*)) - \ln\left(\frac{\mu_0(s_{i^*+1}, r) f_1(s_{i^*+1}) + \mu_0(s_{i^*}, r) f_1(s_{i^*})}{\mu_0(s_{i^*+1}, r) + \mu_0(s_{i^*}, r)}\right)]]. \end{aligned} \tag{26}$$

By the MLRP property,

$$\frac{\mu_0(s_{i^*+1}, r_j)}{\mu_0(s_{i^*}, r_j)}$$

is strictly increasing in j , and consequently if $f_1(s_{i^*+1}) > f_1(s_{i^*})$,

$$\frac{f_1(s_{i^*})}{s_{i^*+1} - s_{i^*}} [\ln(f_1(s_{i^*}^*)) - \ln\left(\frac{\mu_0(s_{i^*+1}, r) f_1(s_{i^*+1}) + \mu_0(s_{i^*}, r) f_1(s_{i^*})}{\mu_0(s_{i^*+1}, r) + \mu_0(s_{i^*}, r)}\right)]]$$

is strictly decreasing in j . Note that this result uses the assumption that $f_1(s_{i^*}) > 0$. If we instead had $f_1(s_{i^*}) = 0$, adopting the standard convention that $0 \ln(0) = 0$, this quantity is constant (and equal to zero). If instead $f_1(s_{i^*+1}) < f_1(s_{i^*})$, this quantity is strictly increasing in j . Returning

the case in which $f_1(s_{i^*+1}) > f_1(s_{i^*}) > 0$, it follows by the covariance rule that

$$\begin{aligned} \frac{\partial}{\partial \epsilon} D_{SR}(\mu'(\epsilon) || \mu) |_{\epsilon=0} < \\ \left(\sum_{j=1}^{|R|} \mu_0(s_{i^*}, r_j) (r_j - E^{\mu_0}[r' | s_{i^*}]) \right) \times \\ \left(\sum_{j=1}^{|R|} \mu_0(s_{i^*}, r_j) \frac{f_1(s_{i^*}^*)}{s_{i^*+1} - s_{i^*}} (\ln(f_1(s_{i^*}^*)) - \ln\left(\frac{\mu_0(s_{i^*+1}, r_j) f_1(s_{i^*+1}) + \mu_0(s_{i^*}, r_j) f_1(s_{i^*}^*)}{\mu_0(s_{i^*+1}, r_j) + \mu_0(s_{i^*}, r_j)}\right)) \right) \end{aligned}$$

which yields

$$\frac{\partial}{\partial \epsilon} D_{SR}(\mu'(\epsilon) || \mu) |_{\epsilon=0} < 0.$$

It follows that there exists an ϵ sufficiently close to zero such that

$$D_{SR}(\mu'(\epsilon) || \mu) < D_{SR}(\mu'(0) || \mu) = D_{SR}(\eta_R(\mu'(\epsilon), \mu) || \mu).$$

By the same argument, $D_{SR}(\mu'(\epsilon) || \mu) > D_{SR}(\eta_R(\mu'(\epsilon), \mu) || \mu)$ if $f_1(s_{i^*+1}) < f_1(s_{i^*})$. Likewise, $\frac{\partial}{\partial \epsilon} D_{SR}(\mu'(\epsilon) || \mu) |_{\epsilon=0} = 0$ if $f_1(s_{i^*+1}) > f_1(s_{i^*}) = 0$.

We now use this argument to prove that C_{SR} (the cost function associated with D_{SR}) is nowhere-R-monotone. Assume μ_0 satisfies the MLRP property above, define $\bar{\alpha}$ as the constant aggregate action strategy described above, and let $\mu = \phi[\mu_0, \bar{\alpha}]$. This fixes the prior for which we prove nowhere-R-monotonicity.

Take as given any non-zero-information, non-R-measurable strategy $\pi \in \Pi(\mu)$. Let $\pi_0 \in \Pi_0(\mu_0)$ be the strategy induced from π by $(a, \mu') \mapsto (a, \gamma_{-\bar{\alpha}}[\mu'])$.

By the assumption that π is a non-R-measurable strategy and the s-measurability of $\bar{\alpha}$, for each $(a, \mu'_0) \in \text{supp}(\pi_0)$ there is a function $f : S \rightarrow \mathbb{R}_+$ such that $f = \frac{d\mu'_0}{d\mu_0}$. Let $f(s; \mu'_0)$ denote this function.

For any $\omega \in [0, 1)$ and $i \in \{0, \dots, |S| - 1\}$, define the set $F_{i,\omega}^- \subset \text{supp}(\pi_0)$ as the set of (a, μ'_0) such that the associated f satisfies $f(s_{i+1}) < (1-\omega)f(s_i)$ and $f(s_i) \geq \omega$, adopting the convention that $f(s_0) = 1$ for some $s_0 < s_1$.

Define $i(\omega)$ as the smallest value of i for which $F_{i,\omega}^-$ has strictly positive measure under π_0 , with $i(\omega) = |S|$ if no such i exists. By the assumption that π is not zero-information, there exists an $\omega^* \in (0, 1)$ and $i^* = i(\omega^*) < |S|$ such that F_{i^*,ω^*}^- has strictly positive measure, as otherwise by Bayes-consistency π must be a zero-information strategy.

Let $F^- = F_{i^*,0}^- \supset F_{i^*,\omega^*}^-$. Define the set $F^+ \subset \text{supp}(\pi_0)$ as the set of (a, μ'_0) such that the associated f satisfies $f(s_{i^*+1}) \geq f(s_{i^*})$, Note that this set also has positive measure under π_0 , by Bayes-consistency.

Define the non-negative constants

$$\kappa^+ = E^{\pi_0}[f(s_{i^*}; \mu'_0) \mathbf{1}\{(a, \mu'_0) \in F^+\}]$$

and

$$\kappa^- = E^{\pi_0}[f(s_{i^*}; \mu'_0) \mathbf{1}\{(a, \mu'_0) \in F^-\}].$$

By Bayes-consistency,

$$\kappa^+ + \kappa^- = 1,$$

as $\text{supp}(\pi_0) = F^+ \cup F^-$.

Now observe, by the definition of i^* , that $\kappa^+ > 0$. If $i^* = 1$, then by definition $f(s_{i^*}; \mu'_0) \geq 1$ for all $(a, \mu'_0) \in F^+$. If $i^* > 1$, then $\kappa^+ = 0$ would imply that there is a positive measure of (a, μ'_0) such that $f(s_{i^*}; \mu'_0) = 0$. But in this case, there must have been some $i^{**} < i^*$ with a positive measure of $F_{i^{**}, \omega^{**}}$, a contradiction.

Consider the following mapping defined by some $\epsilon \geq 0$:

$$(a, \mu'_0) \rightarrow \begin{cases} (a, \mu'_0(-\epsilon\kappa^+)) & (a, \mu'_0) \in F^-, \\ (a, \mu'_0(\epsilon\kappa^-)) & (a, \mu'_0) \in F^+, \\ (a, \mu'_0) & \text{otherwise,} \end{cases}$$

where $\mu'_0(\epsilon) = \gamma_{-\bar{A}}[\mu'(\epsilon)]$, with $\mu'(\epsilon)$ defined as above. Let $\pi_0(\epsilon)$ be the measure induced from π_0 by this mapping, and let $\pi(\epsilon)$ be the measure induced from $\pi_0(\epsilon)$ by $(a, \mu'_0) \mapsto (a, \phi[\mu'_0, \bar{\alpha}])$.

Observe that

$$\begin{aligned} E^{\pi_0(\epsilon)}[\mu'_0(s, r)|s = s_{i^*}, r] &= E^{\pi_0}[\mu'_0(s_i^*, r) \mathbf{1}\{(a, \mu'_0) \in F^+\}|s = s_{i^*}, r](1 + \epsilon\kappa^-(r - E^{\mu_0}[r'|s = s_{i^*}])) \\ &\quad + E^{\pi_0}[\mu'_0(s_i^*, r) \mathbf{1}\{(a, \mu'_0) \in F^-\}|s = s_{i^*}, r](1 - \epsilon\kappa^+(r - E^{\mu_0}[r'|s = s_{i^*}])). \end{aligned}$$

This can be rewritten as

$$\begin{aligned} E^{\pi_0(\epsilon)}[\mu'_0(s, r)|s = s_{i^*}, r] &= \mu_0(s_i^*, r) E^{\pi_0}[f(s_i^*; \mu'_0) \mathbf{1}\{(a, \mu'_0) \in F^+\}](1 + \epsilon\kappa^-(r - E^{\mu_0}[r'|s = s_{i^*}])) \\ &\quad + \mu_0(s_i^*, r) E^{\pi_0}[f(s_i^*; \mu'_0) \mathbf{1}\{(a, \mu'_0) \in F^-\}](1 - \epsilon\kappa^+(r - E^{\mu_0}[r'|s = s_{i^*}])), \end{aligned}$$

which is

$$E^{\pi_0(\epsilon)}[\mu'_0(s, r)|s = s_{i^*}, r] = \mu_0(s_i^*, r)(\kappa^+ + \kappa^-) + \epsilon\mu_0(s_i^*, r)(r - E^{\mu_0}[r'|s = s_{i^*}])(\kappa^+\kappa^- - \kappa^-\kappa^+).$$

Note also by construction that $E^{\pi_0(\epsilon)}[\mu'_0(s, r)|s, r] = \mu_0(s_i^*, r)$ for any $s \neq s_{i^*}$. It follows that $\pi_0(\epsilon)$ is Bayes-consistent with μ_0 for all $\epsilon \geq 0$, and hence that $\pi(\epsilon)$ is Bayes-consistent with μ .

By the the continuous differentiability of D_{SR} in ϵ and the compactness of $A \times \mathcal{U}$, we can

differentiate under the integral sign (by multi-variate version of the Leibniz integral rule),

$$\begin{aligned} \frac{\partial}{\partial \epsilon} C_{SR}(\pi(\epsilon), \mu)|_{\epsilon=0} &= \kappa^+ E^{\pi(a, \mu')} [\mathbf{1}\{(a, \gamma_{\bar{A}}\{\mu'\}) \in F^-\} \frac{\partial}{\partial \epsilon} D_{SR}(\phi\{\mu'_0(-\epsilon), \bar{\alpha}\}||\mu)|_{\epsilon=0}] \\ &\quad + \kappa^- E^{\pi(a, \mu')} [\mathbf{1}\{(a, \gamma_{\bar{A}}\{\mu'\}) \in F^+\} \frac{\partial}{\partial \epsilon} D_{SR}(\phi\{\mu'_0(\epsilon), \bar{\alpha}\}||\mu)|_{\epsilon=0}]. \end{aligned}$$

By the results above, for $(a, \gamma_{\bar{A}}\{\mu'\}) \in F^+$,

$$\frac{\partial}{\partial \epsilon} D_{SR}(\phi\{\mu'_0(\epsilon), \bar{\alpha}\}||\mu)|_{\epsilon=0} \leq 0,$$

and for $(a, \gamma_{\bar{A}}\{\mu'\}) \in F^-$,

$$\frac{\partial}{\partial \epsilon} D_{SR}(\phi\{\mu'_0(\epsilon), \bar{\alpha}\}||\mu)|_{\epsilon=0} > 0,$$

and hence this derivative is weakly negative.

Observe that

$$\begin{aligned} E^{\pi(a, \mu')} [\mathbf{1}\{(a, \gamma_{\bar{A}}\{\mu'\}) \in F^-\} \frac{\partial}{\partial \epsilon} D_{SR}(\phi\{\mu'_0(-\epsilon), \bar{\alpha}\}||\mu)|_{\epsilon=0}] < \\ E^{\pi(a, \mu')} [\mathbf{1}\{(a, \gamma_{\bar{A}}\{\mu'\}) \in F_{i^*, \omega^*}^-\} \frac{\partial}{\partial \epsilon} D_{SR}(\phi\{\mu'_0(-\epsilon), \bar{\alpha}\}||\mu)|_{\epsilon=0}], \end{aligned}$$

and hence by $\kappa^+ > 0$ it is sufficient to prove that

$$\frac{\partial}{\partial \epsilon} D_{SR}(\phi\{\mu'_0(\epsilon), \bar{\alpha}\}||\mu)|_{\epsilon=0}$$

is bounded below by some $\bar{d} > 0$ on the set F_{i^*, ω^*}^- , observing that F_{i^*, ω^*}^- is a positive measure set by construction.

Using (26) for this set,

$$\begin{aligned} \frac{\partial}{\partial \epsilon} D_{SR}(\mu'(\epsilon)||\mu)|_{\epsilon=0} = \\ - \sum_{j=1}^{|R|} \mu_0(s_{i^*}, r_j) (r_j - E^{\mu_0}[r' | s_{i^*}]) \frac{f(s_{i^*})}{s_{i^*+1} - s_{i^*}} \left[\ln \left(\frac{\frac{\mu_0(s_{i^*+1}, r_j)}{\mu_0(s_{i^*}, r_j)} \frac{f(s_{i^*+1})}{f(s_{i^*})} + 1}{\frac{\mu_0(s_{i^*+1}, r_j)}{\mu_0(s_{i^*}, r_j)} + 1} \right) \right]. \end{aligned}$$

Note that

$$\begin{aligned} \frac{\partial}{\partial x} \sum_{j=1}^{|R|} \mu_0(s_{i^*}, r_j) (r_j - E^{\mu_0}[r' | s_{i^*}]) \frac{f(s_{i^*})}{s_{i^*+1} - s_{i^*}} \left[\ln \left(\frac{\frac{\mu_0(s_{i^*+1}, r_j)}{\mu_0(s_{i^*}, r_j)} x + 1}{\frac{\mu_0(s_{i^*+1}, r_j)}{\mu_0(s_{i^*}, r_j)} + 1} \right) \right] = \\ \sum_{j=1}^{|R|} \mu_0(s_{i^*}, r_j) (r_j - E^{\mu_0}[r' | s_{i^*}]) \frac{f(s_{i^*})}{s_{i^*+1} - s_{i^*}} \frac{\frac{\mu_0(s_{i^*+1}, r_j)}{\mu_0(s_{i^*}, r_j)}}{\frac{\mu_0(s_{i^*+1}, r_j)}{\mu_0(s_{i^*}, r_j)} x + 1}, \end{aligned}$$

and that

$$\frac{\frac{\mu_0(s_{i^*+1}, r_j)}{\mu_0(s_{i^*}, r_j)}}{\frac{\mu_0(s_{i^*+1}, r_j)}{\mu_0(s_{i^*}, r_j)} x + 1}$$

is increasing in j for all $x < 1$. It follows by the covariance rule that

$$\frac{\partial}{\partial x} \sum_{j=1}^{|R|} \mu_0(s_{i^*}, r_j) (r_j - E^{\mu_0}[r' | s_{i^*}]) \frac{f(s_{i^*})}{s_{i^*+1} - s_{i^*}} \left[\ln \left(\frac{\frac{\mu_0(s_{i^*+1}, r_j)}{\mu_0(s_{i^*}, r_j)} x + 1}{\frac{\mu_0(s_{i^*+1}, r_j)}{\mu_0(s_{i^*}, r_j)} + 1} \right) \right] \geq 0,$$

and hence, observing that $f(s_{i^*})(1 - \omega^*) > f(s_{i^*+1}) \geq 0$ on F_{i^*, ω^*}^- , that

$$\frac{\partial}{\partial \epsilon} D_{SR}(\mu'(\epsilon) | \mu) |_{\epsilon=0} \geq \sum_{j=1}^{|R|} \mu_0(s_{i^*}, r_j) (r_j - E^{\mu_0}[r' | s_{i^*}]) \frac{f(s_{i^*})}{s_{i^*+1} - s_{i^*}} \left[\ln \left(\frac{\frac{\mu_0(s_{i^*+1}, r_j)}{\mu_0(s_{i^*}, r_j)} (1 - \omega^*) + 1}{\frac{\mu_0(s_{i^*+1}, r_j)}{\mu_0(s_{i^*}, r_j)} + 1} \right) \right].$$

Noting that

$$\ln \left(\frac{\frac{\mu_0(s_{i^*+1}, r)}{\mu_0(s_{i^*}, r)} (1 - \omega^*) + 1}{\frac{\mu_0(s_{i^*+1}, r)}{\mu_0(s_{i^*}, r)} + 1} \right)$$

is itself decreasing in r , and that $f(s_{i^*}) \geq \omega^*$, $\bar{d} > 0$ is a uniform bound lower bound for $\frac{\partial}{\partial \epsilon} D_{SR}(\mu'(\epsilon) | \mu) |_{\epsilon=0}$ on the set F_{i^*, ω^*}^- , where

$$\bar{d} = - \sum_{j=1}^{|R|} \mu_0(s_{i^*}, r_j) (r_j - E^{\mu_0}[r' | s_{i^*}]) \frac{\omega^*}{s_{i^*+1} - s_{i^*}} \left[\ln \left(\frac{\frac{\mu_0(s_{i^*+1}, r_j)}{\mu_0(s_{i^*}, r_j)} (1 - \omega^*) + 1}{\frac{\mu_0(s_{i^*+1}, r_j)}{\mu_0(s_{i^*}, r_j)} + 1} \right) \right].$$

Consequently,

$$\frac{\partial}{\partial \epsilon} C_{SR}(\pi(\epsilon), \mu) |_{\epsilon=0} < -\kappa^+ \bar{d} < 0.$$

By construction, $(a, \mu') \mapsto (a, \eta_{-R|R}[\mu' | \mu])$ induces $\pi = \pi(0)$ from $\pi(\epsilon)$, and nowhere-R-monotonicity follows.

E.4 Proof of Proposition 1

We prove existence in the standard way: by using the theorem of the maximum to establish the upper hemi-continuity of best reply correspondences, and then use a fixed point theorem. This proof refers several of the additional lemmas found in Technical Appendix Section E.1.

The agent's problem is

$$\pi_0^* \in \arg \max_{\pi_0 \in \Pi_0(\mu_0)} J_0(\pi_0, \bar{\alpha}),$$

where $J_0(\pi_0, \bar{\alpha})$ is defined as

$$J_0(\pi_0, \bar{\alpha}) = E^{\pi_0} [V(a, \phi[\mu'_0, \bar{\alpha}]) - D(\phi[\mu'_0, \bar{\alpha}] | \phi[\mu_0, \bar{\alpha}])],$$

as in Lemma 9. We can invoke the theorem of the maximum, Lemma 15, to show that the best-reply correspondence $\Pi_0^* : \bar{\mathcal{A}} \Rightarrow \Pi_0(\mu_0)$ is non-empty, compact-valued, convex-valued, and upper hemi-continuous.

By Lemma 10, the function $M : \Pi_0(\mu_0) \rightarrow \bar{\mathcal{A}}$ that defines the mean-consistency condition ($\bar{\alpha} = M[\pi_0]$) is continuous.

It follows by Aliprantis and Border [2006] 17.23 that the composition of M and Π_0^* , $M \circ \Pi_0^* : \Pi_0(\mu_0) \Rightarrow \Pi_0(\mu_0)$, is a non-empty, upper hemi-continuous, and compact- and convex-valued. By the infinite-dimensional version of the Kakutani fixed point theorem (Aliprantis and Border [2006] 17.55), $M \circ \Pi_0^*$ has a fixed point, π_0^* . By construction, for $\bar{\alpha}^* = M[\pi_0^*]$, π_0^* is a best response to $\bar{\alpha}^*$ and $(\pi_0^*, \bar{\alpha}^*)$ satisfies mean-consistency, and thus $(\pi_0^*, \bar{\alpha}^*)$ is an equilibrium under the alternative formulation (Definition 4).

Define π^* as the measure induced from π_0^* by $(a, \mu'_0) \rightarrow (a, \phi[\mu'_0, \bar{\alpha}^*])$. It follows that $(\pi^*, \bar{\alpha}^*)$ is an equilibrium under Definition 3.

E.5 Proof of Proposition 2

By Lemma 14, the planner's problem is concave on the set of the mean-consistent pairs $(\pi_0, \bar{\alpha})$, and hence by Lemma 13 all equilibria are global maxima. If there were multiple equilibria with different values of $\bar{\alpha}$, then by Lemma 14, a convex combination would generate higher utility in the planner's problem, contradicting optimality. It follows that all equilibria share a common aggregate action function.

Note that the result in fact holds under weaker conditions. Define, for any $\bar{\alpha}' \in \bar{\mathcal{A}}$,

$$f(\bar{\alpha}) = \sup_{\pi_0 \in \Pi_0(\mu_0) : M[\pi_0] = \bar{\alpha}} E^{\pi_0} [g(a; s) - D(\phi[\mu'_0, \bar{\alpha}'] || \phi[\mu_0, \bar{\alpha}'])],$$

and observe that $f(\bar{\alpha})$ is concave by the convexity of $\{(\bar{\alpha} \in \bar{\mathcal{A}}, \pi_0 \in \Pi_0(\mu_0)) : M[\pi_0] = \bar{\alpha}\}$ and the linearity of the objective. If $f(\bar{\alpha}) + E^{\mu_0} [G(\bar{\alpha}(s, r); s)]$ is strictly concave, then

$$J_0^*(\bar{\alpha}) = \sup_{\pi_0 \in \Pi_0(\mu_0) : M[\pi_0] = \bar{\alpha}} E^{\pi_0} [g(a; s) - D(\phi[\mu'_0, \bar{\alpha}'] || \phi[\mu_0, \bar{\alpha}'])] + E^{\mu_0} [G(\bar{\alpha}(s, r); s)]$$

is strictly concave. By invariance in $\bar{\mathcal{A}}$, this problem is equivalent the planner's problem (setting $\bar{\alpha}' = \bar{\alpha}$), and therefore the solution to the planner's problem involves a unique $\bar{\alpha}$. Strict concavity in G is thus a sufficient but not necessary condition.

E.6 Additional Lemma

Lemma 16. *Let $(\pi_0^*, \bar{\alpha}^*) \in \Pi_0(\mu_0) \times \bar{\mathcal{A}}$ be any mean-consistent strategy profile. Suppose that for some $(s_0, r_0) \in S \times R$, there exists a π_0^* -positive-measure set $\Omega_0 \subset \text{supp}(\pi_0^*)$ such that for all $(a, \mu'_0) \in \Omega_0$, $a > \bar{\alpha}^*(s_0, r_0)$. Then there exists a $\delta > 0$ and a family of strategies $\{\pi_{0,\epsilon} \in \Pi_0(\mu_0)\}_{|\epsilon| < \delta}$*

such that each $\pi_{0,\epsilon}$ is induced from π_0^* by a mapping $m_\epsilon : A \times \mathcal{U}_0 \rightarrow A \times \mathcal{U}_0$ is differentiable in ϵ in the neighborhood of $\epsilon = 0$ and that satisfies, for some $x > 0$,

$$E^{\pi_{0,\epsilon}}[a|s, r] = \bar{\alpha}^*(s, r) + \epsilon x \mathbf{1}\{(s, r) = (s_0, r_0)\}.$$

Proof. See the technical appendix, [H.9](#). □

E.7 Proof of Proposition 3

This proof refers several of the additional lemmas found in Technical Appendix Section [E.1](#), and uses the perturbation lemma above (Lemma [16](#)). Note that the objective function J_0 is defined in Lemma [9](#) and the mean-consistency function M is defined in Lemma [10](#).

We first prove the “if” claim. If $(\pi_0^*, \bar{\alpha}^*)$ is a constrained efficient strategy profile, then by mean-consistency $\bar{\alpha}^* = M[\pi_0^*]$, and by optimality any other $\pi'_0 \in \Pi_0(\mu_0)$ must satisfy

$$\lim_{\epsilon \rightarrow 0^+} \frac{1}{\epsilon} (J_0((1 - \epsilon)\pi_0^* + \epsilon\pi'_0, M[(1 - \epsilon)\pi_0^* + \epsilon\pi'_0]) - J_0(\pi_0^*, M[\pi_0^*])) \leq 0.$$

It follows from Lemma [12](#) that

$$J_0(\pi'_0, M[\pi_0^*]) - J_0(\pi_0^*, M[\pi_0^*]) \leq 0,$$

which is the optimality condition of the agent’s problem given $\bar{\alpha}^* = M[\pi_0^*]$. Therefore, $(\pi_0^*, \bar{\alpha}^*)$ is an equilibrium. Note that this direction of the proof did not require that π_0^* have interior support.

We next prove the “only if” claim. Suppose that $(\pi_0^*, \bar{\alpha}^*)$ is a constrained efficient strategy profile and an equilibrium, and for which

$$\nabla_{\bar{\alpha}} C_0(\pi_0^*, \mu_0, \bar{\alpha}^*) \neq \mathbf{0}, \tag{27}$$

Let us apply Lemma [16](#) defined above, and suppose that for some (s_0, r_0) satisfying

$$\nabla_{\bar{\alpha}} C_0(\pi_0^*, \mu_0, \bar{\alpha}^*)(s_0, r_0) \neq 0,$$

there exists a π_0^* -positive-measure set $\Omega_0 \subset \text{supp}(\pi_0^*)$ such that for all $(a, \mu'_0) \in \Omega_0$, $a > \bar{\alpha}^*(s_0, r_0)$.

By the continuous differentiability of D and of the utility function, $J_0(\pi_{0,\epsilon}, \bar{\alpha}^*)$ is differentiable with respect to ϵ (by Lemma [16](#), $m_\epsilon(a, \mu'_0)$ is differentiable with respect to ϵ). Consequently, by optimality in the agent’s problem (which must hold in equilibrium),

$$\frac{\partial}{\partial \epsilon} J_0(\pi_{0,\epsilon}, \bar{\alpha}^*)|_{\epsilon=0} = 0.$$

Now consider $\frac{\partial}{\partial \epsilon} J(\pi_{0,\epsilon}, M[\pi_{0,\epsilon}])|_{\epsilon=0}$. By the continuous differentiability of D , this must equal

$$\frac{\partial}{\partial \epsilon} J_0(\pi_{0,\epsilon}, M[\pi_{0,\epsilon}])|_{\epsilon=0} = \left(\frac{\partial}{\partial \epsilon} M[\pi_{0,\epsilon}]|_{\epsilon=0} \right) \cdot \nabla_{\bar{\alpha}} C_0(\pi_0^*, \mu_0, \bar{\alpha}^*)$$

and by construction this is non-zero, contradicting optimality in the planner's problem.

It must therefore be the case that, for $(s_0, r_0) \in S \times R$ with $\nabla_{\bar{\alpha}} C_0(\pi_0^*, \mu_0, \bar{\alpha}^*)(s_0, r_0) \neq 0$, no set Ω_0 as described in Lemma 16 exists. But in this case, mean-consistency requires that the action $a = \bar{\alpha}(s_0, r_0)$ occur with probability one under π_0^* conditional on (s_0, r_0) .

If this action occurs with probability one unconditionally under π_0^* , then by the definition of a divergence we must have $C_0(\pi_0^*, \mu_0, \bar{\alpha}) = 0$ for all $\bar{\alpha}$, contradicting (27). Instead suppose some actions $a \in A \setminus \{\bar{\alpha}(s_0, r_0)\}$ occur with positive probability. In this case, the posteriors associated with such actions must place zero probability on (s_0, r_0) , contradicting the assumption that the posteriors lie in the interior of the simplex.

We therefore conclude that in all constrained-efficient equilibria satisfying the interiority assumptions, $\nabla_{\bar{\alpha}} C_0(\pi_0^*, \mu_0, \bar{\alpha}^*) = \mathbf{0}$.

E.8 Proof of Proposition 4

We first prove that $\bar{\alpha}$ lies in the relative interior of \bar{A} , which is equivalent to showing that $\min_{a \in A} a = a_{min} < \bar{\alpha}(s, r) < a_{max} = \max_{a \in A} a$ for all $(s, r) \in S \times R$.

Suppose not, and that (without loss of generality), for some $(s_0, r_0) \in S \times R$,

$$E^{\pi_0}[a|s_0, r_0] = a_{max}.$$

This implies that

$$\int_{A \times \mathcal{U}_0} (a_{max} - a) \mu'_0(s_0, r_0) d\pi_0(a, \mu'_0) = 0.$$

Therefore, for any $\epsilon > 0$, the set of (a, μ'_0) with $(a_{max} - a) \mu'_0(s_0, r_0) > \epsilon$ has measure zero under π_0 .

By the assumption that $\bar{\alpha}$ is not constant, there exists some $(s, r) \in S \times R$ with $\bar{\alpha}(s, r) \neq a_{max}$. It follows that

$$\int_{A \times \mathcal{U}_0} (a_{max} - a) \mu'_0(s, r) d\pi_0(a, \mu'_0) = \mu_0(s, r)(a_{max} - \bar{\alpha}(s, r)).$$

Therefore, the set of (a, μ'_0) with $(a_{max} - a) \mu'_0(s, r) \geq \mu_0(s, r)(a_{max} - \bar{\alpha}(s, r)) > 0$ has positive measure under π_0 .

If a π_0 -positive measure subset of this set exists with $\frac{\mu'_0(s_0, r_0)}{\mu'_0(s, r)} > \delta$ for any $\delta > 0$, a contradiction arises. But if a π_0 -positive measure subset exists with $\frac{\mu'_0(s_0, r_0)}{\mu'_0(s, r)} = 0$, this violates the interior support assumption. We conclude that $E^{\pi_0}[a|s_0, r_0] < a_{max}$, and by a symmetric argument $E^{\pi_0}[a|s_0, r_0] > a_{min}$. Note also that, by the assumption that $\bar{\alpha}$ is not constant, mean-consistency implies that π_0 cannot be a zero-information strategy. It follows from the definition of nowhere-

\bar{A} -invariance that local invariance cannot hold, and therefore by Proposition 3 $(\pi_0, \bar{\alpha})$ cannot be a constrained-efficient equilibrium.

E.9 Proof of Proposition 5

The first part of this proposition follows from the proof of Proposition 3, noting that the “if” direction of that proof did not require an interiority assumption and that invariance implies local invariance.

By Lemma 14, if G is concave in \bar{a} for all $s \in S$, the planner’s problem is concave on the set of the mean-consistent pairs $(\pi_0, \bar{\alpha})$, and hence by Lemma 13 all equilibria are global maxima.

E.10 Additional Lemmas

Lemma 17. *Let $\bar{\alpha} \in \bar{A}$ be an s -measurable aggregate action function whose image is a non-singleton subset of the relative interior of \bar{A} , and let $\mu'' \in \mathcal{U}$ be a measure satisfying $\mu'' \ll \phi[\mu_0, \bar{\alpha}]$ and $\mu'' = \eta_{-R|R}[\mu''|\phi[\mu_0, \bar{\alpha}]]$. Then there exists a $\rho \in (0, 1)$, $\bar{\epsilon} > 0$, and $a_0, a_1 \in A$ such that, for all $\epsilon \in [0, \bar{\epsilon})$, there exists an s -measurable $\pi_{\epsilon, \rho} \in \Pi(\phi[\mu_0, \bar{\alpha}])$ with finite support such that $(\bar{\alpha}, \pi_{\epsilon, \rho})$ satisfies mean-consistency, $\pi_{\epsilon, \rho}(a_1, \phi[\mu_0'', \bar{\alpha}]) = \epsilon\rho$, and $\pi_{\epsilon, \rho}(a_0, \phi[\mu_0, \bar{\alpha}]) = \epsilon(1 - \rho)$. Moreover, for each $a \in A$, $(a, \mu') \in \text{supp}(\pi_{\epsilon, \rho})$ for at most one $\mu' \in \mathcal{U}$.*

Proof. See the technical appendix, H.5. □

Lemma 18. *Suppose the divergence D is not monotone in R on the set of s -measurable priors. Then there exists an s -measurable, mean-consistent $(\pi \in \Pi(\phi[\mu_0, \bar{\alpha}]), \bar{\alpha} \in \bar{A})$ such that π has finite support, for each $a \in A$, $(a, \mu') \in \text{supp}(\pi)$ for at most one $\mu' \in \mathcal{U}$, and a $\hat{\pi} \in \Pi(\phi[\mu_0, \bar{\alpha}])$ such that π is induced from $\hat{\pi}$ by $(a, \mu') \mapsto (a, \eta_{-R|R}[\mu'|\mu])$ exists with*

$$C(\hat{\pi}, \phi[\mu_0, \bar{\alpha}]) < C(\pi, \phi[\mu_0, \bar{\alpha}]).$$

Proof. See technical appendix, H.6. □

Lemma 19. *Let D be a divergence monotone in R , and let $(\bar{\pi}, \bar{\alpha})$ be any s -measurable, mean-consistent strategy profile such that $\bar{\pi}$ has finite support and, for each $a \in A$, $(a, \mu') \in \text{supp}(\bar{\pi})$ for at most one $\mu' \in \mathcal{U}$. There exists, for any $\chi > 0$ and $r_0 \in R$, a mean-critical utility function*

$$v(a, \bar{a}, s) = g(a; s) - \frac{\chi}{2} |\bar{a} - \bar{\alpha}(s, r_0)|^2 - \chi(a - \bar{a}) \cdot (\bar{a} - \bar{\alpha}(s, r_0))$$

such that $(\pi, \bar{\alpha})$ is an equilibrium.

Proof. See the technical appendix, H.7. □

E.11 Proof of Proposition 6

Suppose there exists an s-measurable BNE $(\pi, \bar{\alpha})$ in which π has non-zero information, and that C is nowhere-R-monotone for all $\mu \in \mathcal{U}$ such that $\gamma_{-\bar{A}}[\mu] = \mu_0$. Then by definition there must exist a $\hat{\pi} \in \Pi(\mu)$ such that $C(\hat{\pi}, \mu) < C(\pi, \mu)$ and such that π is induced from $\hat{\pi}$ by $(a, \mu') \mapsto (a, \eta_{-R|R}[\mu'|\mu])$.

By the s-measurability of $\bar{\alpha}$, for all $\mu' \in \mathcal{U}$ with $\mu' \ll \mu$,

$$V(a, \mu') = V(a, \eta_{-R|R}[\mu'|\mu]),$$

and consequently

$$E^{\hat{\pi}}[V(a, \mu')] = E^{\pi}[V(a, \mu')].$$

By $C(\hat{\pi}, \mu) < C(\pi, \mu)$, $(\pi, \bar{\alpha})$ cannot be an equilibrium.

E.12 Proof of Proposition 7

We first prove the “if” part of the result. This essentially reproduces our existence proof (Proposition 1) on the space of s-measurable aggregate action functions and strategies.

It is convenient to define a strategy $\pi_0 \in \Pi_0(\mu_0)$ as s-measurable if, for all $(a, \mu'_0) \in \text{supp}(\pi_0)$, there exists a function (Radon-Nikodym derivative) $f_{\mu'_0} : S \rightarrow \mathbb{R}$ such that $\frac{d\mu'_0(s,r)}{d\mu_0(s,r)} = f_{\mu'_0}(s)$ for all $s \in S$ and $r \in R$. For s-measurable aggregate action functions $\bar{\alpha}$, this definition is equivalent to the requirement that the strategy $\pi \in \Pi(\phi[\mu_0, \bar{\alpha}])$ induced from π_0 by $(a, \mu'_0) \mapsto (a, \phi[\mu'_0, \bar{\alpha}])$ be non-R-measurable.

Define $\Pi_0^S(\mu_0) \subseteq \Pi_0(\mu_0)$ as the set of s-measurable strategies.

Lemma 20. *The set $\Pi_0^S(\mu_0)$ is non-empty, compact, and convex.*

Proof. See technical appendix, H.8. □

Using this result and the continuity of the expected utility function $J_0(\pi_0, \bar{\alpha})$ defined in Lemma 9, we can invoke the theorem of the maximum. Define $\bar{A}^S \subset \bar{A}$ as the set of s-measurable aggregate action functions. Let $\Pi^* : \bar{A}^S \Rightarrow \Pi^S(\mu_0)$ be the best reply correspondence in the agent’s problem restricted to s-measurable strategies,

$$\Pi_S^*(\bar{\alpha}) = \{\pi \in \arg \max_{\pi' \in \Pi_0^S(\mu_0)} J(\pi'_0, \bar{\alpha})\}.$$

By the maximum theorem (see Aliprantis and Border [2006] 17.31), the linearity of J in π_0 , and the convexity of $\Pi_0^S(\mu_0)$, Π_S^* is non-empty, compact-valued, convex-valued, and is upper hemicontinuous. By Lemma 10, the function $M : \Pi_0(\mu_0) \rightarrow \bar{A}$ defined by $M[\pi_0](s, r) = E^{\pi_0}[a|s, r]$ for

all $(s, r) \in S \times R$ is continuous. Moreover, if $\pi_0 \in \Pi_0^S(\mu_0)$, then for all $s \in S$ and $r, r' \in R$,

$$\begin{aligned} E^{\pi_0}[a|s, r] &= \int_{\text{supp}(\pi_0)} a f_{\mu'_0}(s) d\pi_0(a, \mu'_0) \\ &= E^{\pi_0}[a|s, r'] \end{aligned}$$

and consequently the s-measurability of π_0 guarantees the s-measurability of $M(\pi_0)$.

It follows by [Aliprantis and Border \[2006\]](#) 17.23 that the composition of M and Π_S^* , $M \circ \Pi_S^* : \Pi_0^S(\mu_0) \Rightarrow \Pi_0^S(\mu_0)$, is a non-empty, upper hemi-continuous, and compact- and convex-valued. By the infinite-dimensional version of the Kakutani fixed point theorem ([Aliprantis and Border \[2006\]](#) 17.55), $M \circ \Pi^*$ has a fixed point, π_0^* . By construction, for $\bar{\alpha}^* = M[\pi_0^*]$, π_0^* is a best response to $\bar{\alpha}^*$ on the set of s-measurable strategies, $(\pi_0^*, \bar{\alpha}^*)$ is mean-consistent, and both $\bar{\alpha}^*$ and π_0^* are s-measurable. Define π^* as the measure induced from π_0^* by $(a, \mu'_0) \rightarrow (a, \phi[\mu'_0, \bar{\alpha}^*])$ and observe by construction that $\pi^* \in \Pi(\phi[\mu_0, \bar{\alpha}^*])$ is a best response to $\bar{\alpha}^*$ on the set of s-measurable strategies.

By the s-measurability of $\bar{\alpha}^*$, for all $\mu' \in \mathcal{U}$ with $\mu' \ll \phi[\mu_0, \bar{\alpha}^*]$,

$$V(a, \mu') = V(a, \eta_{-R|R}[\mu'|\mu]).$$

Consequently, for any $\pi \in \Pi(\phi[\mu_0, \bar{\alpha}^*])$, the s-measurable $\hat{\pi} \in \Pi(\phi[\mu_0, \bar{\alpha}^*])$ induced from π by $(a, \mu') \mapsto (a, \eta_{-R|R}[\mu'|\phi[\mu_0, \bar{\alpha}^*]])$ satisfies

$$E^{\hat{\pi}}[V(a, \mu')] = E^{\pi}[V(a, \mu')],$$

and by monotonicity in R for the prior $\phi[\mu_0, \bar{\alpha}^*]$, $E^{\hat{\pi}}[D(\mu'|\phi[\mu_0, \bar{\alpha}^*])] \leq E^{\pi}[D(\mu'|\phi[\mu_0, \bar{\alpha}^*])]$. It follows that every non-s-measurable π is weakly dominated by an s-measurable one, and consequently π^* is a best response without the restriction to s-measurable strategies, and therefore $(\pi^*, \bar{\alpha}^*)$ is an equilibrium.

We next prove the “only if” part of the claim. Let $\mu = \phi\{\mu_0, \bar{\alpha}\}$, and suppose D is not monotone in R for all s-measurable $\bar{\alpha}$. Then by [Lemma 18](#), there exists an s-measurable $\bar{\alpha}$ and strategies π and $\hat{\pi}$ as described in the lemma, with

$$C(\hat{\pi}, \mu) < C(\pi, \mu).$$

Note that π is non-R-measurable, and $\hat{\pi}$ is not (π is induced from $\hat{\pi}$ by $(a, \mu') \mapsto (a, \eta_{-R|R}[\mu'|\mu])$ and $\hat{\pi} \neq \pi$), and that $(\pi, \bar{\alpha})$ are mean-consistent and satisfy certain technical properties (which are relevant for [Lemma 19](#)).

First, let us observe by continuity ([Assumption 1](#)) that if this inequality holds, it holds for some $\bar{\alpha} \in \bar{A}$ whose image is contained in the relative interior of \bar{A} . Let us define, for all $\mu' \in \mathcal{U}$ with $\mu' \ll \mu$,

$$D_R(\mu'|\mu) = D(\eta_{-R|R}[\mu'|\mu]|\mu)$$

and observe by definition that D_R is an R-monotone divergence, and that by the s-measurability of π ,

$$E^\pi[D(\mu' || \mu)] = E^\pi[D_R(\mu' || \mu)].$$

By Lemma 19, for the divergence D_R , there exists for any $\chi > 0$ a utility function v of the sort described in that lemma, with

$$v(a, \bar{a}, s) = g(a; s) - \frac{\chi}{2}(\bar{a} - \bar{\alpha}(s, r_0))^2 - \chi(a - \bar{a}) \cdot (\bar{a} - \bar{\alpha}(s, r_0)),$$

such that $(\pi, \bar{\alpha})$ is an equilibrium given D_R . By the argument given above, $E^{\hat{\pi}}[V(a, \mu')] = E^\pi[V(a, \mu')]$ and consequently $(\pi, \bar{\alpha})$ cannot be an equilibrium given D .

Moreover, because π is a best response under D_R to $\bar{\alpha}$, any other non-R-measurable $\pi' \in \Pi(\mu)$ must satisfy

$$\begin{aligned} E^{\pi'}[V(a, \mu') - D(\mu' || \mu)] &= E^{\pi'}[V(a, \mu') - D_R(\mu' || \mu)] \\ &\leq E^\pi[V(a, \mu') - D_R(\mu' || \mu)]. \end{aligned}$$

Consequently,

$$E^{\pi'}[V(a, \mu') - D(\mu' || \mu)] < E^{\hat{\pi}}[V(a, \mu') - D(\mu' || \mu)],$$

and therefore no non-R-measurable $\pi' \in \Pi(\mu)$ is a best reply to $\bar{\alpha}$, and consequently no s-measurable equilibrium involving $\bar{\alpha}$ exists.

Suppose there exists some other s-measurable equilibrium $(\tilde{\pi}, \tilde{\alpha})$. We claim that, for sufficiently large values of χ , any equilibrium must lie in the neighborhood of $\bar{\alpha}$. To see this, suppose that for some $\epsilon > 0$ and $s \in S$,

$$|\tilde{\alpha}(s, r_0) - \bar{\alpha}(s, r_0)| > \epsilon.$$

Note that we can rewrite the utility function as

$$v(a, \tilde{\alpha}(s, r_0), s) = g(a; s) + \frac{\chi}{2}(\tilde{\alpha}(s, r_0) - \bar{\alpha}(s, r_0))^2 - \chi(a - \bar{\alpha}(s, r_0)) \cdot (\tilde{\alpha}(s, r_0) - \bar{\alpha}(s, r_0)).$$

Thus, by mean-consistency,

$$E^{\tilde{\pi}}[v(a, \tilde{\alpha}(s, r_0), s) | s, r] < E^{\tilde{\pi}}[g(a; s) | s, r] - \frac{\chi}{2}\epsilon^2.$$

If we consider instead the policy $\pi^+ \in \Pi(\phi[\mu_0, \tilde{\alpha}])$ induced from π by $(a, \mu') \rightarrow (a, \eta_{-\bar{A}|\bar{A}}[\mu' | \phi[\mu_0, \tilde{\alpha}]])$, we have by construction $E^{\pi^+}[a | s, r] = E^\pi[a | s, r] = \bar{\alpha}(s, r)$ and therefore

$$\begin{aligned} E^{\pi^+}[v(a, \tilde{\alpha}(s, r_0), s) | s, r] &= E^\pi[g(a; s) | s, r] + \frac{\chi}{2}(\tilde{\alpha}(s, r_0) - \bar{\alpha}(s, r_0))^2 \\ &> E^\pi[g(a; s) | s, r] + \frac{\chi}{2}\epsilon^2. \end{aligned}$$

For sufficiently large χ , it follows that π^+ achieves a higher expected utility than $\tilde{\pi}$ in response

to $\tilde{\alpha}$ (as $C(\pi^+, \phi[\mu_0, \tilde{\alpha}])$ is by Assumption 1 finite), and consequently $(\tilde{\pi}, \tilde{\alpha})$ cannot be an equilibrium.

Now consider a sequence χ_n with $\lim_{n \rightarrow \infty} \chi_n = \infty$, and suppose that along this sequence of utility functions there exists an s -measurable equilibrium $(\pi_{0,n}, \bar{\alpha}_n) \in \Pi_0(\mu_0) \times \bar{\mathcal{A}}$. (Note that if no such sequence exists, the proof is complete). By the argument above, $\lim_{n \rightarrow \infty} \bar{\alpha}_n = \bar{\alpha}$. By the theorem of the maximum (Lemma 15), the optimal policy correspondence for the agents is upper hemi-continuous, and consequently $\lim_{n \rightarrow \infty} \pi_{0,n}$ must be a mean-consistent, non-R-measurable best response to $\bar{\alpha}$. But we have already shown that no such policy exists, and consequently for sufficient large χ no s -measurable equilibrium exists.

E.13 Proof of Corollary 1

Take an s -measurable $\bar{\alpha}$ such that D is not monotone in R given $\mu = \phi[\mu_0, \bar{\alpha}]$ as given. The utility function constructed in the necessity part of proof of Proposition 6 is, for some $r_0 \in R$ and $\chi > 0$,

$$v(a, \bar{a}, s) = g(a; s) - \frac{\chi}{2}(\bar{a} - \bar{\alpha}(s, r_0))^2 - \chi(a - \bar{a}) \cdot (\bar{a} - \bar{\alpha}(s, r_0)).$$

Defining $G(\bar{a}; s) = -\frac{\chi}{2}(\bar{a} - \bar{\alpha}(s, r_0))^2$ yields

$$v(a, \bar{a}, s) = g(a; s) + G(\bar{a}; s) + (a - \bar{a}) \cdot \frac{\partial}{\partial \bar{a}} G(\bar{a}, s),$$

and hence v is mean-critical. By Proposition 1, an equilibrium exists, and by the proof of Proposition 6, no s -measurable equilibrium exists.

By the strict concavity of G and Propositions 2 and 5, all equilibria share a common aggregate action function and are constrained efficient.

F Proofs of Appendix Results

F.1 Proof of Lemma 3

The “if” part is essentially immediate:

$$\nabla_{\bar{a}} v(a, \bar{a}, s) = (a - \bar{a}) \cdot H(\bar{a}; s),$$

where $H(\cdot)$ is the Hessian of G . The mean-critical property follows by taking expectations.

We next prove the “only if” part of the lemma. Consider a point-mass σ_a with full support on some $a \in A$ given each $s \in S$. By the definition of a mean-critical utility function, we must have

$$\nabla_{\bar{a}} v(a, \bar{a}, s)|_{\bar{a}=a} = 0.$$

In case of concavity in \bar{a} , it follows that

$$a \in \arg \max_{\bar{a} \in \bar{A}} v(a, \bar{a}, s).$$

Define the function

$$f(a, \bar{a}, s) = v(a, a, s) - v(a, \bar{a}, s).$$

It follows immediately that

$$a \in \arg \min_{\bar{a} \in \bar{A}} f(a, \bar{a}, s),$$

and $f(a, a, s) = 0$ by construction.

In the case of convexity in \bar{a} ,

$$a \in \arg \min_{\bar{a} \in \bar{A}} v(a, \bar{a}, s),$$

and

$$f(a, \bar{a}, s) = v(a, \bar{a}, s) - v(a, a, s)$$

is a non-negative function with $f(a, a, s) = 0$.

Consequently, on the domain $A \times A \times S$, $f(a, \bar{a}, s)$ is a weakly positive, and is continuously twice-differentiable by the assumption of regularity.

Now consider a measure σ consisting of two point masses, $\sigma(s) = \alpha\delta_{a_1} + (1 - \alpha)\delta_{a_2}$ for each $s \in S$. By mean-consistency,

$$\int_A a d\sigma(a) \in \arg \min_{\bar{a} \in \bar{A}} \int_A f(a, \bar{a}; s) d\sigma(a).$$

By theorem 4 of [Banerjee et al. \[2005\]](#) (which is proven using only two-point measures; see also the discussion in that paper on restrictions to subspaces of \mathbb{R}^N), it follows that

$$f(a, \bar{a}, s) = F(a, s) - F(\bar{a}, s) - (a - \bar{a}) \cdot \nabla F(\bar{a}, s)$$

for some function $F : A \times S \rightarrow \mathbb{R}$ convex in A for all $s \in S$. In the concave in \bar{a} case, defining

$$g(a, s) = v(a, a, s) + F(a, s)$$

and $G(\bar{a}, s) = F(\bar{a}, s)$ yields

$$\begin{aligned} v(a, \bar{a}, s) &= v(a, a, s) - f(a, \bar{a}, s) \\ &= g(a, s) + G(\bar{a}, s) + (a - \bar{a}) \cdot \nabla G(\bar{a}, s). \end{aligned}$$

In the convex in \bar{a} case, defining $g(a, s) = v(a, a, s) - G(a, s)$ and $G(\bar{a}, s) = -F(\bar{a}, s)$ yields

$$\begin{aligned} v(a, \bar{a}, s) &= v(a, a, s) + f(a, \bar{a}, s) \\ &= g(a, s) + G(\bar{a}, s) - (a - \bar{a}) \cdot \nabla G(\bar{a}, s). \end{aligned}$$

E.2 Proof of Proposition 8

This proof refers several of the additional lemmas found in Technical Appendix Section E.1, and uses the perturbation lemma above (Lemma 16).

The planner's problem can be rewritten as $\max_{\pi_0 \in \Pi_0(\mu_0)} J_0(\pi_0, M[\pi_0])$, where J_0 is defined as in Lemma 9 and $M[\pi_0]$ denotes the aggregate action function implied by π_0 , as defined in Lemma 10. The agent's problem is $\max_{\pi_0 \in \Pi_0(\mu_0)} J_0(\pi_0, \bar{\alpha})$, taking $\bar{\alpha} \in \bar{A}$ as given.

Suppose the proposition is false; that is, suppose there exists a differentiable divergence D , invariant in \bar{A} , such that $(\pi_0^*, \bar{\alpha}^*)$ is a constrained efficient BNE with $\text{supp}(\pi_0) \subset A \times \text{int}(\mathcal{U}_0)$ and, for some $(s_0, r_0) \in S \times R$,

$$\frac{\partial}{\partial \bar{a}} E^{\pi_0^*} [v(a, \bar{a}, s) | s_0, r_0] |_{\bar{a}=\bar{\alpha}(s_0, r_0)} \neq 0.$$

Let us apply Lemma 16, and suppose that there exists a π_0^* -positive-measure set $\Omega_0 \subset \text{supp}(\pi_0^*)$ such that for all $(a, \mu'_0) \in \Omega_0$, $a > \bar{\alpha}^*(s_0, r_0) = E^{\pi_0^*} [a | s_0, r_0]$.

By the differentiability of D and of the utility function, $J_0(\pi_0, \epsilon, \bar{\alpha})$ is differentiable with respect to ϵ (by Lemma 16, $m_\epsilon(a, \mu'_0)$ is differentiable with respect to ϵ). Consequently, by optimality in the agent's problem (which must hold in equilibrium),

$$\frac{\partial}{\partial \epsilon} J_0(\pi_0, \epsilon, \bar{\alpha}) |_{\epsilon=0} = 0.$$

Now consider $\frac{\partial}{\partial \epsilon} J(\pi_0, \epsilon, M[\pi_0, \epsilon]) |_{\epsilon=0}$. By the invariance of D in \bar{A} , this must equal

$$\begin{aligned} \frac{\partial}{\partial \epsilon} J(\pi_0, \epsilon, M[\pi_0, \epsilon]) |_{\epsilon=0} &= E^{\mu_0} \left[\left(\frac{\partial}{\partial \epsilon} M[\pi_0, \epsilon](s, r) \right) |_{\epsilon=0} \cdot \frac{\partial}{\partial \bar{a}} E^{\pi_0^*} [v(a, \bar{a}, s) | s, r] |_{\bar{a}=\bar{\alpha}(s, r)} \right] \\ &= \mu_0(s_0, r_0) x \frac{\partial}{\partial \bar{a}} E^{\pi_0^*} [v(a, \bar{a}, s) | s_0, r_0] |_{\bar{a}=\bar{\alpha}(s_0, r_0)}, \end{aligned}$$

where $x > 0$ is the constant defined in Lemma 16. By construction this is non-zero, contradicting optimality in the planner's problem.

It must therefore be the case that, for $(s_0, r_0) \in S \times R$ with $\frac{\partial}{\partial \bar{a}} E^{\pi_0^*} [v(a, \bar{a}, s) | s_0, r_0] |_{\bar{a}=\bar{\alpha}(s_0, r_0)} \neq 0$, no set Ω_0 as described in Lemma 16 exists. But in this case, mean-consistency requires that the action $a = \bar{\alpha}(s_0, r_0)$ occur with probability one under π_0^* conditional on (s_0, r_0) .

If this action occurs with probability one unconditionally under π_0^* , then we must have

$$\frac{\partial}{\partial \bar{a}} E^{\mu_0} [v(\bar{\alpha}(s_0, r_0), \bar{a}, s) | \bar{a}=\bar{\alpha}(s_0, r_0)] \neq 0,$$

but this contradicts optimality in the planner's problem given optimality in the agent's problem, as above.

Instead suppose some actions $a \in A \setminus \{\bar{\alpha}(s_0, r_0)\}$ occur with positive probability. In this case, the posteriors associated with such actions must place zero probability on (s_0, r_0) , contradicting the assumption that the posteriors lie in the interior of the simplex.

We therefore conclude that in all constrained-efficient BNE satisfying the interiority assump-

tions, $\frac{\partial}{\partial \bar{a}} E^{\pi_0^*} [v(a, \bar{a}, s) | s_0, r_0] |_{\bar{a}=\bar{a}(s_0, r_0)} = 0$.

E.3 Proof of Proposition 10

To ease notation we define $\vec{\alpha} = (\bar{\alpha}_s, \bar{\alpha}_r)$. A strategy profile $(\lambda_0, \tau_0, \vec{\alpha})$ is an equilibrium if (λ_0, τ_0) is a best response and mean-consistency is satisfied. Recall that in this section we use the normalization $\lambda_0' \Sigma_0 \lambda_0 = 1$. Under this normalization,

$$\alpha(\lambda_0, \tau_0, \vec{\alpha}) = (\psi(\vec{\alpha})' \Sigma_0 \lambda_0) \frac{\tau_0}{1 + \tau_0}.$$

Mean-consistency is

$$\vec{\alpha} = \alpha(\lambda_0, \tau_0, \vec{\alpha}) \lambda_0.$$

We also have

$$\psi(\vec{\alpha}) = (1 - \beta)e_1 + \beta \alpha(\lambda_0, \tau_0, \vec{\alpha}) \lambda_0,$$

where $e_1' = (1, 0)$.

Mutual Information As argued in the text, with $\vec{\alpha}$ s -measurable, $\vec{\alpha} = (\bar{\alpha}_s, 0)$, $\lambda_0' = (\sigma_s^{-1}, 0)$ is the best response. In this case we have

$$\psi(\vec{\alpha}) \Sigma_0 \lambda_0 = (1 - \beta + \beta \bar{\alpha}_s) \sigma_s$$

and mean-consistency requires

$$\bar{\alpha}_s = (1 - \beta + \beta \bar{\alpha}_s) \frac{\tau_0}{1 + \tau_0}.$$

Optimality in this case requires

$$\tau_0^* \in \arg \max_{\tau_0} (1 - \beta + \beta \bar{\alpha}_s)^2 \sigma_s^2 \left(1 - \frac{1}{1 + \tau_0}\right) + \theta \ln\left(\frac{1}{1 + \tau_0}\right),$$

The first-order condition for τ_0^* is

$$(1 - \beta + \beta \bar{\alpha}_s)^2 \sigma_s^2 \leq \theta(1 + \tau_0^*),$$

with equality if $\tau_0^* > 0$. If $\theta \geq (1 - \beta)^2 \sigma_s^2$, an equilibrium with $(\bar{\alpha}_s = 0, \tau_0 = 0)$ exists.

Suppose $\theta < (1 - \beta)^2 \sigma_s^2$. Define the function $g : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ by

$$g(\tau_0, \bar{\alpha}_s) = \begin{bmatrix} \theta^{-1} (1 - \beta + \beta \bar{\alpha}_s)^2 \sigma_s^2 - 1 \\ (1 - \beta + \beta \bar{\alpha}_s) \frac{\tau_0}{1 + \tau_0} \end{bmatrix}.$$

Let $K \subset \mathbb{R}^2$ be the interval $[0, \theta^{-1} \sigma_s^2] \times [0, 1]$, a convex, compact set. Observe that for any $(\tau_0, \bar{\alpha}_s) \in K$, $g(\tau_0, \bar{\alpha}_s) \in K$. The function g is continuous, and consequently by Brouwer's fixed point

theorem, a fixed point $(\tau_0^*, \bar{\alpha}_s^*)$ exists. By construction, $(\lambda_0^* = (\sigma_s^{-1}, 0)', \bar{\alpha}^* = (\bar{\alpha}_s^*, 0), \tau_0^*)$ is an s-measurable equilibrium.

Now observe that in any equilibrium $(\lambda_0^*, \tau_0^*, \bar{\alpha}^*)$, optimality requires (as argued in the text)

$$\frac{\tau_0^*}{1 + \tau_0^*} (\psi(\bar{\alpha}^*)' \Sigma_0 \lambda_0^*) \psi(\bar{\alpha}^*)' = \nu^* \lambda_0^*,$$

for some multiplier $\nu^* \geq 0$. Under our normalization, this implies

$$\lambda_0^* = \frac{\psi(\bar{\alpha}^*)}{(\psi(\bar{\alpha}^*)' \Sigma_0 \psi(\bar{\alpha}^*))^{\frac{1}{2}}}.$$

Mean-consistency requires

$$\psi(\bar{\alpha}) = (1 - \beta)e_1 + \beta (\psi(\bar{\alpha}^*)' \Sigma_0 \lambda_0^*) \frac{\tau_0^*}{1 + \tau_0^*} \lambda_0^*,$$

and combining these equations yields

$$\psi(\bar{\alpha}^*) = (1 - \beta)e_1 + \beta \frac{\tau_0^*}{1 + \tau_0^*} \psi(\bar{\alpha}^*).$$

It follows by $\beta < 1$ that any equilibrium must feature $\psi(\bar{\alpha}^*) \propto e_1$, or equivalently $\bar{\alpha}_r^* = 0$.

Fisher Information For the Fisher information, we build on results in [Hébert and Woodford \[2021\]](#). Those authors adopt the normalization $|\lambda_0| = \theta^{-\frac{1}{2}}$ (note that a factor of two is omitted due to a difference in the definition of θ). They show that if $\theta < |\Sigma_0 \psi(\bar{\alpha})|^2$, agents will respond by gathering information, and otherwise agents will not gather information. Let us call $\hat{\lambda}_0, \hat{\tau}_0$ the optimal policies under the [Hébert and Woodford \[2021\]](#) normalization.

Consequently, if $\theta \geq ((1 - \beta)\sigma_s^2)^2$, an equilibrium with no information gathering exists. Suppose in what follows that $\theta < ((1 - \beta)\sigma_s^2)^2$. If agents acquire information, [Hébert and Woodford \[2021\]](#) show that

$$\begin{aligned} \hat{\lambda}_0 &= (\theta \Sigma_0^{-1} + \hat{\tau}_0 I)^{-1} \psi(\bar{\alpha}), \\ \theta &= |(\Sigma_0^{-1} + \theta^{-1} \hat{\tau}_0 I)^{-1} \psi(\bar{\alpha})|^2. \end{aligned}$$

Observe that

$$\psi(\bar{\alpha})' \Sigma_0 \hat{\lambda}_0 = \theta \hat{\lambda}_0' \hat{\lambda}_0 + \hat{\tau}_0 \hat{\lambda}_0' \Sigma_0 \hat{\lambda}_0 > 0,$$

which leads to the result in [Hébert and Woodford \[2021\]](#) that

$$E[(s, r) \cdot \psi(\bar{\alpha}) | \omega] = \hat{\tau}_0 \omega.$$

Mean-consistency therefore requires

$$\psi(\vec{\alpha}) = (1 - \beta)e_1 + \beta\hat{\tau}_0\hat{\lambda}_0.$$

Combing these equations,

$$(1 - \beta)\hat{\tau}_0\hat{\lambda}_0 = (1 - \beta)e_1 - \theta\Sigma_0^{-1}\hat{\lambda}_0,$$

or

$$\hat{\lambda}_0 = \theta^{-1}(1 - \beta)\Sigma_0(e_1 - \hat{\tau}_0\hat{\lambda}_0). \quad (28)$$

This can be rewritten as

$$(\theta\tilde{\Sigma}_0^{-1} + \hat{\tau}_0I)\hat{\lambda}_0 = e_1,$$

where $(\tilde{\Sigma}_0)^{-1} = (1 - \beta)^{-1}\Sigma_0^{-1}$. That is, in equilibrium agents solve the [Hébert and Woodford \[2021\]](#) without strategic considerations (i.e. the $\beta = 0$ case), but acting as if the variance-covariance matrix is scaled by $(1 - \beta)$. It follows immediately that an equilibrium with information acquisition exists if

$$\theta < |(1 - \beta)\Sigma_0 e_1|^2 = ((1 - \beta)\sigma_s^2)^2.$$

Observe, by the argument in the text or (28) that $\hat{\lambda}'_0 \neq (\theta^{-\frac{1}{2}}, 0)$ due to the off-diagonal elements of Σ_0 . Consequently, the equilibrium cannot be s-measurable.

E4 Proof of 11

Let us recall the first-order condition and mean-consistency condition:

$$(1 - \beta + \beta\bar{\alpha}_s)^2 = (1 - \beta)^2(1 + \tau_0\sigma_s^2),$$

$$\bar{\alpha}_s = (1 - \beta + \beta\bar{\alpha}_s)\left(1 - \frac{1}{1 + \tau_0\sigma_s^2}\right).$$

We must have $\tau_0\sigma_s^2 \geq 0$, and if this quantity is strictly positive, $\text{sgn}(\bar{\alpha}_s) = \text{sgn}(1 - \beta + \beta\bar{\alpha}_s)$. We can rewrite the mean-consistency condition as

$$(1 + \tau_0\sigma_s^2)(1 - \beta)(1 - \bar{\alpha}_s) = (1 - \beta + \beta\bar{\alpha}_s).$$

It follows that $\beta \in (0, 1)$ requires $\bar{\alpha}_s \geq 0$ in equilibrium.

We can substitute out $1 + \tau_0\sigma_s^2$ to find

$$\frac{(1 - \beta + \beta\bar{\alpha}_s)^2}{(1 - \beta)}(1 - \bar{\alpha}_s) = (1 - \beta + \beta\bar{\alpha}_s)$$

and therefore

$$(1 - \beta + \beta\bar{\alpha}_s)(1 - \bar{\alpha}_s) = 1 - \beta,$$

which is

$$(2\beta - 1)\bar{\alpha}_s - \beta\bar{\alpha}_s^2 = 0.$$

Therefore $\bar{\alpha}_s = 0$ and $\tau_0 = 0$ is an equilibrium. There can also be an equilibrium if

$$\bar{\alpha}_s = \frac{2\beta - 1}{\beta},$$

with the latter being possible and non-zero only when $\beta \in (\frac{1}{2}, 1)$. In this case,

$$1 + \tau_0\sigma_s^2 = \frac{\beta^2}{(1 - \beta)^2}.$$

Let us compare these two equilibria. The general version of the expected payoff inclusive of information costs is

$$\begin{aligned} & E^{\pi_0}[E^{\mu_0}[v(a, \bar{a}, s)]] - C_{MI}(\lambda_0, \tau_0, \Sigma_0, \bar{\alpha}_s, \bar{\alpha}_r) \\ &= \frac{\tau_0\psi(\bar{\alpha}_s, \bar{\alpha}_r)'\Sigma_0\lambda_0}{1 + \tau_0\lambda_0'\Sigma_0\lambda_0}(\psi(\bar{\alpha}_s, \bar{\alpha}_r)'\Sigma_0\lambda_0) \\ & - \psi(\bar{\alpha}_s, \bar{\alpha}_r)'\Sigma_0\psi(\bar{\alpha}_s, \bar{\alpha}_r) - \beta(1 - \beta)E^{\mu_0}[(s - \bar{\alpha}_s s - \bar{\alpha}_r r)^2] \\ & - \theta \ln(1 + \tau_0\lambda_0'\Sigma_0\lambda_0). \end{aligned}$$

With $\theta = (1 - \beta)^2\sigma_s^2$, $\lambda_0 = (1, 0)$, and $\psi(\bar{\alpha}_s, \bar{\alpha}_r) = (1 - \beta + \bar{\alpha}_s, 0)'$, define

$$J(\tau_0, \bar{\alpha}_s) = E^{\pi_0}[E^{\mu_0}[u(a, \bar{a}, s)]] - C_{MI}(\lambda_0, \tau_0, \Sigma_0, \bar{\alpha}_s, \bar{\alpha}_r)$$

as

$$\begin{aligned} J(\tau_0, \bar{\alpha}_s) &= (1 - \beta + \beta\bar{\alpha}_s)^2 \frac{\tau_0(\sigma_s^2)^2}{1 + \tau_0\sigma_s^2} \\ & - (1 - \beta + \beta\bar{\alpha}_s)^2\sigma_s^2 - \beta(1 - \beta)(1 - \bar{\alpha}_s)^2\sigma_s^2 \\ & - (1 - \beta)^2\sigma_s^2 \ln(1 + \tau_0\sigma_s^2). \end{aligned}$$

When $\bar{\alpha}_s = \tau_0 = 0$,

$$J(0, 0) = -(1 - \beta)\sigma_s^2.$$

In the other equilibrium,

$$\begin{aligned} J\left(\frac{2\beta - 1}{\beta}, \sigma_s^{-2}\left(\frac{\beta^2}{(1 - \beta)^2} - 1\right)\right) &= \beta^2\sigma_s^2 \frac{\beta^2 - (1 - \beta)^2}{\beta^2} \\ & - \beta^2\sigma_s^2 - \frac{(1 - \beta)}{\beta}(1 - \beta)^2\sigma_s^2 \\ & - (1 - \beta)^2\sigma_s^2 \ln\left(\frac{\beta^2}{(1 - \beta)^2}\right). \end{aligned}$$

which simplifies to

$$J\left(\frac{2\beta-1}{\beta}, \sigma_s^{-2}\left(\frac{\beta^2}{(1-\beta)^2} - 1\right)\right) = -(1-\beta)^2 \sigma_s^2 \left(\beta^{-1} + \ln\left(\frac{\beta^2}{(1-\beta)^2}\right)\right).$$

Defining $x = \frac{1-\beta}{\beta}$,

$$1 - \beta = 1 - \frac{1}{1+x} = \frac{x}{1+x},$$

$$f(x) = \frac{J\left(\frac{2\beta-1}{\beta}, \sigma_s^{-2}\left(\frac{\beta^2}{(1-\beta)^2} - 1\right)\right) - J(0,0)}{\beta(1-\beta)\sigma_s^2} = 1 - x^2 + 2x \ln(x).$$

For $x \in (0, 1]$, which corresponds to $\beta \in [\frac{1}{2}, 1)$, $f(1) = 0$ and

$$f'(x) = 2(1-x + \ln(x)) \leq 0.$$

Consequently,

$$J\left(\frac{2\beta-1}{\beta}, \sigma_s^{-2}\left(\frac{\beta^2}{(1-\beta)^2} - 1\right)\right) \geq J(0,0).$$

G Proof for the Continuous State Case

This section provides proofs for the claims in Proposition 9. The proofs are essentially notational adaptations of the proofs of Propositions 1, 6, 5, 6, and 7 in the main text.

G.1 Proof of Existence and Sufficiency of (Nowhere-)Monotonicity in R

Endow \mathcal{F}_0 with the sup-norm topology and \mathcal{U}_0 with the weak* topology. Endow $\bar{\Pi}_0$ with the weak* topology and $\bar{\mathcal{A}}^+$ with the sup-norm topology.

Lemma 21. *The function $\phi_F : \mathcal{F}_0 \rightarrow \mathcal{U}_0$ defined by $d\phi_F[f_0](s, r) = f_0(s, r)d\mu_0(s, r)$ is continuous.*

Proof. If $f_{0,n} \rightarrow f_0$ in the sup-norm topology, then for all continuous $h : S \times R \rightarrow \mathbb{R}$,

$$\int_{S \times R} h(s, r) f_{0,n}(s, r) d\mu_0(s, r) \rightarrow \int_{S \times R} h(s, r) f_0(s, r) d\mu_0(s, r)$$

by Aliprantis and Border [2006] 15.7 and the continuity of the product of continuous functions. \square

It is convenient to define a strategy $\bar{\pi}_0 \in \bar{\Pi}_0$ as s-measurable if, for all $(a, f_0) \in \text{supp}(\bar{\pi}_0)$, $f_0(s, r) = f_0(s, r')$ for all $s \in S$ and $r, r' \in R$. For s-measurable aggregate action functions $\bar{\alpha}$, this definition is equivalent to the requirement that the strategy $\pi \in \Pi(\phi[\mu_0, \bar{\alpha}])$ induced from π_0 by $(a, \mu'_0) \mapsto (a, \phi[\phi_F[f_0], \bar{\alpha}])$ be non-R-measurable.

Define $\bar{\Pi}_0^S \subseteq \bar{\Pi}_0$ as the set of s-measurable strategies.

Lemma 22. *The spaces \mathcal{F}_0 and \mathcal{U}_0 are metrizable in the sup-norm and weak* topologies, respectively. The sets $\bar{\Pi}_0$ and $\bar{\Pi}_0^S$ are non-empty, compact, and convex.*

Proof. See the technical appendix, [H.12](#) □

Define the objective function

$$\bar{J}(\bar{\pi}_0, \bar{\alpha}) = E^{\bar{\pi}_0} [V(a, \phi[\phi_F[f_0], \bar{\alpha}]) - D^{\mathcal{F}}(f_0 || \mu_0; \bar{\alpha})]. \quad (29)$$

The following lemma shows that this function is continuous.

Lemma 23. *The function $\bar{J}(\bar{\pi}_0, \bar{\alpha})$ is continuous.*

Proof. Observe that $A \times \mathcal{F}_0$ and $A \times \mathcal{U}_0$ are metrizable. By the continuity of ϕ_F and ϕ (Lemma 8, which applies to the continuous state case as well), the function $V(a, \phi[\phi_F[f_0], \bar{\alpha}]) = E^{\phi[\phi_F[f_0], \bar{\alpha}]} [u(a, \bar{a}, s)]$ is continuous on $A \times \mathcal{F}_0$ by the continuity of u and the definition of the weak* topology (see [Aliprantis and Border \[2006\]](#) 15.7). By assumption, $D^{\mathcal{F}}(f_0 || \mu_0; \bar{\alpha})$ is continuous. It follows almost immediately by [Aliprantis and Border \[2006\]](#) 15.7 that \bar{J} is continuous (see the proof of Lemma 9 for details). □

We next invoke the theorem of the maximum. Let $\bar{\Pi}^* : \bar{\mathcal{A}} \Rightarrow \bar{\Pi}_0$ be the best reply correspondence in the agent's problem,

$$\bar{\Pi}^*(\bar{\alpha}) = \{\bar{\pi}_0 \in \arg \max_{\bar{\pi}'_0 \in \bar{\Pi}_0} \bar{J}(\bar{\pi}'_0, \bar{\alpha})\}.$$

Define $\bar{\mathcal{A}}^S \subset \bar{\mathcal{A}}$ as the set of s-measurable aggregate action functions, and let $\bar{\Pi}_S^* : \bar{\mathcal{A}}^S \Rightarrow \bar{\Pi}_0^S$ be the best reply correspondence in the agent's problem restricted to s-measurable strategies,

$$\bar{\Pi}_S^*(\bar{\alpha}) = \{\bar{\pi}_0 \in \arg \max_{\bar{\pi}'_0 \in \bar{\Pi}_0^S} \bar{J}(\bar{\pi}'_0, \bar{\alpha})\}$$

Lemma 24. *The correspondences $\bar{\Pi}^*$ and $\bar{\Pi}_S^*$ are non-empty, compact-valued, convex-valued, and upper hemi-continuous.*

Proof. See [Aliprantis and Border \[2006\]](#) 17.31 (the maximum theorem). Convexity follows from the linearity of J in π_0 and the convexity of $\bar{\Pi}_0$ and $\bar{\Pi}_0^S$. □

Lemma 25. *The function $M : \bar{\Pi}_0 \rightarrow \bar{\mathcal{A}}^+$ defined by $M[\bar{\pi}_0](s, r) = \int_{\text{supp}(\bar{\pi}_0)} a f_0(s, r) d\bar{\pi}_0(a, f_0)$ for all $(s, r) \in S \times R$ is continuous. If $\bar{\pi}_0 \in \bar{\Pi}_0^S$, then $M[\bar{\pi}_0] \in \bar{\mathcal{A}}^S$.*

Proof. See the technical appendix, [H.13](#). □

It follows that the composition of M and $\bar{\Pi}^*$, $M \circ \bar{\Pi}^* : \bar{\Pi}_0 \Rightarrow \bar{\Pi}_0$, is a non-empty, upper hemi-continuous, and compact- and convex-valued. By the infinite-dimensional version of the Kakutani fixed point theorem ([Aliprantis and Border \[2006\]](#) 17.55), $M \circ \bar{\Pi}^*$ has a fixed point,

$\bar{\pi}_0^*$. By construction, for $\bar{\alpha}^* = M[\bar{\pi}_0^*]$, $\bar{\pi}_0^*$ is a best response to $\bar{\alpha}^*$, and $(\bar{\pi}_0^*, \bar{\alpha}^*)$ satisfies mean-consistency. It follows that $(\bar{\pi}_0^*, \bar{\alpha}^*)$ is an equilibrium, and hence that an equilibrium exists.

It also follows that the composition of M and $\bar{\Pi}_S^*$, $M \circ \bar{\Pi}_S^* : \Pi_0^S \Rightarrow \Pi_0^S$, is a non-empty, upper hemi-continuous, and compact- and convex-valued. By the infinite-dimensional version of the Kakutani fixed point theorem (Aliprantis and Border [2006] 17.55), $M \circ \bar{\Pi}_S^*$ has a fixed point, $\bar{\pi}_0^S$. By construction, for $\bar{\alpha}^S = M[\bar{\pi}_0^S]$, $\bar{\pi}_0^S$ is a best response to $\bar{\alpha}^S$ on the set of s-measurable strategies, $(\bar{\pi}_0^S, \bar{\alpha}^S)$ is mean-consistent, and both $\bar{\alpha}^S$ and $\bar{\pi}_0^S$ are s-measurable.

Let us now assume monotonicity in R . By the s-measurability of $\bar{\alpha}^S$, for all $\mu' \in \mathcal{U}$ with $\mu' \ll \phi[\mu_0, \bar{\alpha}^S]$,

$$V(a, \mu') = V(a, \eta_{-R|R}[\mu'|\mu]).$$

Consequently, for any $\pi \in \Pi(\phi[\mu_0, \bar{\alpha}^S])$, the s-measurable $\hat{\pi} \in \Pi(\phi[\mu_0, \bar{\alpha}^*])$ induced from π by $(a, \mu') \mapsto (a, \eta_{-R|R}[\mu'|\phi[\mu_0, \bar{\alpha}^S]])$ satisfies

$$E^{\hat{\pi}}[V(a, \mu')] = E^{\pi}[V(a, \mu')],$$

and by monotonicity in R for the prior $\phi[\mu_0, \bar{\alpha}^S]$, $E^{\hat{\pi}}[D(\mu'|\phi[\mu_0, \bar{\alpha}^S])] \leq E^{\pi}[D(\mu'|\phi[\mu_0, \bar{\alpha}^S])]$. It follows that every non-s-measurable π is weakly dominated by an s-measurable one.

By assumption, for any $\bar{\pi}_0 \in \bar{\Pi}_0$ that induces some $\pi \in \Pi(\phi[\mu_0, \bar{\alpha}^S])$ via the mapping $(a, f_0) \mapsto (a, \bar{\phi}[f_0, \bar{\alpha}^S])$, there exists an s-measurable $\bar{\pi}'_0 \in \bar{\Pi}_0$ that induces $\pi' \in \Pi(\phi[\mu_0, \bar{\alpha}^S])$ via this mapping, where π' is induced from π by the mapping $(a, \mu') \mapsto (a, \eta_{-R|R}[\mu'|\phi[\mu_0, \bar{\alpha}^S]])$. We must have $\bar{J}(\bar{\pi}_0, \bar{\alpha}) \leq \bar{J}(\bar{\pi}'_0, \bar{\alpha})$, and therefore $\bar{\pi}_0^S$ must be a best response without restricting to s-measurable strategies. It follows that $(\bar{\pi}_0^S, \bar{\alpha}^S)$ is an equilibrium.

By an argument identical to the proof of Proposition 6, if instead C is nowhere- R -monotone, an s-measurable equilibrium with non-zero information acquisition cannot exist.

G.2 Proof of Sufficiency of Invariance in \bar{A}

The planner's problem can be rewritten as $\max_{\bar{\pi}_0 \in \bar{\Pi}_0(\mu_0)} \bar{J}(\bar{\pi}_0, M[\bar{\pi}_0])$, where \bar{J} is defined in (29) above and $M[\bar{\pi}_0]$ is defined as in Lemma 25. The agent's problem is $\max_{\bar{\pi}_0 \in \bar{\Pi}_0(\mu_0)} \bar{J}(\bar{\pi}_0, \bar{\alpha})$, taking $\bar{\alpha} \in \bar{A}^+$ as given.

The set $\bar{\Pi}_0$ is convex and compact in the weak* topology, $M[\bar{\pi}_0]$ is continuous, and $\bar{J}(\bar{\pi}_0, \bar{\alpha})$ is continuous (see Lemmas 22, 23, and 25 in the proof of above). Consequently, a maximizer of the planner's problem exists.

Let $(\bar{\alpha}^+, \bar{\pi}_0^+)$ be a maximum of the planner's problem with $\bar{\alpha}^+ = M[\bar{\pi}_0^+]$. The function $\bar{J}(\bar{\pi}_0, M[\bar{\pi}_0])$ is directionally differentiable in $\bar{\pi}_0$ at $\bar{\pi}_0^+$:

$$\lim_{\epsilon \rightarrow 0^+} \frac{1}{\epsilon} (\bar{J}((1 - \epsilon)\bar{\pi}_0^+ + \epsilon\bar{\pi}'_0, M[(1 - \epsilon)\bar{\pi}_0^+ + \epsilon\bar{\pi}'_0]) - \bar{J}(\bar{\pi}_0^+, M[\bar{\pi}_0^+])) = \bar{J}(\bar{\pi}'_0, M[\bar{\pi}_0^+]) - \bar{J}(\bar{\pi}_0^+, M[\bar{\pi}_0^+])$$

for all $\bar{\pi}'_0 \in \bar{\Pi}_0$. This result follows from the observation that

$$E^{\bar{\pi}_0}[V(a, \phi[\phi_F[f_0], \bar{\alpha}])]$$

is directionally differentiable in $\bar{\alpha}$ (which follows from the mean-critical property, and note this derivative is zero) and the linearity of conditional expectations,

$$\lim_{\epsilon \rightarrow 0^+} \frac{1}{\epsilon} (M[(1 - \epsilon)\bar{\pi}_0^+ + \epsilon\bar{\pi}'_0] - M[\bar{\pi}_0^+]) = M[\bar{\pi}'_0] - M[\bar{\pi}_0^+],$$

along with invariance in \bar{A} (which ensures that the expected divergence is directionally differentiable with respect to $\bar{\alpha}$, with a derivative of zero).

Because $\bar{\pi}_0^+$ is maximal, we must have $\bar{J}(\bar{\pi}'_0, M[\bar{\pi}_0^+]) - \bar{J}(\bar{\pi}_0^+, M[\bar{\pi}_0^+]) \leq 0$ for all $\bar{\pi}'_0 \in \bar{\Pi}_0(\mu_0)$. Therefore, by $\bar{\alpha}^+ = M[\bar{\pi}_0^+]$, $\bar{\pi}_0^+$ is a best response in the agent's problem to $\bar{\alpha}^+$, and thus $(\bar{\alpha}^+, \bar{\pi}_0^+)$ is an equilibrium.

G.3 Proof of Uniqueness

By the assumption of invariance,

$$D(\bar{\phi}[f_0, \bar{\alpha}] || \phi[\mu_0, \bar{\alpha}]) = D(\bar{\phi}[f_0, \bar{\alpha}'] || \phi[\mu_0, \bar{\alpha}'])$$

for any arbitrary $\bar{\alpha}' \in \bar{\mathcal{A}}^+$.

The planner's problem can be rewritten as $\max_{\bar{\pi}_0 \in \bar{\Pi}_0(\mu_0)} \bar{J}(\bar{\pi}_0, M[\bar{\pi}_0])$, where \bar{J} is defined in (29) above and $M[\bar{\pi}_0]$ is defined as in Lemma 25. The agent's problem is $\max_{\bar{\pi}_0 \in \bar{\Pi}_0(\mu_0)} \bar{J}(\bar{\pi}_0, \bar{\alpha})$, taking $\bar{\alpha} \in \bar{\mathcal{A}}$ as given.

The set $\bar{\Pi}_0$ is convex and compact in the weak* topology, $M[\bar{\pi}_0]$ is continuous, and $\bar{J}(\bar{\pi}_0, \bar{\alpha})$ is continuous (see Lemmas 22, 23, and 25 in the proof of above). Consequently, a maximizer of the planner's problem exists.

Plugging in the functional form of a mean-critical utility function, for any mean-consistent $(\bar{\pi}_0, \bar{\alpha})$,

$$\bar{J}(\bar{\pi}_0, \bar{\alpha}) = E^{\bar{\pi}_0}[g(a; s) - D(\bar{\phi}[f_0, \bar{\alpha}'] || \phi[\mu_0, \bar{\alpha}'])] + E^{\mu_0}[G(\bar{\alpha}(s, r); s)].$$

If G is concave, J is concave over mean-consistent pairs: if $(\bar{\pi}_{0,1}, \bar{\alpha}_1)$ and $(\bar{\pi}_{0,2}, \bar{\alpha}_2)$ are mean-consistent pairs, then for $\lambda \in (0, 1)$, $(\lambda\bar{\pi}_{0,1} + (1 - \lambda)\bar{\pi}_{0,2}, \lambda\bar{\alpha}_1 + (1 - \lambda)\bar{\alpha}_2) \in \bar{\Pi}_0(\mu_0) \times \bar{\mathcal{A}}$ is mean-consistent and delivers weakly higher utility. If G is strictly concave, utility is strictly higher whenever $\bar{\alpha}_1 \neq \bar{\alpha}_2$ on a μ_0 -positive-measure set.

We next show that, with a divergence D invariant in \bar{A} and a mean-critical utility function, every equilibrium is a critical point of the planner's problem. Let $(\bar{\pi}_0^*, \bar{\alpha}^*)$ be an equilibrium. By the definition of the agent's problem, $\bar{\pi}_0^*$ is a critical value (a maximum) of $\bar{J}(\cdot, \bar{\alpha}^*) : \bar{\Pi}_0(\mu_0) \rightarrow \mathbb{R}$. By the definition of mean-consistency and the mean-critical property of the utility function, for

all $\bar{\alpha}'' \in \bar{\mathcal{A}}^+$,

$$\lim_{\epsilon \rightarrow 0^+} \frac{1}{\epsilon} \bar{J}(\bar{\pi}_0^*, \bar{\alpha}^* + \epsilon(\bar{\alpha}'' - \bar{\alpha}^*)) = 0$$

It follows that $(\bar{\pi}_0^*, \bar{\alpha}^*)$ is a critical value of $\bar{J}(\pi_0, \bar{\alpha})$, and hence that $(\bar{\pi}_0^*, \bar{\alpha}^*)$ is a critical value of the planner's problem.

If the planner's problem is concave, which occurs when G is concave, all critical points are maxima, and hence all equilibria are constrained-efficient. If G is strictly concave, two constrained efficient equilibria with different aggregate action functions on a μ_0 -positive-measure set cannot exist, as otherwise a convex combination would generate strictly higher utility, contradicting the assumption of constrained efficiency.

Lastly, observe that $\bar{\alpha}_1$ and $\bar{\alpha}_2$ are continuous functions (by the Arzela-Ascoli theorem and the compactness of $\bar{\mathcal{A}}^+$). Consequently, if they are equal μ_0 -almost-everywhere, they are in fact equal everywhere in the support of μ_0 (which, by assumption, is everywhere).

H Additional Technical Proofs

H.1 Proof of Lemma 7

Note: all parenthetical numbers in the proof refer to numbered results in [Aliprantis and Border \[2006\]](#).

S and R are assumed to be compact, separable, and metrizable; consequently, $S \times R$ is compact, separable, and metrizable (2.61). It follows that $\mathcal{U}_0 = \Delta(S \times R)$ is compact, separable, and metrizable (15.11, 15.12), and that $A \times \mathcal{U}_0$ is compact, separable, and metrizable in the product topology.

Consequently, $\Delta(A \times \mathcal{U}_0)$ is a compact, separable, metrizable space in its weak* topology (15.11, 15.12). By the compactness of $\Delta(A \times \mathcal{U}_0)$, $\Pi_0(\mu_0)$ is relatively compact in $\Delta(A \times \mathcal{U}_0)$ (15.21).

Moreover, $f(a, \mu'_0) = \mu'_0$ is a bounded continuous function on $A \times \mathcal{U}_0$, which is a compact set, and consequently for any sequence $\{\pi_n \in \Pi_0(\mu_0)\}$ that converges in the weak* topology to some $\pi_0 \in \Delta(A \times \mathcal{U}_0)$, we must have $\mu_0 = \int_{\text{supp}(\pi_n)} \mu'_0 d\pi_n(a, \mu'_0) \rightarrow \int_{\text{supp}(\pi_0)} \mu'_0 d\pi_0(a, \mu'_0)$, and therefore $\pi \in \Pi_0(\mu_0)$ (15.3). Thus, $\Pi_0(\mu_0)$ is closed (2.40), and hence is compact.

$\Pi_0(\mu_0)$ is non-empty: by definition, the point mass on (a, μ_0) for any $a \in A$ is an element of $\Pi_0(\mu_0)$.

It is also convex: for any $\pi_1, \pi_2 \in \Pi_0(\mu_0)$, $\int_{\text{supp}(\pi_1)} u'_0 d\pi_1(a, \mu'_0) = \mu_0$ and $\int_{\text{supp}(\pi_2)} u'_0 d\pi_2(a, \mu'_0) = \mu_0$, and consequently by the convexity of $\Delta(A \times \mathcal{U}_0)$, $\pi = \alpha\pi_1 + (1 - \alpha)\pi_2$ is an element of $\Pi_0(\mu_0)$ for any $\alpha \in (0, 1)$.

H.2 Proof of Lemma 8

Let $f : S \times R \times \bar{A} \rightarrow \mathbb{R}$ be a continuous and bounded function. Then for any $(\bar{\alpha}_n, \mu_{0,n}) \in \bar{\mathcal{A}} \times \mathcal{U}_0$,

$$\int_{S \times R \times \bar{A}} f(s, r, \bar{a}) d\phi[\mu_{0,n}, \bar{\alpha}_n](s, r, \bar{a}) = \int_{S \times R} f(s, r, \bar{\alpha}_n(s, r)) d\mu_{0,n}(s, r).$$

Suppose $(\bar{\alpha}_n, \mu_{0,n}) \rightarrow (\bar{\alpha}, \mu'_0)$. By definition, for any δ there exists an n_0 such that for all $n \geq n_0$,

$$\sup_{(s,r) \in S \times R} |\bar{\alpha}(s, r) - \bar{\alpha}_n(s, r)| < \delta.$$

By the definition of continuity, for any $\epsilon > 0$, there exists a $\delta > 0$ such that for all $(s, r) \in S \times R$ and all $\bar{a}, \bar{a}' \in \bar{A}$ with $|\bar{a} - \bar{a}'| < \delta$,

$$|f(s, r, \bar{a}) - f(s, r, \bar{a}')| < \epsilon.$$

Consequently, for any $\epsilon > 0$, there exists an n_0 such that

$$\sup_{(s,r) \in S \times R} |f(s, r, \bar{\alpha}(s, r)) - f(s, r, \bar{\alpha}_n(s, r))| < \epsilon,$$

which is to say $f(s, r, \bar{\alpha}_n(s, r)) \rightarrow f(s, r, \bar{\alpha}(s, r))$ in the sup norm topology.

It follows by [Aliprantis and Border \[2006\] 15.7](#) that

$$\int_{S \times R} f(s, r, \bar{\alpha}_n(s, r)) d\mu_{0,n}(s, r) \rightarrow \int_{S \times R} f(s, r, \bar{\alpha}(s, r)) d\mu'_0(s, r),$$

and hence that

$$\int_{S \times R \times \bar{A}} f(s, r, \bar{a}) d\phi[\mu_{0,n}, \bar{\alpha}_n](s, r, \bar{a}) \rightarrow \int_{S \times R \times \bar{A}} f(s, r, \bar{a}) d\phi[\mu'_0, \bar{\alpha}](s, r, \bar{a}).$$

H.3 Proof of Lemma 9

Concavity is immediate by linearity.

The function

$$V(a, \mu') = E^{\mu'}[v(a, \bar{a}, s)]$$

is continuous on $A \times \mathcal{U}$ by the continuity of v and the definition of the weak* topology (see [Aliprantis and Border \[2006\] 15.7](#)). The function $D(\mu' || \mu)$ is continuous on $\{(\mu', \mu) \in \mathcal{U} \times \mathcal{U} : \mu' \ll \mu\}$ by assumption.

By [Lemma 8](#), and continuity of the composition of continuous functions, it follows that

$$h(a, \mu'_0, \bar{\alpha}) = V(a, \phi[\mu'_0, \bar{\alpha}]) - D(\phi[\mu'_0, \bar{\alpha}] || \phi[\mu_0, \bar{\alpha}])$$

is continuous on $X = \{(a, \mu'_0, \bar{\alpha}) \in A \times \mathcal{U}_0 \times \bar{\mathcal{A}}\}$.

Suppose $(\bar{\alpha}_n, \pi_{0,n}) \rightarrow (\bar{\alpha}, \pi_0)$. By definition, for any δ there exists an n_0 such that for all $n \geq n_0$,

$$\sup_{(s,r) \in S \times R} |\bar{\alpha}(s,r) - \bar{\alpha}_n(s,r)| < \delta,$$

which is to say $|\bar{\alpha}_n - \bar{\alpha}| < \delta$ in the sup norm. By the definition of continuity, for any $\epsilon > 0$, there exists a $\delta > 0$ such that for all $(a, \mu'_0) \in A \times \mathcal{U}_0$, if $|\bar{\alpha}_n - \bar{\alpha}|$ in the sup norm, then

$$|h(a, \mu'_0, \bar{\alpha}_n) - h(a, \mu'_0, \bar{\alpha})| < \epsilon.$$

Consequently, $h(\cdot, \bar{\alpha}_n(s,r)) \rightarrow h(\cdot, \bar{\alpha}(s,r))$ in the sup norm topology.

It follows by [Aliprantis and Border \[2006\] 15.7](#) that

$$E^{\pi_{0,n}}[h(a, \mu'_0, \bar{\alpha}_n)] \rightarrow E^{\pi_0}[h(a, \mu'_0, \bar{\alpha})]$$

and hence that $J(\cdot)$ is continuous.

H.4 Proof of Lemma 10

Define the function $f_{s,r} : A \times \mathcal{U}_0 \rightarrow \mathbb{R}$ by

$$f_{s,r}[a, \mu'_0] = a \frac{\mu'_0(s,r)}{\mu_0(s,r)}.$$

By the finiteness of $S \times R$, $\mu'_{0,n} \rightarrow \mu'_0$ implies point-wise convergence. Consequently, if $(a_n, \mu'_{0,n}) \rightarrow (a, \mu'_0)$, then $f_{s,r}[a_n, \mu'_{0,n}] \rightarrow f_{s,r}[a, \mu'_0]$ for all $(s,r) \in S \times R$. Thus, $f_{s,r}$ is a continuous and bounded function for all $(s,r) \in S \times R$.

It follows immediately from the definition of the weak* topology (see, e.g., [Aliprantis and Border \[2006\] 15.3](#)) that if $\pi_{0,n} \rightarrow \pi_0$, then $M[\pi_{0,n}](s,r) \rightarrow M[\pi_0](s,r)$ for all $(s,r) \in S \times R$, which is to say that $M[\pi_{0,n}]$ converges point-wise to $M[\pi_0]$. By the finiteness of $S \times R$, point-wise convergence implies uniform convergence and hence $M[\pi_0]$ is continuous.

H.5 Proof of Lemma 17

By $\mu'' \ll \phi[\mu_0, \bar{\alpha}]$, there exists a $\mu''_0 \in \mathcal{U}_0$ such that $\mu'' = \phi[\mu''_0, \bar{\alpha}]$. Define $\mu = \phi[\mu_0, \bar{\alpha}]$.

Define the action

$$a_0 = E^{\mu_0}[\bar{\alpha}(s,r)].$$

Define, for some $\lambda > 1$, $\bar{A}_\lambda = \{a \in \mathbb{R}^L : \exists (s,r) \in S \times R \text{ s.t. } a = \lambda \bar{\alpha}(s,r) + (1-\lambda)a_0\}$. Because the image of $\bar{\alpha}$ lies in the relative interior of \bar{A} , by the finiteness of $S \times R$, for some λ sufficiently close to one, $\bar{A}_\lambda \subset \bar{A}$. Because the image of $\bar{\alpha}$ is not a singleton, a_0 is not in \bar{A}_λ . Define $a_1 \in A \setminus \bar{A}_\lambda$ as an arbitrary element of the relative interior of \bar{A}_λ not equal to a_0 .

Define $f_0 = \frac{d\mu'_0}{d\mu_0}$ and observe that, by the fact that $\bar{\alpha}$ is s -measurable and $\mu'' = \eta_{-R|R}[\mu'|\mu]$, f_0 does not depend on r . Now define a shifted $\bar{\alpha}$ function,

$$\begin{aligned}\tilde{\alpha}(s, r) &= \bar{\alpha}(s, r) + \frac{\epsilon \rho f_0(s)}{1 - \epsilon \rho f_0(s) - \epsilon(1 - \rho)} (\bar{\alpha}(s, r) - a_1) + \frac{\epsilon(1 - \rho)}{1 - \epsilon \rho f_0(s) - \epsilon(1 - \rho)} (\bar{\alpha}(s, r) - a_0) \\ &= \frac{\bar{\alpha}(s, r) - a_1 \epsilon \rho f_0(s) a_1 - \epsilon(1 - \rho) a_0}{1 - \epsilon \rho f_0(s) - \epsilon(1 - \rho)}\end{aligned}$$

for some $\epsilon \geq 0$ and $\rho \in (0, 1)$. By the interiority of the image of $\bar{\alpha}$, for sufficiently small ϵ ($\epsilon \leq \bar{\epsilon}$), $\tilde{\alpha}(s, r)$ remains in \bar{A}_λ . By the fact that $\bar{\alpha}$ is s -measurable, $\tilde{\alpha}(s, r)$ is s -measurable.

By construction, every value in the image of $\tilde{\alpha}$ is a convex combination of elements of \bar{A}_λ . Consequently, for each $s \in S$, there is a probability measure $\nu_s \in \Delta(\bar{A}_\lambda)$ satisfying

$$E^{\nu_s}[a] = \tilde{\alpha}(s, r)$$

for all $r \in R$ (by the s -measurability of $\bar{\alpha}$).

Now define $\nu_{s, \epsilon, \rho} \in \Delta(\bar{A}_\lambda \cup \{a_0, a_1\})$ by

$$\begin{aligned}\nu_{s, \epsilon, \rho}(a) &= (1 - \epsilon \rho f_0(s) - \epsilon(1 - \rho)) \nu_s(a) \\ &\quad + \epsilon \rho f_0(s) \delta_{a_1}(a) + \epsilon(1 - \rho) \delta_{a_0}(a),\end{aligned}$$

where $\delta_{a_1}(a)$ and $\delta_{a_0}(a)$ are the point masses on a_0 and a_1 .

By construction, for all $r \in R$,

$$\begin{aligned}E^{\nu_{s, \epsilon, \rho}}[a] &= (1 - \epsilon \rho f_0(s) - \epsilon(1 - \rho)) E^{\nu_s}[a] \\ &\quad + \epsilon \rho f_0(s) a_1 + \epsilon(1 - \rho) a_0 \\ &= \tilde{\alpha}(s, r).\end{aligned}$$

Define the finite set $A_{\epsilon, \rho}^* = \{a \in \bar{A}_\lambda \cup \{a_0, a_1\} : \exists s \in S \text{ s.t. } \nu_{s, \epsilon, \rho}(a) > 0\}$. Define, for each $a \in A_{\epsilon, \rho}^*$, the posterior measure $\mu_{0, a} \in \mathcal{U}_0$ by

$$\frac{d\mu_{0, a}(s, r)}{d\mu_0(s, r)} = \frac{\nu_{s, \epsilon, \rho}(a)}{\sum_{(s', r') \in S \times R} \mu_0(s', r') \nu_{s', \epsilon, \rho}(a)},$$

observing that this is well-defined by the finiteness of $S \times R$. Observe, by construction, that $\mu_{0, a_1} = \mu_0''$ and $\mu_{0, a_0} = \mu_0$. These posterior measures define a $\pi_{0, \epsilon, \rho} \in \Pi_0(\mu_0)$ that is a finite sum of point masses,

$$d\pi_{0, \epsilon, \rho}(a, \mu'_0) = \sum_{a' \in A_{\epsilon, \rho}^*} \sum_{(s', r') \in S \times R} \mu_0(s', r') \nu_{s', \epsilon, \rho}(a') \delta_{(a', \mu_{0, a'})}(a, \mu'_0).$$

This induces a $\pi_{\epsilon, \rho} \in \Pi(\mu)$ by $(a, \mu'_0) \mapsto (a, \phi[\mu'_0, \bar{\alpha}])$. By construction, $\pi_{\epsilon}(a_1, \mu'') = E^{\mu_0}[\nu_{s, \epsilon, \rho}(a_1)] = \rho$ and $\pi_{\epsilon}(a_0, \mu) = E^{\mu_0}[\nu_{s, \epsilon, \rho}(a_0)] = (1 - \rho)$, and for each $a \in A_{\epsilon, \rho}^*$, $(a, \mu') \in \text{supp}(\pi_{\epsilon, \rho})$ for at most

one μ' .

H.6 Proof of Lemma 18

If D is not R-monotone on the set of s-measurable priors, there exists an s-measurable $\bar{\alpha} \in \bar{\mathcal{A}}$, and $\mu' \in \mathcal{U}$ with $\mu' \ll \phi[\mu_0, \bar{\alpha}]$ such that

$$D(\mu' || \phi[\mu_0, \bar{\alpha}]) < D(\eta_{-R|R}[\mu' | \phi[\mu_0, \bar{\alpha}]] || \phi[\mu_0, \bar{\alpha}]).$$

Define $\mu = \phi[\mu_0, \bar{\alpha}]$.

Observe that if this inequality holds, it holds by continuity (Assumption 1) for some $\bar{\alpha}$ function whose image is strictly in the relative interior of $\bar{\mathcal{A}}$ and that is not a constant function. Assume these properties in what follows.

Define $\mu'' = \eta_{-R|R}[\mu' | \mu]$, and recall by the definition of η_R that $\mu'' \ll \mu$ and hence $\mu'' = \phi[\mu''_0, \bar{\alpha}]$ for some $\mu''_0 \in \mathcal{U}_0$ with $\mu''_0 \ll \mu_0$. We must have $\mu''_0 \neq \mu_0$, as otherwise $\mu'' = \mu$ and $D(\mu'' || \mu) = 0$.

Apply Lemma 17, which defines an $\epsilon > 0$, $\rho \in (0, 1)$, and $\pi_{\epsilon, \rho}$ that is mean-consistent with $\bar{\alpha}$ and exhibits several properties stated in that lemma. Let $\pi_0 \in \Pi_0(\mu_0)$ be the measure induced from $\pi_{\epsilon, \rho}$ by $(a, \mu''') \mapsto (a, \gamma_{-\bar{\mathcal{A}}}[\mu'''])$.

Let us now define a family $\pi_{0, \tau} \in \Pi_0(\mu_0)$ induced from π_0 by the mapping

$$(a, \mu'_0) \mapsto \begin{cases} (a, \mu'_0) & a \notin \{a_0, a_1\} \cup \mu'_0 \notin \{\mu''_0, \mu_0\}, \\ (a, \mu''_0 + \tau(\gamma_{-\bar{\mathcal{A}}}[\mu'] - \mu''_0)) & (a, \mu'_0) = (a_1, \mu''_0), \\ (a, \mu_0 - \frac{\rho}{1-\rho}\tau(\gamma_{-\bar{\mathcal{A}}}[\mu'] - \mu''_0)) & (a, \mu'_0) = (a_0, \mu_0) \end{cases}$$

for some $\tau > 0$. This induces a $\pi_\tau \in \Pi(\mu)$ by $(a, \mu'_0) \mapsto (a, \phi[\mu'_0, \bar{\alpha}])$, and this satisfies Bayes-consistency by $\pi_0(a_1, \mu''_0) = \epsilon\rho$ and $\pi_0(a_0, \mu_0) = \epsilon(1 - \rho)$.

By the differentiability of the divergence (Assumption 1) and the fact that the divergence is minimized when $\mu' = \mu$,

$$\frac{\partial}{\partial \tau} C(\pi_\tau, \mu)|_{\tau=0^+} = \epsilon\rho \frac{\partial}{\partial \tau} D(\phi[\mu''_0 + \tau(\gamma_{-\bar{\mathcal{A}}}[\mu'] - \mu''_0), \bar{\alpha}] || \mu)|_{\tau=0}.$$

By the convexity of the divergence,

$$\frac{\partial}{\partial \tau} D(\phi[\mu''_0 + \tau(\gamma_{-\bar{\mathcal{A}}}[\mu'] - \mu''_0), \bar{\alpha}] || \mu)|_{\tau=0} + D(\mu'' || \mu) \leq D(\mu' || \mu),$$

and therefore by $D(\mu' || \mu) < D(\mu'' || \mu)$,

$$\frac{\partial}{\partial \tau} C(\pi_\tau, \mu)|_{\tau=0^+} < 0.$$

It follows that for τ sufficiently small, $C(\pi_\tau, \mu) < C(\pi, \mu)$. By construction,

$$\begin{aligned}\eta_{-R|R}[\phi[\mu_0'' + \tau(\gamma_{\bar{A}}\{\mu'\} - \mu_0''), \bar{\alpha}]|\mu] &= \eta_{-R|R}[\mu'' + \tau(\mu' - \mu'')|\mu], \\ &= (1 - \tau)\eta_R\{\mu'', \mu\} + \tau\eta_{-R|R}[\mu'|\mu] \\ &= \mu''\end{aligned}$$

and likewise

$$\eta_{-R|R}[\phi[\mu_0 - \frac{\rho}{1-\rho}\tau(\gamma_{\bar{A}}[\mu'] - \mu_0''), \bar{\alpha}]|\mu] = \mu.$$

Therefore, π is induced from π_τ by $(a, \mu') \rightarrow (a, \eta_{-R|R}[\mu'|\mu])$, proving the claim for $\hat{\pi} = \pi_\tau$.

H.7 Proof of Lemma 19

Define $\mu = \phi[\mu_0, \bar{\alpha}]$. Given μ , define the function $D_S : \Delta(S) \rightarrow [0, \infty]$ by

$$D_S(\mu_s; \mu) = D(\eta_{S|R\bar{A}}[\mu_s|\mu]|\mu)$$

where

$$\frac{d\eta_{S|R\bar{A}}[\mu_s|\mu]}{d\mu} = \frac{d\mu_s}{d\gamma_{-R\bar{A}}[\mu]}$$

and $\gamma_{-R\bar{A}} : \Delta(S \times R \times \bar{A}) \rightarrow \Delta(S)$ computes the marginal distribution on S . This operator is like the η operator used in the main text, except that it acts on both the R and \bar{A} dimensions and takes as its first argument an element of $\Delta(S)$.

The function $D_S(\cdot; \mu)$ is convex on $\Delta(S)$, by the convexity of D . Let $\mathcal{M}(S)$ be the space of signed measures of bounded variation on S . The space of probability measures $\Delta(S)$ is a convex subset of this space. By convexity and finiteness (which is implied by Assumption 1), there exists a sub-gradient, meaning for any $\mu_s \in \Delta(S)$ and all $\mu'_s \in \Delta(S)$, there exists a function $h : S \rightarrow \mathbb{R}$ such that

$$D_S(\mu'_s; \mu) \geq D_S(\mu_s; \mu) + E^{\mu'_s}[h(s; \mu_s)] - E^{\mu_s}[h(s; \mu_s)].$$

We use the notation $h(s; \mu_s)$ to indicate that each $\mu_s \in \Delta(S)$ generates a different sub-gradient function. Observe that it is without loss of generality to suppose that $E^{\mu_s}[h(s; \mu_s)] = 0$, as $\hat{h}(s; \mu_s) = h(s; \mu_s) - E^{\mu_s}[h(s; \mu_s)]$ also satisfies the conditions of the sub-gradient.

By R-monotonicity and the s-measurability of $\bar{\alpha}$, for all $\mu' \in \mathcal{U}$ with $\mu' \ll \mu$,

$$D(\mu'|\mu) \geq D_S(\mu_s; \mu) + E^{\mu'_s}[h(s; \mu_s)].$$

Define $A^* \subset A$ as the set of actions with $(a, \mu') \in \text{supp}(\pi)$ for some $\mu' \in \mathcal{U}$, and let $\mu_a \in \mathcal{U}$ be the unique corresponding measure (recall that we have assumed in the statement of the lemma that each action is associated with at most one posterior under π).

Define, for all $a \in A^*$, $s \in S$, and an arbitrary $r_0 \in R$,

$$\tilde{v}(a, \bar{\alpha}(s, r_0), s) = D(\mu_a || \mu) + h(s; \gamma_{-R\bar{A}}[\mu_a]).$$

We extend the utility function \tilde{v} to other values of \bar{a} using, for all $a \in A^*$,

$$\tilde{v}(a, \bar{a}, s) = \tilde{v}(a, \bar{\alpha}(s, r_0), s) - \chi(a - \bar{\alpha}(s, r_0)) \cdot (\bar{a} - \bar{\alpha}(s, r_0)) + \frac{\chi}{2}(\bar{a} - \bar{\alpha}(s, r_0))^2.$$

We extend this utility function to other actions $a \in A \setminus A^*$ by defining

$$\tilde{v}(a, \bar{a}, s) = \min_{a' \in A^*} \tilde{v}(a', \bar{a}, s) - (a' - a)^2.$$

By construction, all actions not in A^* are dominated by some action in A^* . Consequently, with this utility function, any optimal policy will have support only on A^* .

The utility function \tilde{v} is not differentiable at certain points. However, we can define a mollified version of it,

$$v_\delta(a, \bar{a}, s) = \int_A \omega_\delta(a'' - a) \tilde{v}(a'', \bar{a}, s) da'',$$

where $\omega_\delta(z)$ is a smooth symmetric kernel with full support on $z < \delta$. Setting δ sufficiently small (less than $\min_{a, a' \in A^*} |a - a'|$) ensures that, for some constant c_δ that depends on the kernel,

$$v_\delta(a, \bar{\alpha}(s, r_0), s) = \tilde{v}(a, \bar{\alpha}(s, r_0), s) - c_\delta$$

for all $a \in A^*$, and that v_δ is continuously differentiable, while preserving the property that actions not in A^* are dominated.

Defining $v(a, \bar{a}, s) = v_\delta(a, \bar{\alpha}(s, r_0), s) + c_\delta$ results in the utility function

$$v(a, \bar{a}, s) = g(a; s) - \frac{\chi}{2}(\bar{a} - \bar{\alpha}(s, r_0))^2 - \chi(a - \bar{a}) \cdot (\bar{a} - \bar{\alpha}(s, r_0)),$$

where $g(a; s) = v_\delta(a, \bar{\alpha}(s, r_0), s) + c_\delta$.

Let us first show that $(\pi, \bar{\alpha})$ is an equilibrium. Consider the relaxed agent's problem,

$$\max_{\pi_0 \in \Delta(A \times \mathcal{U}_0)} E^{\pi_0} [E^{\mu'_0} [v(a, \bar{\alpha}(s, r_0), s)] - D(\phi[\mu'_0, \bar{\alpha}] || \mu)],$$

which is the agent's problem without the Bayes-consistency constraint. It is immediate a necessary and sufficient condition for optimality is that, for all $(a, \mu'_0) \in \text{supp}(\pi_0^*)$,

$$E^{\mu'_0} [v(a, \bar{\alpha}(s, r_0), s)] - D(\phi[\mu'_0, \bar{\alpha}] || \mu) \geq E^{\mu''_0} [v(a', \bar{\alpha}(s, r_0), s)] - D(\phi[\mu''_0, \bar{\alpha}] || \mu)$$

for all $(a', \mu''_0) \in A \times \mathcal{U}_0$. We will show this equation holds for all $a \in A^*$ and $\mu'_0 = \gamma_{-\bar{A}}[\mu_a]$.

First, note that because the actions in $A \setminus A^*$ are dominated by actions in A^* , no such action is optimal in this problem.

Second, choosing $a' = a$ and defining $\mu'' = \phi[\mu''_0, \bar{\alpha}]$ yields, by the definition of the utility

function,

$$E^{\mu'_0} [v(a, \bar{\alpha}(s, r_0), s)] = D(\mu_a || \mu) + E^{\mu_a} [h(s; \gamma_{-R\bar{A}}[\mu_a])],$$

$$E^{\mu''_0} [v(a', \bar{\alpha}(s, r_0), s)] = D(\mu_a || \mu) + E^{\mu''_0} [h(s; \gamma_{-R\bar{A}}[\mu_a])],$$

and therefore by $E^{\mu_a} [h(s; \gamma_{-R\bar{A}}[\mu_a])] = 0$,

$$0 \geq D(\mu_a || \mu) - D(\mu' || \mu) + E^{\mu'} [h(s; \gamma_{-R\bar{A}}[\mu_a])],$$

which holds by the definition of the sub-gradient.

Third, by construction, for all $a \in a^*$,

$$E^{\mu_a} [v(a, \bar{a}, s)] - D(\mu_a || \mu) = 0,$$

and consequently π is a maximizer of the actual agent's problem. It follows that $(\pi, \bar{\alpha})$ is an equilibrium.

H.8 Proof of Lemma 20

$\Pi_0^S(\mu_0)$ is non-empty: by definition, the point mass on (a, μ_0) for any $a \in A$ is an element of $\Pi_0(\mu_0)$.

It is also convex: for any $\pi_1, \pi_2 \in \Pi_0^S(\mu_0)$, $\int_{\text{supp}(\pi_1)} u'_0 d\pi_1(a, \mu'_0) = \mu_0$ and $\int_{\text{supp}(\pi_2)} u'_0 d\pi_2(a, \mu'_0) = \mu_0$, and consequently $\pi = \alpha\pi_1 + (1 - \alpha)\pi_2$ is an element of $\Pi_0^S(\mu_0)$ for any $\alpha \in (0, 1)$, as $\text{supp}(\pi) = \text{supp}(\pi_1) \cup \text{supp}(\pi_2)$.

$\Pi_0^S(\mu_0)$ is a subset of $\Pi_0(\mu_0)$, which is compact by Lemma 7. It follows that it is sufficient to prove it is closed. Define, for any s-measurable $\bar{\alpha}_s \in \bar{\mathcal{A}}$,

$$f(a, \mu'_0) = D_{KL}(\phi[\mu'_0, \bar{\alpha}_s] || \eta_{-R|R}[\phi[\mu'_0, \bar{\alpha}_s] | \phi[\mu_0, \bar{\alpha}_s]]).$$

Observe that this function is zero for all $(a, \mu'_0) \in \text{supp}(\pi_0)$ if and only if $\pi_0 \in \Pi_0^S(\mu_0)$, and that this function is bounded and continuous. It follows that if a sequence $\{\pi_n \in \Pi_0^S(\mu_0)\}$ converges to some $\pi_0 \in \Pi_0(\mu_0)$, then

$$0 = \int_{A \times \mathcal{U}_0} f(a, \mu'_0) d\pi_n(a, \mu'_0) \rightarrow \int_{A \times \mathcal{U}_0} f(a, \mu'_0) d\pi_0(a, \mu'_0),$$

and hence that $\pi_0 \in \Pi_0^S(\mu_0)$, demonstrating that $\Pi_0^S(\mu_0)$ is a closed subset of a compact set and therefore compact.

H.9 Proof of Lemma 16

Let us think of $\omega = (a, \mu'_0) \in \text{supp}(\pi_0^*)$ as a signal realization. Define the conditional measure

$$d\nu(\omega|s, r) = \frac{\mu'_0(s, r; \omega)}{\mu_0(s, r)} d\pi_0^*(\omega)$$

where $\mu'_0(s, r; (a, \mu''_0)) = \mu''_0(s, r)$. By $a > \bar{\alpha}(s_0, r_0)$ for all $(a, \mu'_0) \in \Omega_0$ and mean-consistency,

$$\nu(\Omega_0|s_0, r_0) = \int_{\Omega_0} d\nu(\omega|s_0, r_0) \in (0, 1).$$

Now define a new family of conditional measures parametrized by $\epsilon \in \mathbb{R}$,

$$d\nu_\epsilon(\omega|s, r) = \begin{cases} d\nu(\omega|s, r) & (s, r) \neq (s_0, r_0), \\ (1 - \epsilon + \frac{\epsilon}{\nu(\Omega_0|s, r)}) d\nu(\omega|s, r) & (s, r) = (s_0, r_0), \omega \in \Omega_0, \\ (1 - \epsilon) d\nu(\omega|s, r) & \text{otherwise.} \end{cases}$$

Observe that ν_ϵ is a valid conditional probability measure for all $\epsilon \in [-\frac{\nu(\Omega_0|s_0, r_0)}{1 - \nu(\Omega_0|s_0, r_0)}, 1]$; define $\delta = \min\{\frac{\nu(\Omega_0|s_0, r_0)}{1 - \nu(\Omega_0|s_0, r_0)}, 1\}$. Define the unconditional signal measure

$$d\pi_\epsilon(\omega) = \begin{cases} d\pi_0^*(\omega)(1 - \epsilon\mu'_0(s_0, r_0; \omega)) & \omega \neq \Omega_0 \\ d\pi_0^*(\omega)(1 + (\frac{\epsilon}{\nu(\Omega_0|s_0, r_0)} - \epsilon)\mu'_0(s_0, r_0; \omega)) & \omega = \Omega_0 \end{cases}$$

and define the associated posteriors measures by

$$\frac{\mu'_{0, \epsilon}(s, r; \omega)}{\mu_0(s, r)} = \begin{cases} \frac{\mu'_0(s, r; \omega)}{1 + (\frac{\epsilon}{\nu(\Omega_0|s_0, r_0)} - \epsilon)\mu'_0(s_0, r_0; \omega)} & (s, r) \neq (s_0, r_0), \omega \in \Omega_0, \\ \frac{(1 - \epsilon + \frac{\epsilon}{\nu(\Omega_0|s_0, r_0)})\mu'_0(s, r; \omega)}{1 + (\frac{\epsilon}{\nu(\Omega_0|s_0, r_0)} - \epsilon)\mu'_0(s_0, r_0; \omega)} & (s, r) = (s_0, r_0), \omega \in \Omega_0, \\ \frac{\mu'_0(s, r; \omega)}{1 - \epsilon\mu'_0(s_0, r_0; \omega)} & (s, r) \neq (s_0, r_0), \omega \in \text{supp}(\pi_0^*) \setminus \Omega_0, \\ \frac{(1 - \epsilon)\mu'_0(s, r; \omega)}{1 - \epsilon\mu'_0(s_0, r_0; \omega)} & (s, r) = (s_0, r_0), \omega \in \text{supp}(\pi_0^*) \setminus \Omega_0. \end{cases}$$

Let $m_\epsilon(a, \mu'_0) = \mu'_{0, \epsilon}(s, r; a, \mu'_0)$. Let $\pi_{0, \epsilon}$ be the measure induced by $\omega \mapsto (a, \mu'_{0, \epsilon}(\omega))$, which by construction is differentiable in ϵ . By construction,

$$E^{\pi_{0, \epsilon}}[a|s, r] = \begin{cases} \bar{\alpha}^*(s, r) & (s, r) \neq (s_0, r_0) \\ (1 - \epsilon)\bar{\alpha}^*(s, r) + \epsilon \frac{\int_{\Omega_0} a(\omega) d\nu(\omega|s_0, r_0)}{\nu(\Omega_0|s_0, r_0)} & \text{otherwise,} \end{cases}$$

and by the definition of Ω_0 ,

$$x = \frac{\int_{\Omega_0} a(\omega) d\nu(\omega|s_0, r_0)}{\nu(\Omega_0|s_0, r_0)} > 0.$$

H.10 Proof of Lemma 4

Define $\pi_0 \in \Pi_0(\mu_0)$ as the measure induced from π by $(a, \mu') \mapsto (a, \gamma_{-\bar{A}}[\mu'])$. By Bayes-consistency,

$$E^{\pi_0}[\mu'_0] = \mu_0.$$

By assumption, Ω is equal to $A \times \mathcal{U}_0$. Let $\bar{\nu}$ be the probability measure on Ω induced from π by the homeomorphism $h : A \times \mathcal{U}_0 \rightarrow \Omega$. By the Radon-Nikodym theorem, there exists a set of functions $h_\omega : S \times R \rightarrow \mathbb{R}_+$ for all $\omega \in \text{supp}(\bar{\nu})$ such that

$$\int_{\Omega} h_\omega(s, r) d\bar{\nu}(\omega) = 1, \mu_0\text{-a.e.}$$

Define

$$\bar{h}(s, r) = 1 - \int_{\Omega} h_\omega(s, r) d\bar{\nu}(\omega)$$

and

$$f_\omega(s, r) = \begin{cases} 1 & \bar{h}(s, r) \neq 0 \\ h_\omega(s, r) & \bar{h}(s, r) = 0. \end{cases}$$

We have $f_\omega(s, r) = h_\omega(s, r)$ μ_0 -a.e, and therefore $d\mu'_0(s, r) = f_\omega(s, r)d\mu_0(s, r)$. Moreover, by construction,

$$\int_{\Omega} f_\omega(s, r) d\bar{\nu}(\omega) = 1.$$

Define

$$\nu(\omega|s, r, \bar{a}) = f_\omega(s, r)$$

for all $\omega \in \text{supp}(\bar{\nu})$, and as zero otherwise. We have, for all $\omega \in \text{supp}(\bar{\nu})$,

$$\bar{f}[\nu, \mu](\omega) = \int_{S \times R} f_\omega(s, r) d\mu_0(s, r) = 1$$

and

$$f^\omega[\nu, \mu](s, r, \bar{a}) = f_\omega(s, r),$$

and therefore $\hat{\pi}[\bar{\nu}, \nu, \mu] = \pi$.

H.11 Proof of Lemma 5

Consider first D_{FI-SR} . Note that with D_{FI-SR} , any signal structure that does not depend on \bar{a} is optimal.

By Bayes-consistency,

$$1 = \int_{\text{supp}(\bar{\pi}_0)} f_0(s, r) d\bar{\pi}_0(a, f_0).$$

By assumption, Ω is homeomorphic (in fact equal to) $A \times \mathcal{F}_0^+$. Let $\bar{\nu}$ be the probability measure on Ω induced from π by the homeomorphism $h : A \times \mathcal{F}_0^+ \rightarrow \Omega$. Define $f_{h(a, f_0)} = f_0$ and define

$$\nu(\omega|s, r, \bar{a}) = f_\omega(s, r)$$

for all $\omega \in \text{supp}(\bar{\nu})$, and as zero otherwise. We have, for all $\omega \in \text{supp}(\bar{\nu})$,

$$\bar{f}\{\nu, \mu\}(\omega) = \int_{S \times R} f_\omega(s, r) d\mu_0(s, r) = 1$$

and

$$f^\omega\{\nu, \mu\}(s, r, \bar{a}) = f_\omega(s, r),$$

and therefore $\hat{\pi}\{\bar{\nu}, \nu, \mu\} = \pi$.

Now consider $D_{FI-S\bar{A}}$. Construct $(\bar{\nu}, \nu)$ as above. Now define

$$h(s) = \frac{\bar{\alpha}_s(s, r)}{1 + \bar{\alpha}_s(s, r)^2} \frac{\partial}{\partial s'} \nu(\omega|s', r, \bar{\alpha}(s, r))|_{s'=1}.$$

Observe by the uniform Lipschitz-continuity of the derivative of $f_0 \in \mathcal{F}_0^+$ and compactness of $S \times R \times \bar{A}$ that $h(s)$ is bounded. Define

$$\begin{aligned} \nu_1(\omega|s, r, \bar{a}) &= \min(\max(\bar{a} - \bar{\alpha}(s, r), -\delta), \delta) \cdot h(s) \\ &\quad + \nu(\omega|s, r, \bar{\alpha}(s, r)). \end{aligned}$$

For sufficiently small δ , by the boundedness of $h(s)$, this is non-negative and the optimality condition (24).

We can verify that

$$\int_{\text{supp}(\bar{\nu})} \nu_1(\omega|s, r, \bar{a}) d\bar{\nu}(\omega) = 1$$

by observing that

$$\frac{\bar{\alpha}_s(s, r)}{1 + \bar{\alpha}_s(s, r)^2} \int_{\text{supp}(\bar{\nu})} \frac{\partial}{\partial s'} \nu(\omega|s', r, \bar{a})|_{s'=s} d\bar{\nu}(\omega) = 0$$

everywhere.

H.12 Proof of Lemma 22

Note: numbers refer to results in [Aliprantis and Border \[2006\]](#).

$S \times R$ is compact, metrizable, and separable, and \mathbb{R} is separable and metrizable; by 3.96 and 3.99 \mathcal{F}_0 is metrizable and separable, and by 15.11 and 15.12 \mathcal{U}_0 is compact, metrizable, and separable.

By assumption, $\mathcal{F}_0^+ \subset \mathcal{F}_0$ is compact. Consequently, $\Delta(A \times \mathcal{F}_0^+)$ is a compact, separable, metrizable space in its weak* topology (15.11, 15.12). By the compactness of $\Delta(A \times \mathcal{F}_0^+)$, $\bar{\Pi}_0$ is relatively compact in $\Delta(A \times \mathcal{F}_0^+)$ (15.21).

Moreover, $h(a, f_0) = f_0$ is a bounded continuous function on $A \times \mathcal{F}_0^+$, which is a compact set, and consequently for any sequence $\{\bar{\pi}_n \in \bar{\Pi}_0\}$ that converges in the weak* topology to some $\bar{\pi}_0 \in \Delta(A \times \mathcal{F}_0^+)$, we must have $1 = \int_{\text{supp}(\bar{\pi}_n)} f_0 d\bar{\pi}_n(a, f_0) \rightarrow \int_{\text{supp}(\bar{\pi}_0)} f_0 d\bar{\pi}_0(a, f_0)$, and therefore

$\bar{\pi}_0 \in \bar{\Pi}_0$ (15.3). Thus, $\bar{\Pi}_0$ is closed (2.40), and hence is compact.

$\bar{\Pi}_0$ is non-empty: by assumption, the point mass on (a, ι) for any $a \in A$, where ι is a function equal to one everywhere, is an element of $\bar{\Pi}_0$.

It is also convex: for any $\bar{\pi}_1, \bar{\pi}_2 \in \bar{\Pi}_0$, $\int_{\text{supp}(\bar{\pi}_1)} f_0 d\pi_1(a, f_0) = 1$ and $\int_{\text{supp}(\bar{\pi}_2)} f_0 d\pi_2(a, f_0) = 1$, and consequently by the convexity of $\Delta(A \times \mathcal{F}_0^+)$, $\pi = \alpha\pi_1 + (1 - \alpha)\pi_2$ is an element of $\bar{\Pi}_0$ for any $\alpha \in (0, 1)$.

Define $\mathcal{F}_{0,S}^+ \subset \mathcal{F}_0^+$ as the subset of functions satisfying $f_0(s, r) = f_0(s, r')$ for all $s \in S$ and $r, r' \in R$. We have $\bar{\Pi}_0^S \subset \Delta(A \times \mathcal{F}_{0,S}^+)$, and by essentially identical arguments, $\bar{\Pi}_0^S$ is non-empty, convex, and compact.

H.13 Proof of Lemma 25

Define the function $f_{s,r} : A \times \mathcal{F}_0^+ \rightarrow \mathbb{R}$ by

$$f_{s,r}[a, f_0] = af_0(s, r).$$

Recall that A is endowed with the Euclidean topology, \mathcal{F}_0^+ with the sup-norm topology, and $A \times \mathcal{F}_0^+$ with the product topology. Observe that

$$\begin{aligned} |a_n f_{0,n}(s, r) - af_0(s, r)| &\leq |a_n - a| \sup_{(s', r') \in S \times R} |f_{0,n}(s', r')| \\ &\quad + |a| \sup_{(s', r') \in S \times R} |f_{0,n}(s', r') - f_0(s', r')| \end{aligned}$$

and consequently, by the compactness of \mathcal{F}_0^+ and A , if $(a_n, f_{0,n}) \rightarrow (a, f_0)$, then $f_{s,r}[a_n, f_{0,n}] \rightarrow f_{s,r}[a, f_0]$ for all $(s, r) \in S \times R$. Consequently, $f_{s,r}$ is a continuous and bounded function for all $(s, r) \in S \times R$.

It follows immediately from the definition of the weak* topology (see, e.g., [Aliprantis and Border \[2006\]](#) 15.3) that if $\bar{\pi}_{0,n} \rightarrow \bar{\pi}_0$, then $M[\bar{\pi}_{0,n}](s, r) \rightarrow M[\bar{\pi}_0](s, r)$ for all $(s, r) \in S \times R$, which is to say that $M[\bar{\pi}_{0,n}]$ converges point-wise to $M[\bar{\pi}_0]$. By assumption, the image of $M[\bar{\pi}_0], \bar{\mathcal{A}}^+$, is compact. Consequently, by the Arzela-Ascoli theorem, the sequence $M[\bar{\pi}_{0,n}]$ has a convergent sub-sequence in the sup-norm topology, which must converge uniformly to $M[\bar{\pi}_0]$. It follows that $M[\bar{\pi}_0]$ is continuous.

Observe that if $\bar{\pi}_0$ is s-measurable, then for all $(a, f_0) \in \text{supp}(\bar{\pi}_0)$, $f_0(s, r) = f_0(s, r')$ for all $s \in S$ and $r, r' \in R$. It follows immediately that $M[\bar{\pi}_0](s, r) = M[\bar{\pi}_0](s, r')$ for all $s \in S$ and $(r, r') \in R$.