Information Acquisition, Efficiency, and Non-Fundamental Volatility[*]

Benjamin Hébert[†]   Jennifer La'O[‡]

April 26, 2021

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 as well as endogenous aggregate actions. We study how properties of information costs relate to properties of equilibria. We provide necessary and sufficient conditions to guarantee zero non-fundamental volatility in equilibrium and another set of necessary and sufficient conditions to guarantee constrained efficient equilibria. Mutual information, the cost typically assumed in rational inattention models, precludes non-fundamental volatility and imposes efficiency.

JEL codes: C72, D62, D83

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

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 [Morris and Shin, 2002, Angeletos and Pavan, 2007]. These games of strategic interaction form the underlying basis for 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 of the game features or does not feature non-fundamental volatility? Second, what properties of the agents' information acquisition costs guarantee that an equilibrium exists that is constrained efficient?

Our framework. We study a class of large games of strategic interaction. 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 as well as with the endogenous mean action.

Agents acquire information endogenously. In particular, we adopt the rational inattention approach to costly information acquisition proposed by Sims [2003]. However, relative to the standard rational inattention framework, we make three significant departures.

First, we do not assume information costs are proportional to mutual information—the typical cost function introduced by Sims [2003] and used widely throughout the rational inattention literature. We instead consider a more general class of cost functions: costs that are "posterior-separable" in the terminology of Caplin, Dean, and Leahy [2021]. 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 [2021], and the neighborhood-based cost functions of Hébert and Woodford [2021]. These 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]).

Second, we introduce exogenous noisy public signals. Noisy public signals are standard components of exogenous information sets in beauty contest games [Morris and Shin, 2002,

Angeletos and Pavan, 2007, Bergemann and Morris, 2013]. In these games, noisy public signals play a major role in determining equilibrium outcomes. As a result, errors in these signals orthogonal to fundamentals manifest as "non-fundamental" volatility in equilibrium.

Noisy public signals are less standard in rational inattention models, which generally assume that agents pay attention only to payoff-relevant states. By allowing for the existence of noisy public signals, we introduce a channel by which agents can choose to coordinate, as well as a potential source of equilibrium "non-fundamental" volatility.

Third, we do not restrict agents to acquiring information only about exogenous states; agents in our model can also learn by observing the endogenous mean action of other agents. This modeling choice is motivated by the observation that agents often rely on information that aggregates the actions of other agents (e.g. government statistics and prices) when making their economic decisions. In this sense our work is similar in spirit to Denti [2020] and Angeletos and Sastry [2021]. Both of these papers consider environments in which rationally inattentive agents can acquire information about endogenous actions; for example, in Angeletos and Sastry [2021] agents can learn from equilibrium prices in a Walrasian market.

Within this context, we answer the questions posed above. What properties of information costs ensure that an equilibrium of the game does or does not exhibit non-fundamental volatility? What properties of information costs guarantee that an equilibrium exists that is constrained efficient? 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 a cost function leads to non-fundamental volatility or inefficiency depends on two key properties: "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 some dimension of the aggregate 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 *monotone in R* 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.

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 *invariant* in R 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.

The forms of partial monotonicity and partial invariance that we introduce are generalizations of the invariance concept described in the literature on information geometry [Chentsov, 1982, Amari and Nagaoka, 2007]. Invariance of the divergence in this sense holds if and only if the agents' behavior is characterized by the "invariance under compression" axiom of Caplin, Dean, and Leahy [2021]. This form of invariance has been applied to particular economic applications by Hébert [2018] and Hébert and Woodford [2019].

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 partial monotonicity of the cost function in noisy public signals is sufficient and (subject to some technical caveats) necessary to ensure the existence of an equilibrium that features zero non-fundamental volatility. With such a cost function, it is always cheaper for agents to condition their signals on the payoff-relevant states and aggregate actions than to condition their actions on noisy public signals. But if conditioning on public signals only increases information costs, then an equilibrium exists in which all agents ignore such signals. This equilibrium features zero non-fundamental volatility.

On the other hand, when the cost function is non-monotonic in noisy public signals, agents find it relatively "easier" or "cheaper" to condition their actions on the noisy public signals than to pay attention to the payoff-relevant states directly. This information cost-saving incentive is what induces agents to condition their actions on public signals, introducing correlated noise across agents and, as a result, non-fundamental volatility in equilibrium.

We find that two conditions are sufficient and (again subject to some technical caveats) necessary for constrained efficiency. First, payoff externalities, which arise from our assumption that the actions of other agents enter each agents' utility function, must be absent at the margin in equilibrium. We characterize the class of utility functions exhibiting this property; we call these "mean-critical" utility functions. This result extends prior results by Angeletos and Pavan [2007] to our more general setting.

Second, information costs must exhibit partial invariance with respect to the endogenous

mean action. We show that, given any mean-critical utility function, this form of partial invariance is a sufficient condition for constrained efficiency. Moreover, a local form of partial invariance on the relevant part of the domain is necessary.

When cost functions are not invariant in the aggregate action, agents' actions affect the ease by which other agents acquire information. For example, suppose agents learn in part by paying attention to the aggregate action, and more extreme aggregate actions are easier to observe than less extreme actions. This is an externality—agents do not internalize how their own strategies affect the information acquisition costs of others. The planner in this context would like to encourage more extreme actions to in order to reduce information acquisition costs.

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 the properties of equilibria.

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 non-monotone in the noisy public signals, and the divergence associated with the other is not 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, it precludes both non-fundamental volatility and inefficiency.

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 B.

Related Literature. 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, and exogenously specified Gaussian signals about exogenous states.

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. As a result, we are able to accommodate non-fundamental volatility, building a bridge between these two seemingly distinct approaches.

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 settings with general payoff functions, and obtain necessary and sufficient conditions for constrained efficiency with endogenous information acquisition. However, we shut down a key channel present in Angeletos and Pavan [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.[1] 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 this respect, our paper complements Angeletos and Sastry [2021]. Both papers take up the question of efficiency with rationally inattentive agents in different settings. 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.[2]

---

[1] Colombo, 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 externalities that arise in our game from the absence of invariance.

[2] The 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

Our focus on games with agents who can acquire information about the endogenous actions of other agents builds on Denti [2020]. 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 [2020] to static games with a continuum of identical players. Relative to Denti [2020], the "largeness" feature of our class of games permits a simpler definition of equilibrium, which is essentially 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 [2020] introduces. Our exercise is similar in spirit to Morris and Yang [2019], in that we relate the properties of information costs to the properties of equilibria in strategic settings.

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 $S \subset \mathbb{R}^n$.[3] We refer to these states as the "fundamentals."

Agents have a payoff function $u : A \times \bar{A} \times S \to \mathbb{R}$; 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 $u(a, \bar{a}, s)$. 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 first present our results in the context of a specific linear-quadratic beauty-contest game, similar to the one studied in Morris and Shin [2002] and Angeletos and Pavan [2007], before generalizing our results to broader classes of utility functions. We initially assume payoffs take

price function is an efficient conveyor of information in the sense of Hayek [1945].

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

the following linear-quadratic form:

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

where $\beta < 1$ is a constant. The first component of (1) is a quadratic loss in the distance between the agent's action a and fundamental s; the second component is a quadratic loss in the distance between the agent's action a and the aggregate action $\bar{a}$. The scalar β governs the extent of strategic complementarity ($\beta > 0$) or substitutability ($\beta < 0$) in this game; the assumption that $\beta < 1$ ensures that a unique symmetric pure-strategy Nash equilibrium exists under complete information given by $a^i = \bar{a} = s$.

We assume the linear-quadratic payoff structure in (1) to focus our analysis on the connection between properties of information acquisition costs and properties of equilibria. Our findings, however, are not specific to the linear-quadratic beauty contest game; Section 7 shows how our results generalize to a larger class of payoff environments.

2.2 Information Acquisition

Agents are rationally inattentive: they acquire information 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 shocks $e \in \mathcal{E}$ that are orthogonal to the fundamentals $s \in S$, and then generating the public signal r through some mapping $\mathcal{E} \times S \to R$. 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 $S \times R$, endowed with the weak* topology. 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 also pay attention to the endogenous aggregate action $\bar{a}$. Let $\bar{\alpha} : S \times R \to \bar{A}$ be a function mapping exogenous states to an aggregate action; this function will be determined in equilibrium by aggregating over the individual agents' strategies.[4] Let $\bar{A}$ be the space of all such aggregate action functions, endowed with the

---

[4]Note that we assume $\bar{\alpha}$ is a deterministic function of (s, r); we will subsequently define an equilibrium in a manner consistent with this assumption.

sup-norm topology.

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 $\phi : \mathcal{U}_0 \times \bar{\mathcal{A}} \to \Delta(S \times R \times \bar{A})$ denote a mapping from any pair $(\mu_0, \bar{\alpha})$ to its induced probability measure, defined by, for all $(s, r, \bar{a}) \in S \times R \times \bar{A}$,

$$d\phi\{\mu_0, \bar{\alpha}\}(s, r, \bar{a}) = d\mu_0(s, r)d\delta_{\bar{\alpha}(s,r)}(\bar{a}), \tag{2}$$

where $\delta_{\bar{\alpha}(s,r)}$ is the point mass at $\bar{\alpha}(s, r)$. 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 $\bar{a} = \bar{\alpha}(s, r)$.

This degeneracy arises from two features of our model: 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, and our choice to consider learning about endogenous outcomes. These features arise in, for example, any model that features Walrasian markets with price-taking and learning via prices (in this context, prices are a function of the aggregate action).

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

$$\mathcal{U} = \{\mu \in \Delta(S \times R \times \bar{A}) : \exists(\mu_0', \bar{\alpha}) \in \mathcal{U}_0 \times \bar{\mathcal{A}} \text{ s.t. } \mu = \phi\{\mu_0', \bar{\alpha}\}\}.$$

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 signals: 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, $\pi \in \Delta(A \times \mathcal{U})$.[5] We require strategies to be "Bayes-consistent" with the prior: the expectation of the posterior under measure π must equal to the prior, $\mu \in \mathcal{U}$.[6]

---

[5]This joint distribution over actions and posteriors can be thought of as summarizing the agent's state-dependent stochastic choice (Caplin and Dean [2015]).

[6]We discuss the relationship between signals, actions, and distributions over posteriors and actions in detail in Appendix Section A, and see also e.g. Proposition 1 of Kamenica and Gentzkow [2011].

Let $\mu' \in \text{supp}(\pi) \subseteq \mathcal{U}$ denote a posterior measure in the support of π. A strategy $\pi \in \Delta(A \times \mathcal{U})$ is "Bayes-consistent" with the prior $\mu \in \mathcal{U}$ if $E^{\pi(a,\mu')}[\mu'] = \mu$,[7] where $E^{\pi(a,\mu')}[\cdot]$ denotes the expectation under measure π, and if all $\mu' \in \text{supp}(\pi)$ are absolutely continuous with respect to μ, which we denote by $\mu' \ll \mu$.[8]

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 will think of $\Pi(\mu)$ as the set of feasible strategies.

Information Costs.   While all strategies in $\Pi(\mu)$ are feasible, some strategies will be 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 [2021], we assume costs take the following "posterior-separable" form.

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

$$C(\pi, \mu) = E^{\pi(a,\mu')}[D(\mu'||\mu)], \tag{3}$$

where $D : \{(\mu', \mu) \in \mathcal{U} \times \mathcal{U} : \mu' \ll \mu\} \to \mathbb{R}_+$ 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 $C(\pi, \mu)$.

The cost function in (3) 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 or 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 (3). Note that the cost of a strategy that gathers no information (all posteriors $\mu' \in \text{supp}(\pi)$ are equal to the prior μ) is zero.

---

[7]Throughout the paper, we use the expectation $E^{\pi(a,\mu')}[f(a,\mu')]$ to denote $\int_{supp(\pi)} f(a,\mu')d\pi(a,\mu')$. Note 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].

[8]Because we have assumed S and R are finite and that μ_0 has full support, absolute continuity here requires only that the posteriors μ' be consistent with the aggregate action strategy $\bar{\alpha}$. Even away from the finite S and R case, the requirement for absolute continuity is not very restrictive, as singular posterior measures relative to the prior must by Bayes-consistency be zero measure events under π.

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

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

where $V(a,\mu') \equiv E^{\mu'(s,r,\bar{a})}[u(a,\bar{a},s)]$ denotes the agent's expected utility from taking action a under posterior measure $\mu' \in \mathcal{U}$. Caplin, Dean, and Leahy [2021] call $V(a,\mu') - D(\mu'||\mu)$ the net utility function.

To ensure that the agent's problem is well-behaved, we assume the divergence is continuous and differentiable.

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 Fréchet differentiable on $(\mu'_0,\bar{\alpha}) \in int(\mathcal{U}_0 \times \bar{\mathcal{A}})$.*[9]

2.4 Equilibrium Definition

We first discuss how the aggregate action function $\bar{\alpha} \in \bar{\mathcal{A}}$ is determined, and then define an equilibrium of this game.

Let $\xi \equiv (\pi, \bar{\alpha})$ denote a symmetric strategy profile, consisting of a strategy $\pi \in \Delta(A \times \mathcal{U})$ 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 $(s,r) \in S \times R$,

$$\bar{\alpha}(s,r) = E^{\pi(a,\mu')}[a|s,r], \tag{4}$$

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

---

[9]Because $S \times R$ is finite, we could regard μ_0 and $\bar{\alpha}$ as vectors in $\mathbb{R}^{|S| \times |R|}$ and define differentiability in the usual way. Equivalently, consistent with the topologies we have assumed, observe that $int(\mathcal{U}_0 \times \bar{\mathcal{A}})$ is a subset of $ca(2^{S \times R}) \times C^b(S \times R)$, where $ca(2^{S \times R})$ is the space of signed measures of bounded variation on the power set of $S \times R$ endowed with the weak* topology and $C^b(S \times R)$ is the set of bounded continuous functions endowed with the sup-norm topology. $ca(2^{S \times R})$ and C^b are Banach spaces (see, e.g., Aliprantis and Border [2006] chapter 10), and hence Fréchet differentiability is well-defined.

Because we have assumed S and R are finite, defining this conditional expectation is straightforward.[10] 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. We model 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 a symmetric equilibrium in our game as a strategy profile in which the strategies π are individually optimal, the prior μ is consistent with the aggregate action function $\bar{\alpha} \in \bar{\mathcal{A}}$, and $\bar{\alpha}$ is consistent with the mean action generated by π.

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

$$\pi \in \arg \max_{\pi \in \Pi(\mu)} E^{\pi(a,\mu')}[V(a, \mu') - D(\mu'||\mu)]. \tag{5}$$

Our 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.[11] Our equilibrium concept builds on Denti [2020], 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 $S \times R \times \bar{A}$. 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 $\bar{\alpha} \in \bar{\mathcal{A}}$ with $\mu' = \phi\{\mu_0', \bar{\alpha}\}$ and $\mu = \phi\{\mu_0, \bar{\alpha}\}$. 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.

---

[10]With 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 $E^{\pi(a,\mu')}[a|s,r] = \int_{supp(\pi)} a f_0\{\mu', \mu\}(s,r) d\pi(a, \mu')$.

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

Definition 4. A strategy profile $(\pi_0, \bar{\alpha})$ is mean-consistent if $\bar{\alpha}(s, r) = E^{\pi_0(a,\mu'_0)}[a|s, r]$. A strategy profile $(\pi_0, \bar{\alpha})$ is an equilibrium if it is mean-consistent and

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

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 $(a, \mu'_0) \mapsto (a, \phi\{\mu'_0, \bar{\alpha}\})$.[12] Similarly, given some $\pi \in \Pi(\phi\{\mu_0, \bar{\alpha}\})$, we can construct the corresponding π_0.[13] In what follows, we will move back and forth between these spaces as convenient.

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.

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, 2021]. 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, 2021], 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 are general in the following sense: Hébert and Woodford [2019] show that all differentiable information cost functions can be locally approximated by a differentiable posterior-separable cost function. Fourth, posterior-separable cost functions can be studied via Lagrangian methods [Caplin, Dean, and Leahy, 2021], which facilitates certain proofs.

---

[12]Here and throughout the paper, we say a measure π is induced from a measure π_0 by some mapping when π is the push-forward measure of π_0 under that mapping.

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

That said, many of our results could 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 signals, $r \in R$. 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.

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.[14] 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 third departure: we incorporate learning of this type by allowing agents to pay attention to the endogenous aggregate action. In particular, we define the cost function in (3), and its associated divergence, on priors and posteriors over $S \times R \times \bar{A}$.[15]

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.

---

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

[15]Rational expectations are maintained by requiring that the prior over $S \times R \times \bar{A}$ is consistent with the aggregate action function (2) and that the aggregate action function is consistent with the mean (4).

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 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 its 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 any posterior constant. For example, interpreting $\bar{a}$ as a price, we might expect that larger price movements are less costly to observe. See Section 8.4 for a concrete example and Appendix Section A for a more in depth analysis on this point.

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 D, and properties of the equilibrium. In this section we begin by providing three examples of the (reduced-form) divergence D. We then introduce and define two properties of divergences that are instrumental for our later results; we call these properties "monotonicity in R" and "invariance in $\bar{A}$."

3.1 Examples of Posterior-Separable Cost Functions

Mutual Information.   Mutual information is posterior-separable: 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

$$D_{KL}(\mu'||\mu) = E^{\mu'(s,r,\bar{a})} \left[\ln \left(\frac{d\mu'(s,r,\bar{a})}{d\mu(s,r,\bar{a})} \right) \right] = \sum_{s \in S} \sum_{r \in R} \mu'_0(s,r) \ln \left(\frac{\mu'_0(s,r)}{\mu_0(s,r)} \right). \tag{6}$$

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, the Kullback-Leibler divergence does not depend on $\bar{\alpha}$.

Cost Functions with Perceptual Distance.   Several recent papers (e.g. Hébert and Woodford [2021], Pomatto et al. [2020], Walker-Jones [2019], Bloedel and Zhong [2020]) have described information costs that, unlike mutual information, incorporate a notion of "perceptual distance." We use as an example 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 $S = \{s_1, \ldots, s_{|S|}\} \subset \mathbb{R}$ with the s_i strictly increasing in i, and $R = \{r_1, \ldots, r_{|R|}\} \subset \mathbb{R}$ with r_j strictly increasing in j. The neighborhood-based cost function we consider 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

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

and

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

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.

It is also useful when D_N is viewed as a reduced form divergence arising from a choice over how best to acquire the information. Suppose as an example that the agent learns about s in part by observing $\bar a$. To simplify, assume that $|R| = 1$ so that we may ignore the second part of the D_N cost function. In this case, we can define, for $j = 1$,

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

as the perceptual distance between two states. We denote this particular divergence as $D_{S\bar A}$, to indicate that the agent can observe S and $\bar A$ but not R. The perceptual distance $d_{ij}^s(\bar\alpha)$ cap-

tures 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). This seems appealing if the agent pays attention to the aggregate action and extreme actions are easier to observe. In Appendix Section A 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 8.4.

We can also define versions of the neighborhood-based cost function for which the perceptual distances do not depend on the aggregate action function. 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 this particular divergence as D_{SR}. Like the KL divergence, D_{SR} does not depend on $\bar{\alpha}$. However, as we will soon show, D_{SR} and D_{KL} differ with respect to whether agents will attend to r.

Because S and R are finite, it is straightforward to show that the examples D_{KL}, $D_{S\bar{A}}$, and D_{SR} satisfy our continuity and differentiability assumptions.

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

Proof. See the appendix, E.2. $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\square$

3.2 Monotonicity and Invariance in R and $\bar{A}$

The goal of this section is to introduce two properties of divergences. These properties describe how a divergence might respond to certain transformations of its inputs: the prior and posterior. We begin by defining the particular transformations we have in mind.

First, let $\gamma_R : \Delta(S \times R \times \bar{A}) \to \Delta(S \times \bar{A})$ 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.[16] Next, define an operator $\eta_R : \Delta(S \times R \times \bar{A}) \times \Delta(S \times R \times \bar{A}) \to \Delta(S \times R \times \bar{A})$. 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 $\eta_R\{\mu_1, \mu_2\} \in \Delta(S \times R \times \bar{A})$, defined by

$$\frac{d\eta_R\{\mu_1, \mu_2\}}{d\mu_2} = \frac{d\gamma_R\{\mu_1\}}{d\gamma_R\{\mu_2\}}. \tag{7}$$

The operator η_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.[17] The measure $\eta_R\{\mu_1, \mu_2\}$ is similar to μ_1 in that they share a common marginal distribution on $S \times \bar{A}$, but matches μ_2 with respect to the conditional distribution of $r \in R$. In this sense, it is closer to μ_2 than μ_1 was originally.

---

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

[17]Equivalently, this transformation replaces the marginal distribution on $S \times \bar{A}$ of μ_2 with that of μ_1.

Example. In this example we illustrate the mechanics of the γ_R and η_R operators. For simplicity, we assume $\bar{\alpha}(s, r) = \bar{a}_0$ is constant. We also assume 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 $\mu_1, \mu_2 \in \mathcal{U}$.

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

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

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

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

	$\bar{a} = \bar{a}_0$
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	0.30	0.30
s_2	0.10	0.30

Table 4. $\eta_R\{\mu_1, \mu_2\} \in \mathcal{U}$

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, $d\gamma_R\{\mu_1\}/d\gamma_R\{\mu_2\}$, 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 $\eta_R\{\mu_1, \mu_2\} \in \mathcal{U}$ 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 η_R defined in (7) to a posterior and prior probability measure, μ' and μ, and consider how this affects the divergence. We will call a divergence D "monotone in R" if making the posterior more similar to the prior in the sense described by η_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 $\mu' \ll \mu$,

$$D(\mu'||\mu) \geq D(\eta_R\{\mu', \mu\}||\mu). \tag{8}$$

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

This property compares the divergence of 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 conditional distribution on r given $(s, \bar{a})$ of the posterior with that of the prior reduces their divergence, then we say that the divergence is monotone in R.

Invariance in R. We next define another concept, 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 $\gamma_R\{\mu\} \ll \gamma_R\{\mu''\}$,

$$D(\eta_R\{\mu', \mu\}||\mu) = D(\eta_R\{\mu', \mu''\}||\eta_R\{\mu, \mu''\}). \tag{9}$$

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

The above property compares the divergence of 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 over $r \in R$, the particulars of this conditional distribution would be immaterial for the divergence.

Invariance in R and monotonicity in R are different properties of divergences, and one property does not imply the other. Monotonicity requires that replacing the posterior's conditional distributions of $r \in R$ with that of the prior reduces their divergence. Invariance and monotonicity together require the divergence to shrink to the same value for *all possible* conditional distributions:

$$D(\mu'||\mu) \geq D(\eta_R\{\mu', \mu\}||\mu) = D(\eta_R\{\mu', \mu''\}||\eta_R\{\mu, \mu''\})$$

for all $\mu, \mu', \mu'' \in \mathcal{U}$ with $\mu' \ll \mu$ and $\gamma_R\{\mu\} \ll \gamma_R\{\mu''\}$.

Invariance in $\bar{A}$. We have thus far defined monotonicity and invariance with respect to only one dimension: $r \in R$. We can, however, define these concepts with respect to the other dimensions as well. We next define invariance in $\bar{A}$.

Let $\gamma_{\bar{A}} : \Delta(S \times R \times \bar{A}) \to \Delta(S \times R)$ 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 $\eta_{\bar{A}} : \Delta(S \times R \times \bar{A}) \times \Delta(S \times R \times \bar{A}) \to \Delta(S \times R \times \bar{A})$ 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 $\eta_{\bar{A}}\{\mu_1, \mu_2\} \in \Delta(S \times R \times \bar{A})$ defined by

$$\frac{d\eta_{\bar{A}}\{\mu_1, \mu_2\}}{d\mu_2} = \frac{d\gamma_{\bar{A}}\{\mu_1\}}{d\gamma_{\bar{A}}\{\mu_2\}}. \tag{10}$$

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 $\gamma_{\bar{A}}\{\mu\} \ll \gamma_{\bar{A}}\{\mu''\}$,

$$D(\eta_{\bar{A}}\{\mu', \mu\}||\mu) = D(\eta_{\bar{A}}\{\mu', \mu''\}||\eta_{\bar{A}}\{\mu, \mu''\}). \tag{11}$$

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 in fact always share a common conditional distribution over $\bar{a}$, determined by the $\bar{\alpha}$ function. As a result invariance in $\bar{A}$ determines whether or not information costs depend on the aggregate action function $\bar{\alpha}$.

We will show that monotonicity in R and invariance in $\bar{A}$ are key properties that will determine the properties of equilibria we study (non-fundamental volatility and efficiency).

3.3 Cost Functions, Monotonicity, and Invariance

We next describe how the key properties of R-monotonicity and $\bar{A}$-invariance for the divergences D determine related properties of the posterior-separable information costs C.

Lemma 2. *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 $\pi' \in \Pi(\mu)$ induced from π by $(a, \mu') \mapsto (a, \eta_R\{\mu', \mu\})$ 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 $\pi'' \in \Pi(\eta_{\bar{A}}\{\mu, \mu''\})$ induced from π by $(a, \mu') \mapsto (a, \eta_{\bar{A}}\{\mu', \mu''\})$ satisfy

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

Proof. See the Appendix, Section E.3. □

These properties of the cost function C can be viewed as definitions of R-monotonicity and $\bar{A}$-invariance in the context of cost functions as opposed to divergences. These results also suggest how the definitions of R-monotonicity and $\bar{A}$-invariance could be extended to all cost functions including those that are not necessarily posterior-separable.[18] Note that all of the cost functions we consider are Blackwell-monotone (by the convexity of D); as some of our examples below illustrate, Blackwell-monotonicity of the cost function does not by itself imply a divergence monotone in R or invariant in $\bar{A}$.

---

[18]Under 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.

3.4 Examples

Mutual Information. The KL divergence is both monotone in R and invariant in $\bar{A}$. Consider any $\mu = \phi\{\mu_0, \bar{\alpha}\}$ and $\mu' = \phi\{\mu'_0, \bar{\alpha}\}$. Let $g : S \times \bar{A} \to \mathbb{R}_+$ be a density satisfying

$$g(s, \bar{a}) = \frac{d\gamma_R\{\mu'\}(s, \bar{a})}{d\gamma_R\{\mu\}(s, \bar{a})}.$$

The Kullback-Leibler divergence in this context can be written as

$$D_{KL}(\mu'||\mu) = \sum_{s \in S} \sum_{\bar{a} \in \{\bar{a}' \in \bar{A}:g(s,\bar{a}')>0\}} \gamma_R\{\mu'\}(s, \bar{a}) \ln(g(s, \bar{a}))$$

$$+ \sum_{s \in S} \sum_{\bar{a} \in \{\bar{a}' \in \bar{A}:g(s,\bar{a}')>0\}} \sum_{r \in \{r' \in R:\bar{\alpha}(s,r')=\bar{a}\}} \gamma_R\{\mu'\}(s, \bar{a}) \frac{\mu'(s, r, \bar{a})}{\gamma_R\{\mu'\}(s, \bar{a})} \ln\left(\frac{\mu'(s, r, \bar{a})}{\mu(s, r, \bar{a})g(s, \bar{a})} \right).$$

Applying Jensen's inequality in the R dimension yields

$$D_{KL}(\mu'||\mu) \geq \sum_{s \in S} \sum_{\bar{a} \in \{\bar{a}' \in \bar{A}:g(s,\bar{a}')>0\}} \gamma_R\{\mu'\}(s, \bar{a}) \ln(g(s, \bar{a})) = D(\eta_R\{\mu', \mu\}||\mu),$$

which is the definition of monotonicity in R.

Invariance in $\bar{A}$ of mutual information follows immediately from the lack of dependence of the KL divergence on the aggregate action strategy $\bar{\alpha}$.

Neighborhood Cost function with D_{SR}. Consider the example neighborhood cost function defined above with associated divergence D_{SR}. This divergence, like the KL divergence, does not depend on $\bar{\alpha}$; it is therefore invariant in $\bar{A}$ by construction.[19] The divergence D_{SR} differs however from the KL divergence in the following way.

Lemma 3. *Assume $|S| > 1$ and $|R| > 1$. The divergence D_{SR} is not monotone in R.*

Proof. See the appendix, E.4. □

The intuition behind this result is as follows. Suppose that

$$\mu'_0(s, r) = \mu_0(s, r)g_1(s) \exp(\epsilon(r - \bar{r}(s))),$$

where g_1 is a strictly increasing function of s and $\bar{r}(s) = \epsilon^{-1} \ln(E^{\mu_0}[\exp(\epsilon r)|s])$, and suppose for simplicity that $\bar{\alpha}$ is constant across (s, r). We show in the proof that under these assumptions, if

---

[19]This 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.

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 for this class of posteriors.

To understand the intuition behind 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 in this case is attempting to differentiate high values of s from low values of s ($g_1(s)$ 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 (because it reduces the cost of learning via s) but a second-order cost (because 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.

Neighborhood Cost function with $D_{S\bar{A}}$. Consider finally the example neighborhood cost function defined above with associated divergence $D_{S\bar{A}}$. Monotonicity in R is trivially satisfied: we have assumed $|R| = 1$ when defining $D_{S\bar{A}}$. However, the perceptual distance between two states depends on the aggregate action function $\bar{\alpha}$. This divergence, therefore, is not invariant in $\bar{A}$.

Relation to the standard definition of invariance. The special nature of mutual information—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.[20] Invariant divergences have been described in the information geometry literature by Chentsov [1982] and Amari and Nagaoka [2007] and employed in economics by Hébert [2018] and Hébert and Woodford [2019]. Caplin, Dean, and Leahy [2021] 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 described 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 the example of mutual information demonstrates.

---

[20]Our 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.

4 Equilibrium Existence and Uniqueness

We proceed by establishing equilibrium existence.

Proposition 1. *A symmetric BNE exists.*

Proof. See the more general Proposition 5 below. □

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).

We next consider equilibrium uniqueness. Recall that the payoff parameter β controls the degree of strategic interaction due to payoffs in the linear-quadratic game: $\beta > 0$ indicates strategic complementarity in actions, and $\beta < 0$ indicates strategic substitutability. Our restriction that $\beta < 1$ guarantees a unique equilibrium under exogenous information.

With endogenous information acquisition, however, this restriction is not sufficient. 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 $\beta < 0$ and the divergence D is invariant in $\bar{A}$, then there is a unique aggregate action function $\bar{\alpha}$ common to all symmetric BNE.*

Proof. See the more general Proposition 10 below. □

Relative to the restriction that $\beta < 1$, Proposition 2 calls for two additional conditions. One is a stricter condition on payoffs: that actions be strategic substitutes ($\beta < 0$). The second is a condition on costs: that D be invariant in $\bar{A}$. Both conditions essentially rule out strategic complementarity in 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. Specifically, we have not ruled out the possibility that multiple strategies $\pi \in \Pi(\phi\{\mu_0, \bar{\alpha}\})$ are both best responses to and mean-consistent with $\bar{\alpha}$. However, for our purposes, this weaker form of uniqueness will suffice.

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 with the perceptual distance

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

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 endogenous information acquisition can still lead to multiplicity (in the spirit of 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 ($0 \leq \beta < 1$). 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 setting, section 8.3, 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 weak strategic substitutability in payoffs combined with invariance in $\bar{A}$ is sufficient to generate a form of uniqueness; this example illustrates the difficulty of strengthening the result.

We next consider whether these equilibria will exhibit non-fundamental volatility.

5 Non-Fundamental Volatility

Under what conditions does an equilibrium vary with the public signal realization $r \in R$ conditional on the fundamental state $s \in S$? Put another way, under what conditions does the realization of the noise $e \in \mathcal{E}$ affect equilibrium outcomes?

To answer this question, we 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 8. 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 strategy $\pi \in \Pi(\mu)$ is **s-measurable** if, for all $(a, \mu') \in supp(\pi)$, there exists a function $g : S \times \bar{A} \to \mathbb{R}_+$ such that $g = \frac{d\mu'}{d\mu}$. A symmetric BNE $(\pi, \bar{\alpha})$ is **s-measurable** if π and $\bar{\alpha}$ are both s-measurable.

When the aggregate action varies with $r \in R$ conditional on $s \in S$, it varies with the payoff-irrelevant states $\epsilon \in \mathcal{E}$. It thereby exhibits non-fundamental volatility. When instead $\bar{\alpha}$ is s-measurable, the aggregate action 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 s-measurable, i.e. that they do not condition on r.[21] In what follows, we will look for equilibria that are s-measurable in this stronger sense.

---

[21]It is possible for an equilibrium to have an s-measurable $\bar{\alpha}$ and yet for the strategy π to be non-s-measurable, 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. These cases are in a certain sense non-generic; see Section C of the online appendix.

S-measurable equilibria. Lemma 2 above can be understood as stating that if the divergence is monotone in R, then for any strategy, there is an s-measurable strategy that generates the same posteriors on $\Delta(S \times \bar{A})$ and is less costly. Because the agent does not care about r per se, only how it affects $\bar{a}$, monotonicity in R implies that the set of optimal strategies always includes an s-measurable strategy. Consequently, if other agents use strategies that generate an s-measurable $\bar{\alpha}$, one of an individual agent's best responses will involve a strategy that does not depend on r in any way (either directly, because the strategy is s-measurable, or via $\bar{a}$, because $\bar{\alpha}$ is s-measurable). The same fixed-point logic we employed earlier to prove equilibrium existence can be used to show that when D is monotone in R, an s-measurable equilibrium exists.

Proposition 3. *If the divergence D is monotone in R, then an s-measurable symmetric BNE of the game exists. If in addition D is invariant in $\bar{A}$ and $\beta < 0$, then all symmetric BNE feature an s-measurable $\bar{\alpha}$.*

Proof. See the more general Propositions 6 and 10 below. ☐

The logic above shows that if D is monotone in R, an equilibrium with an s-measurable $\bar{\alpha}$ exists. Proposition 2 shows that if D is invariant in $\bar{A}$ and $\beta < 0$, there is a unique $\bar{\alpha}$ function across all equilibria. Consequently, if all of these conditions hold, all equilibria share the same s-measurable $\bar{\alpha}$ function. Thus, the standard mutual information cost function will always lead, in settings with strategic substitutability, to equilibria with zero non-fundamental volatility.

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.[22] If in addition D is invariant in $\bar{A}$ and $\beta < 0$, then all equilibria exhibit zero non-fundamental volatility.

Necessity. In Section 7, we generalize Proposition 3 to a larger class of payoff functions. We furthermore show that a version of monotonicity in R is a *necessary* condition to guarantee the existence of s-measurable equilibria for all utility functions. That is, with cost functions such

---

[22]In 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.

as D_{SR}, non-fundamental volatility can arise in equilibrium. We provide an example with non-fundamental volatility in the linear-quadratic-Gaussian setting (Section 8.1).

The intuition for necessity is essentially the converse of the argument for sufficiency. If the cost function is not monotone in R, then s-measurable strategies are not the least costly. Therefore, the best response of an agent to an s-measurable aggregate action function $\bar{\alpha}$ is a strategy that conditions on $r \in R$. It follows that an s-measurable equilibrium cannot exist.

In this case, the exogenous states $r \in R$ resemble the "noisy public signals" featured in the exogenous information literature. When cost functions are non-monotone in R, agents find it cheaper to condition actions on r (Lemma 2). This incentive is what introduces correlated noise in agents' actions and generates non-fundamental volatility in equilibrium.

Multiple equilibria. When there are multiple equilibria with different aggregate action functions $\bar{\alpha}$, monotonicity in R guarantees one of these equilibrium has zero non-fundamental volatility, but is not enough to rule out other equilibria that depend on r. To see why, consider as an example the standard rational inattention cost function D_{KL}, modified so that the agents can observe the noise $e \in \mathcal{E}$ without cost. Such a cost function is monotone in R, and yet can exhibit multiple equilibria due to strategic complementarity (as discussed in Section 4). From these equilibria, we can construct sunspot equilibria that use $e \in \mathcal{E}$ as a sunspot to select between the various s-measurable equilibria. The resulting sunspot equilibria are not s-measurable, because they depend on r conditional on s (via their dependence on e).[23]

6 Efficiency

We next consider efficiency. To do so, we adapt the constrained efficiency concept of Angeletos and Pavan [2007] to environments with endogenous information acquisition.[24] 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 9. 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 **con-**

---

[23]This argument does not really depend on the agents being able to observe e freely. If e were costly to observe, agents would make mistakes and the "sunspot" equilibrium would not be a randomization over s-measurable equilibria; nevertheless, a non-s-measurable equilibrium might exist.

[24]See 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.

strained efficient if it solves

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

subject to mean-consistency, $\bar{\alpha}(s,r) = E^{\pi_0(a,\mu_0')}[a|s,r], \forall (s,r) \in S \times R$. 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 the effects of individual agents' actions on the aggregate action function $\bar{\alpha}$, and internalizes its effects on welfare, whereas individual agents do not.

The aggregate action function $\bar{\alpha}$ potentially affects welfare in two distinct ways. The first is through payoffs, $V(a, \phi\{\mu_0', \bar{\alpha}\})$. The second is through information costs, $D(\phi\{\mu_0', \bar{\alpha}\} || \phi\{\mu_0, \bar{\alpha}\})$. We discuss both channels below.

Payoff Externalities. Observe by the definition of V that

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

A payoff externality appears to exist in our setting via the aggregate action function $\bar{\alpha}$, which enters the utility function via the utility function's dependence on the aggregate action $\bar{a}$. That is, each agent, when choosing her own action, does not consider how her action affects the aggregate action and hence the utility of other agents. Pollution is a classic example of this kind of externality.

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

$$\frac{\partial}{\partial \bar{a}} E^{\pi_0(a,\mu_0')}[u(a, \bar{a}, s)|s,r]|_{\bar{a}=\bar{\alpha}(s,r)} = 2\beta E^{\pi_0(a,\mu_0')}[a - \bar{\alpha}(s,r)|s,r] = 0.$$

Thus, in equilibrium, small changes to the aggregate action function $\bar{\alpha}$ have no direct effect on the expected payoff of the agents, which is to say there is in fact no payoff externality in our setting. Angeletos and Pavan [2007] develop this result in the context of exogenous signals, and call this property "efficiency in the use of information."

The key property of the utility function u that generates this result is that the mean action $\bar{\alpha}(s,r) = E^{\pi_0(a,\mu_0')}[a|s,r]$ is a critical point of the function $h(\bar{a}) \equiv E^{\pi_0(a,\mu_0')}[u(a, \bar{a}, s)|s,r]$ under any probability measure π_0. In Section 7, we show how this "mean-critical" property can be defined and extended to a general class of utility functions, and discuss its interpretation.

26

Information Cost Externalities. Information costs can depend on $\bar{\alpha}$, as with for example the cost function $D_{S\bar{A}}$. However, when D is invariant in $\bar{A}$, as with our example cost functions D_{SR} and D_{KL}, this channel is shut down. With no externalities in payoffs and no effect of $\bar{\alpha}$ on information costs, the first-order conditions of the planner and of each agent coincide. This leads to the following result:

Proposition 4. *If the divergence D is invariant in $\bar{A}$, then a constrained efficient symmetric BNE exists. If in addition $\beta \leq 0$, all symmetric BNE are constrained efficient.*

Proof. See the more general Propositions 7 and 10 below. $\square$

To prove the 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 a symmetric BNE.

The argument above in fact shows something stronger: if D is invariant in $\bar{A}$, every equilibrium is a critical point of the planner's problem. That is, in every equilibrium, the first-order conditions of the planner's problem are satisfied. This does not mean every equilibrium is constrained efficient– some equilibria might be minima or saddle points of the planner's problem.

However, if the planner's problem is concave, then every critical point is a maximum, and as a result every equilibrium is efficient. A sufficient (but not necessary) condition for concavity in our setting is $\beta \leq 0$. By Proposition 2, when $\beta < 0$, the aggregate action function is the same in all equilibria, and as a result all equilibria are equivalent from the planner's perspective. When $\beta = 0$, we have not ruled out equilibria with different aggregate action functions, but all such equilibria must generate the same expected utility in the planner's problem.

Necessity and Interpretation. In Section 7, we generalize Proposition 4 to a larger class of payoff functions. We furthermore show that a local form of invariance in $\bar{A}$ is a *necessary* condition for the existence of a constrained efficient equilibrium.

When D is not invariant in $\bar{A}$, as in our example $D_{S\bar{A}}$, information costs are affected by $\bar{\alpha}$, at least for some parameter values. 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 welfare by dictating a strategy that puts relatively more weight on easily observable actions, effectively lowering information acquisition costs.[25] We construct an example illustrating this point in the linear-

---

[25]A 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.

quadratic-Gaussian setting (Section 8.2). In this example, more extreme actions are more easily observed, and as a result the planner would prefer that each agent choose a more extreme action given her beliefs than the agent would choose for herself.

This information cost 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. This possibility is also ruled out if agents can learn by observing endogenous actions, but learning via the actions of others and arriving at the same information without observing the actions of other 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.[26]

7 General Utility Functions

The linear-quadratic payoff function assumed in (1) exhibits three key properties that we have relied on thus far. First, it is continuous. Second, it exhibits the aforementioned mean-critical property that shuts down externalities in payoffs. Third, the degree of strategic substitutability or complementarity is summarized by the parameter β.

In this section we show how our previous results can be generalized to a broader class of continuous utility functions. We continue to assume that the action space $A \subset \mathbb{R}$ is convex and compact, and hence that $\bar{A} = A$.[27] We begin by defining the class of continuous utility functions.

Definition 10. A utility function $u : A \times \bar{A} \times S \to \mathbb{R}$ is continuous if it is continuous on $A \times \bar{A}$ for all $s \in S$.

With continuous utility functions, we prove equilibrium existence.

Proposition 5. *If the utility function is continuous, a symmetric BNE (Definition 3) exists.*

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

7.1 Non-Fundamental Volatility

We next provide a necessary and sufficient condition for the existence of an s-measurable equilibrium for the class of continuous utility functions.

Proposition 6. *A s-measurable symmetric BNE exists for all continuous utility functions if and only if, for all s-measurable $\bar{\alpha} \in \bar{A}$, the divergence D is monotone in R given $\mu = \phi\{\mu_0, \bar{\alpha}\}$.*

---

[26]This has been shown in some specific contexts. See e.g. Online Appendix A 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.

[27]In the technical appendix, Section B, which extends our existence, uniqueness, and sufficiency results to the continuous state case, we also allow for a multi-dimensional action space $A \subset \mathbb{R}^N$. Assuming a one-dimensional action space avoids certain technical complications in the proofs of Propositions 8 and 9.

Proof. See the online appendix, section E.7. ◻

Proposition 6 demonstrates that the R-monotonicity condition of Proposition 3 is stronger than necessary. R-monotonicity need not hold given *all* priors; it is sufficient for the divergence to be monotone in R given priors induced from the exogenous prior μ_0 by an s-measurable $\bar{\alpha}$ function. If this property holds, the agent's best response to any s-measurable $\bar{\alpha}$ is s-measurable, and by the same fixed point argument that proves existence, an s-measurable equilibrium exists.

When the cost function is not monotone in R on this set of priors, then for some s-measurable $\bar{\alpha}$ the agent's best response is a non-s-measurable strategy π. Consequently, this $\bar{\alpha}$ can never be part of an s-measurable equilibrium. To prove Proposition 6, we construct a continuous utility function such that, if an s-measurable equilibrium were to exist, it could exist only at this value of $\bar{\alpha}$; consequently, for this utility function, no s-measurable equilibrium exists. It follows by Proposition 5 that a non-s-measurable equilibrium exists in this case.

We conclude that monotonicity in R on s-measurable priors is necessary and sufficient to ensure the existence of an s-measurable symmetric BNE for all continuous utility functions.

7.2 Efficiency

As discussed in Section 6, the linear-quadratic function assumed in (1) exhibits the "mean-critical" property that shuts down payoff externalities. We next define this property on the larger class of continuous utility functions as follows.

Definition 11. For a given $\pi_0 \in \Pi_0$, a utility function $u : A \times \bar{A} \times S \to \mathbb{R}$ is **mean-critical given** π_0 if it is differentiable on $A \times \bar{A}$ for all $s \in S$ and if $E^{\pi_0(a,\mu_0')}[a|s,r]$ is a critical point of $h(\bar{a}) = E^{\pi_0(a,\mu_0')}[u(a,\bar{a},s)|s,r]$. A utility function is **mean-critical** if it is mean-critical given π_0 for all measures $\pi_0 \in \Delta(A \times \mathcal{U}_0)$.

We next generalize the existence result of Proposition 4 to this larger class of payoff functions.

Proposition 7. *A constrained-efficient symmetric BNE exists if the utility function u is mean-critical and the reduced form divergence D is invariant in $\bar{A}$.*

Proof. See the appendix, E.8. ◻

Mean-criticality of payoffs and invariance in $\bar{A}$ of information costs are sufficient for the existence of a constrained efficient equilibrium. Mean-criticality eliminates payoff externalities, while invariance in $\bar{A}$ eliminates externalities in information costs. Consequently, the solution to the planner's problem is an equilibrium.

Weaker versions of these properties are also necessary. Intuitively, if the invariance and mean-critical properties hold locally at every profile $(\pi_0, \bar{\alpha})$ that is a solution to the planner's problem, then every such profile will also be an equilibrium. Violations of invariance or of

the mean-critical property at profiles $(\pi_0, \bar{\alpha})$ that are never solutions to the planner's problem are therefore irrelevant for the question of constrained efficiency. For example, profiles $(\pi_0, \bar{\alpha})$ violating mean-consistency play no role in determining whether or not equilibria are efficient.[28]

Let us first define a local version of invariance in $\bar{A}$. Recall that we have assumed (in Assumption 1) that the divergence D is differentiable.

Definition 12. 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 Frechét derivative with respect to $\bar{\alpha}$.

Note that D is invariant in $\bar{A}$ if and only if C_0 is locally invariant everywhere.

We next show that if a divergence guarantees constrained efficiency for all mean-critical utility functions, C_0 must be locally invariant at every interior solution to the planner's problem. We use the notation $int(\mathcal{U}_0)$ to denote the interior of the simplex.

Proposition 8. *If a constrained-efficient symmetric BNE exists for all mean-critical utility functions, then C_0 is locally invariant in $\bar{A}$ for all $\pi_0 \in \Pi_0(\mu_0)$ and $\bar{\alpha} \in \bar{\mathcal{A}}$ such that, for some mean-critical utility function, $(\pi_0, \bar{\alpha})$ is a constrained efficient strategy profile and $supp(\pi_0) \subset A \times int(\mathcal{U}_0)$.*

Proof. See the appendix, E.10. □

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

Consider the example cost function $D_{S\bar{A}}$. It is locally invariant in $\bar{A}$ only if no information is acquired; otherwise, making actions more extreme reduces information costs. Consequently, with $D_{S\bar{A}}$, any equilibrium with a mean-critical utility function that features information gathering and an interior $\bar{\alpha}$ is constrained inefficient.

Our next result proves necessity of mean-criticality given a divergence invariant in $\bar{A}$.

Proposition 9. *If a constrained-efficient symmetric BNE exists for all divergences D invariant in $\bar{A}$, then u is mean-critical given all $\pi_0 \in \Pi_0(\mu_0)$ such that, for some divergence D invariant in $\bar{A}$ and $\bar{\alpha} \in \bar{\mathcal{A}}$, $(\pi_0, \bar{\alpha})$ is a constrained efficient strategy profile and $supp(\pi_0) \subset A \times int(\mathcal{U}_0)$.*

Proof. See the appendix, E.11. □

---

[28] Another example: if D is monotone in R, non-s-measurable $(\pi, \bar{\alpha})$ might never solve the planner's problem.

If for all divergences invariant in $\bar{A}$ a constrained efficient equilibrium exists, then payoffs must exhibit mean-criticality on the relevant part of the domain. However, if mean-criticality fails on the relevant part of the domain, then the marginal social benefit of actions at the planner's optimum will not coincide with the marginal private benefit. It follows that the solution to the planner's problem will not be an equilibrium.

7.3 The Mean-Critical Representation and Uniqueness

With the linear-quadratic payoff function of (1) we found that if actions are strategic substitutes ($\beta < 0$) and the divergence is invariant in $\bar{A}$, then there is a unique aggregate action function $\bar{\alpha}$ common to all symmetric BNE (Proposition 2). We also found that if $\beta \leq 0$ and the divergence is invariant in $\bar{A}$, then all equilibria are constrained efficient (Proposition 4).

We extend these results to a class of sufficiently well-behaved mean-critical utility functions. First, we build on results in Banerjee et al. [2005] to provide an explicit functional form for what we call "regular" mean-critical utility functions.

Definition 13. A utility function $u : A \times \bar{A} \times S \rightarrow \mathbb{R}$ 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 everywhere concave or everywhere convex in $\bar{A}$.

The linear-quadratic function in (1) 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).

Lemma 4. *A regular utility function u 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*

$$u(a, \bar{a}, s) = g(a, s) + G(\bar{a}, s) + (a - \bar{a})\frac{\partial}{\partial \bar{a}}G(\bar{a}, s). \tag{12}$$

Proof. See the technical appendix, Lemma 6, for a multi-dimensional version of this lemma ($A \subset \mathbb{R}^n$), and the technical appendix, E.15, for a proof. $\qquad\square$

The "if" part of the proof is straightforward. The first-order partial derivative of the payoff function in (12) with respect to $\bar{a}$ is given by[29]

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

---

[29]Note by regularity that G is twice-differentiable.

31

Taking expectations demonstrates this utility function is mean-critical. The "only if" part builds on the proofs in Banerjee et al. [2005]. The linear-quadratic utility function in (1) exhibits this functional form, with $G(\bar{a}; s) = \beta\bar{a}^2$ and $g(a; s) = -(1-\beta)(a-s)^2 - \beta a^2$.[30]

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 (12) and the multi-dimensional version described in the appendix have 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 (12). 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, which show that the pecuniary externalities present in market settings do not lead to inefficiency in the presence of full risk-sharing.[31]

With regular mean-critical utility functions, we can generalize our earlier results. Observe that $\frac{\partial}{\partial a}u(a, \bar{a}, s)$ is decreasing in $\bar{a}$ if and only if $\frac{\partial}{\partial \bar{a}}G(\bar{a}, s)$ is decreasing in $\bar{a}$, which is to say if G is concave. That is, G concave generates weak strategic substitutability (as with $\beta \leq 0$ earlier), and G strictly concave generates strict strategic substitutability (as with $\beta < 0$ earlier).

Proposition 10. *If the utility function is regular and mean-critical and the associated G is strictly concave in $\bar{a} \in \bar{A}$, then for any divergence D invariant in $\bar{A}$, there is a unique aggregate action function $\bar{\alpha}$ common to all symmetric BNE. If the utility function is regular and mean-critical and the associated G function is weakly concave in $\bar{a} \in \bar{A}$, then for any divergence D invariant in $\bar{A}$, all symmetric BNE are constrained efficient.*

Proof. See the appendix, E.12. □

The connection in our results between efficiency and uniqueness is likely a by-product of our approach to the proofs as opposed a more general result. That is, utility functions that are not mean-critical and exhibit strategic substitutability (appropriately defined) will likely lead to unique $\bar{\alpha}$ but constrained inefficient equilibria. We should also emphasize that strategic complementarity will not always lead to multiplicity– if the convexity of the G is not "too strong," the

---

[30]Equation (12) can be seen as a generalization of the Angeletos and Pavan [2007] conditions to non-linear-quadratic games, and the multi-dimensional version in our appendix as a further generalization. That is, the mean-critical property applies in games with exogenous information.

[31]In 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.

social planner's problem will still be strictly concave in $\bar{\alpha}$ and the equilibrium aggregate action function will be unique. In our proof of Proposition 10, we describe a condition (involving g and D) that is equivalent to strict concavity in $\bar{\alpha}$ for the planner's problem.

We demonstrated earlier that monotonicity in R on s-measurable priors is necessary to ensure an s-measurable symmetric BNE exists for all continuous utility functions. This result in fact holds on the class of mean-critical utility functions exhibiting strategic substitutability. 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. □

8   The Linear-Quadratic Gaussian Setting

In this section we give examples illustrating our general results adapted to the linear-quadratic-Gaussian (LQG) setting. Our goal is to provide examples that can understood using first order conditions in a familiar setting, as opposed to providing a second set of formal results.

We use the linear-quadratic utility function of (1), but now assume that A, S, and R are the real lines.[32] 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 $(\bar{\alpha}_s, \bar{\alpha}_r) \in \mathbb{R}^2$.

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].[33] 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

$$\omega^i = (s, r) \cdot \lambda_0 + \varepsilon^i, \qquad \text{with} \qquad \varepsilon^i \sim \mathcal{N}(0, \tau_0^{-1}), \tag{13}$$

[32]Assuming 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.

[33]An alternative approach, common in the literature (see, e.g., Colombo et al. [2014]), is to impose this Gaussian signals assumption, define cost functions on Gaussian signals, and study the linear quadratic Gaussian problem.

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 (1), the optimal action given any posterior $\mu_0' \in \mathcal{U}_0$ is

$$a^*(\mu_0') = E^{\mu_0'(s,r)}[(1 - \beta)s + \beta(\bar{\alpha}_s s + \bar{\alpha}_r r)] = E^{\mu_0'(s,r)}[(s, r) \cdot \psi(\bar{\alpha}_s, \bar{\alpha}_r)],$$

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

$$\psi(\bar{\alpha})' \equiv (1 - \beta + \beta\bar{\alpha}_s, \beta\bar{\alpha}_r). \tag{14}$$

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

$$(\bar{\alpha}_s, \bar{\alpha}_r)' = \frac{\tau_0 \psi(\bar{\alpha}_s, \bar{\alpha}_r)' \Sigma_0 \lambda_0}{1 + \tau_0 \lambda_0' \Sigma_0 \lambda_0} \lambda_0. \tag{15}$$

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 Gaussian signals of the form assumed in (13) 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].

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

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

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:

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

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:

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

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 τ_0; linearity in precision 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 $|\bar{\alpha}_s|$.

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 (13). 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

$$E^{\mu_0'(s,r)}[u(a^*(\mu_0'), \bar{\alpha}(s,r), s)] = -V^{\mu_0'(s,r)}[(s,r)\cdot\psi(\bar{\alpha}_s, \bar{\alpha}_r)] - \beta(1-\beta)E^{\mu_0'(s,r)}[(s-\bar{\alpha}_s s - \bar{\alpha}_r r)^2]$$

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

$$E^{\pi_0(a,\mu_0')}[u(a, \bar{a}, s)] = \frac{\tau_0}{1+\tau_0\lambda_0'\Sigma_0\lambda_0}(\psi(\bar{\alpha}_s, \bar{\alpha}_r)'\Sigma_0\lambda_0)^2 + f(\bar{\alpha}_s, \bar{\alpha}_r),$$

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 symmetric BNE if mean consistency (15) holds and

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

8.1 Non-Fundamental Volatility

We begin by considering 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 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 (detailed in the technical appendix, Section D), there exists a linear equilibrium featuring zero non-fundamental volatility: $\bar{\alpha}_r^* = 0$.[34]

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: $(\lambda_0^*)' \neq (\sigma_s^{-1}, 0)$.[35] 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

---

[34]In fact, it's relatively straightforward to show that all linear equilibria are s-measurable. However, linearity is key to the result– without linearity, best responses are not Gaussian, and linearity rules out the kind of equilibria discussed at the end of Section 5 (in regards to the interaction of multiplicity and non-fundamental volatility).

[35]To 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.

non-zero information acquisition, $\tau_0^* > 0$, there exists non-fundamental volatility: $\bar{\alpha}_r \neq 0$. 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.

What explains these differences? The cost function C_{MI} is associated with a divergence that is monotone in R, and C_{FI-SR} is not.[36] In our general setting, monotonicity in R is sufficient for the existence of an s-measurable equilibrium; our result with C_{MI} is an analog in the LQG setting. In our general setting, monotonicity in R is also in a sense necessary. With the cost function C_{FI-SR}, monotonicity in R is violated if any information is acquired, and consequently the equilibrium fails to be s-measurable in this case.

8.2 Constrained Efficiency

We next consider 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 $\lambda_0' = (1, 0)$. We maintain these assumptions throughout the remainder of this sub-section and the next two sub-sections.

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:

$$\bar{\alpha}_s = (1 - \beta + \beta\bar{\alpha}_s)\left(1 - \frac{1}{1 + \tau_0\sigma_s^2}\right). \tag{16}$$

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 (16). 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

---

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

of the payoff function and the invariance of mutual information with respect to $\bar{A}$.

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 $\bar{\alpha}_s^* = 1 - (1 + \tau_0^*\sigma_s^2)^{-1}$. 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 $\bar{\alpha}_s^* > 0$,

$$(\sigma_s^2)^2(1 - \bar{\alpha}_s^*)^2(1 + (\bar{\alpha}_s^*)^2) = \theta. \tag{17}$$

Consider now the 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 planner's first-order condition is given by

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

with equality if $\bar{\alpha}_s^{**} > 0$. This condition coincides with the equilibrium condition (17) only if no information is gathered, in which case $\bar{\alpha}_s^{**} = \bar{\alpha}_s^* = 0$. This occurs when θ is sufficiently high: $\theta \geq (\sigma_s^2)^2$.[37] When instead $\theta < (\sigma_s^2)^2$, all equilibria are inefficient: $\bar{\alpha}_s^{**} > \bar{\alpha}_s^*$.

In this example, the cost function $C_{FI-S\bar{A}}$ is non-invariant in $\bar{A}$; in fact, 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. This implies that more extreme aggregate actions are easier or "cheaper" for agents to observe than less extreme aggregate actions—an effect not internalized by individual agents. 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.

8.3 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

---

[37]This 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. Consequently, the solution to the planner's problem is always an equilibrium in this case provided that the utility function is mean-critical at the relevant point.

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

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

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

$$(1 - \beta + \beta\bar{\alpha}_s)^2 \sigma_s^2 = \theta(1 + \tau_0^* \sigma_s^2). \tag{18}$$

An equilibrium $(\tau_0^*, \bar{\alpha}_s^*)$ jointly satisfies (18) and mean-consistency (16). 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 (18) and (16); in this equilibrium no information is acquired.

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 (16) and (18), one in which $\bar{\alpha}_s^* = 0$ and another in which $\bar{\alpha}_s^* > 0$.[38]

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.[39] 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. This example therefore illustrates the reasoning behind the qualifiers of Proposition 4. 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: 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.

_____

[38] The 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.

[39] Our 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.

8.4 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 A, 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

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

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

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

In a linear equilibrium ($\bar{\alpha}(s,r) = \bar{\alpha}_s s$), the signal $\bar{\omega}$ is equivalent, from an informational perspective, to receiving the signal $\omega = ((1,\bar{\alpha}_s)\cdot\lambda)s + \bar{\varepsilon}^i$. 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 $\bar{\alpha}$, 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.

9 Conclusion

In this paper we have studied the relationship between information cost functions and properties of equilibria in large games with strategic interaction and rationally inattentive agents. We have demonstrated that there is a close connection between certain properties of information cost functions—monotonicity in R and invariance in $\bar{A}$—and whether or not an equilibrium is efficient and/or exhibits non-fundamental volatility.

Efficiency holds when information cost functions are invariant in endogenous actions and in the absence of payoff externalities. We characterize the class of utility functions that rule out payoff externalities; we call these mean-critical utility functions. Our result generalizes the characterization of Angeletos and Pavan [2007] to non-linear-quadratic games with endogenous information acquisition. Equilibrium non-fundamental volatility exists when information costs are non-monotone in payoff-irrelevant dimensions of the exogenous state space.

Our results indicate that the choice of an information cost function is an important determinant of equilibria properties. The standard rational inattention cost function, mutual information, leads to both efficiency and zero non-fundamental volatility. In contrast, alternative cost functions such as the neighborhood-based cost functions proposed by Hébert and Woodford [2021] can lead to non-fundamental volatility and 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. *mimeo*, 2021.

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. *Cornell University mimeo*, 2020.

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. Rational inattention when decisions take time. *Unpublished manuscript*, October 2019.

Benjamin Hébert and Michael Woodford. Neighborhood-based information costs. Technical report, National Bureau of Economic Research, 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 Matêjka, 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.

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. *Available at SSRN 2889861*, 2019.

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 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 $S \times R$.

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 $\nu : \Omega \times S \times R \times \bar{A} \to \mathbb{R}_+$ 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 $\bar{\nu}$, $\bar{f}\{\nu, \bar{\mu}\} : \Omega \to \mathbb{R}_+$, 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 14. A **cost function defined on signals** is a cost function $\bar{C} : \mathcal{V} \times \mathcal{U} \to [0, \infty]$.

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

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

where $\sigma : \Omega \to \Delta(A)$ 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 $\sigma : \Omega \to \Delta(A)$, defines via the product measure a measure on $A \times \Omega$. From this product measure, the mapping $(a, \omega) \mapsto (a, \hat{u}\{f^\omega\{\nu, \mu\}, \mu\})$ induces a strategy $\hat{\pi}\{\bar{\nu}, \nu, \mu\} \subset \Delta(A \times \mathcal{U}(\mu))$. By construction, this strategy satisfies Bayes-consistency: for all $(s, r, \bar{a}) \in S \times R \times \bar{A}$,

$$\int_{supp(\pi)} d\mu'(s, r, \bar{a}) d\hat{\pi}\{\bar{\nu}, \nu, \mu\}(a, \mu') = \int_{supp(\bar{\nu})} f^\omega\{\nu, \mu\}(s, r, \bar{a}) \bar{f}\{\nu, \mu\}(\omega) d\bar{\nu}(\omega) d\mu(s, r, \bar{a})$$

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

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

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

where $V(a, \mu') = E^{\mu'(s, r, \bar{a})} [u(a, \bar{a}, s)]$ as in the main text. Note here that if $\hat{\pi}\{\bar{\nu}, \nu, \mu\} = \pi$ in the weak* topology, then $E^{\hat{\pi}\{\bar{\nu}, \nu, \mu\}(a, \mu')} [V(a, \mu')] = E^{\pi(a, \mu')} [V(a, \mu')]$. 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 5. *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, G.9. $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\square$

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

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

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 $\pi = \Phi(\pi_0)$. 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,

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

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

$$C(\bar{\nu}, \nu, \mu) = \int_{supp(\bar{\nu})} \bar{D}(f^{\omega}\{\nu, \mu\}||\mu)\bar{f}\{\nu, \bar{\mu}\}(\omega)d\bar{\nu}(\omega), \tag{19}$$

where $\bar{D} : \{(f, \mu) \in \mathcal{F} \times \mathcal{U} : f \in \mathcal{F}(\mu)\} \to [0, \infty]$ 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

46

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 B).

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

$$d\bar{\phi}\{f_0, \bar{\alpha}\}(s, r, \bar{a}) = f_0(s, r)d\mu_0(s, r)d\delta_{\bar{\alpha}(s,r)}(\bar{a}), \tag{20}$$

which is the analog of (2) 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}, $D_{S\bar{A}}$, 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 (19)). The corresponding divergence, the KL divergence, can be defined on densities as

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

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 5 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

$$\bar{D}_{FI}(f||\mu) = \int_{S \times R \times \bar{A}} ||\nabla \ln [f(s,r,\bar{a})]||^2 f(s,r,\bar{a}) d\mu(s,r,\bar{a}), \tag{21}$$

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

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

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

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

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,

$$D_{FI-SR}(\phi\{\mu_0', \bar{\alpha}\}||\phi\{\mu_0, \bar{\alpha}\}) = D^{\mathcal{F}}_{FI-SR}(f_0||\mu_0; \bar{\alpha})$$

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 B 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 $f_0 : S \to (0,\infty)$. 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.[40]

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,

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

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

$$f_a(s,r,\bar{\alpha}(s,r)) = \frac{\bar{\alpha}_s(s,r)}{1 + \bar{\alpha}_s(s,r)^2} f_{0,s}(s), \tag{22}$$

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

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

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,

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

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

---

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

in the technical appendix, Section B 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

B   Results for Continuous States

In this section we describe how some of our results can be extended 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 5 and 10), our sufficient condition for zero non-fundamental volatility (the "if" part of Proposition 6), and our sufficient condition for efficiency (Proposition 7) to the continuous state case. Our necessity proofs (the "only if" part of Proposition 6, and Propositions 8 and 9) 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 $S \times R \to \mathbb{R}$ satisfying, for all $f_0 \in \mathcal{F}_0$,

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

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-

1

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

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

The agents problem is to solve

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

which is a version of the alternative formulation of the agent's problem (Definition (4)) for densities. Mean-consistency requires

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

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 $\bar{\pi}_0 \in \bar{\Pi}_0$.

A symmetric BNE 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

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

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 $(a, f_0) \in supp(\bar{\pi}_0)$, there exists a function $g : S \times \bar{A} \to \mathbb{R}_+$ such that $g(s, \bar{a}) d\phi\{\mu_0, \bar{\alpha}\}(s, r, \bar{a}) = d\bar{\phi}\{f_0, \bar{\alpha}\}(s, r, \bar{a})$ for all $(s, r, \bar{a}) \in S \times R \times \bar{A}$.

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 $\bar{\phi}\{f_0', \bar{\alpha}\} = \eta_R\{\bar{\phi}\{f_0, \bar{\alpha}\}, \phi\{\mu_0, \bar{\alpha}\}\}$.

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 η_R and $\eta_{\bar{A}}$ 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.

Lastly, we adapt our definition mean-critical preferences to the relevant strategy space. A utility function $u : A \times \bar{A} \times S \to \mathbb{R}$ is mean-critical if it is differentiable on $A \times \bar{A}$ for all $s \in S$ and if, for all strategies $\bar{\pi}_0 \in \Delta(A \times \mathcal{F}_0^+)$, $E^{\bar{\pi}_0(a, f_0)}[a|s, r]$ is a critical point of $h(\bar{a}) = E^{\bar{\pi}_0(a, f_0)}[u(a, \bar{a}, s)|s, r]$. We adapt our characterization of mean-critical preferences (Lemma 4) to the infinite-dimensional setting with a multi-dimensional action space A as follows.

Lemma 6. *A regular utility function u is mean-critical if and only if there exists a function $G : \bar{A} \times S \to \mathbb{R}$ and a function $g : A \times S \to \mathbb{R}$ such that*

$$u(a, \bar{a}, s) = g(a; s) + G(\bar{a}; s) + (a - \bar{a}) \cdot \nabla G(\bar{a}; s), \tag{23}$$

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

Proof. See the appendix, E.15. $\square$

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

Proposition 11. *In the continuous state setting, assuming $u : A \times \bar{A} \times S \to \mathbb{R}$ is a continuous function,*

1. *A symmetric BNE exists,*

2. *If the divergence D is monotone in R, an s-measurable symmetric BNE exists,*

3. *If u is mean-critical and the divergence D is invariant in $\bar{A}$, a constrained efficient symmetric BNE exists.*

4. *If the divergence D is invariant in $\bar{A}$ and u is mean-critical with an associated G function that is concave, then all symmetric BNE are constrained efficient. If in addition G is strictly concave, then in any pair of symmetric BNEs $(\bar{\pi}_{0,1}, \bar{\alpha}_1)$ and $(\bar{\pi}_{0,2}, \bar{\alpha}_2)$, $\bar{\alpha}_1 = \bar{\alpha}_2$.*

Proof. See the technical appendix, section F. $\square$

B.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 $f_0'(s) = E^{\mu_0(s,r)}[f_0(s,r)|r]$ 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,

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

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., Matêjka 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 $f_0'(s) = E^{\mu_0(s,r)}[f_0(s,r)|r]$ 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 $\epsilon^{-1} K^2$, and consequently 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

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

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

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.

We conclude with the following proposition, which shows that for both of these cost functions, the optimal posteriors are achievable given a sufficiently rich space of signal realizations. 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 7. *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*

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

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*

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

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

C   Non-S-Measurability and Generic Non-Fundamental Volatility

Let us suppose that, given some continuous utility function u and divergence D satisfying our assumptions on information costs, there exists an equilibrium $(\pi, \bar{\alpha})$ in which $\bar{\alpha}$ is s-measurable but π is not s-measurable (Definition 8).

If with probability one under π there exists a function $g : S \times \bar{A} \to \mathbb{R}_+$ such that $g \in \frac{d\mu'}{d\mu}$ (s-measurability is violated on a zero-measure set), then an s-measurable equilibrium also exists, as s-measurability is violated only in a trivial sense. Let us therefore consider the case in which s-measurability is violated on a positive-measure set.

Let us now define a sequence of functions $f_\epsilon : A \to A$ with the following properties: $f_\epsilon(a)$ is strictly increasing and differentiable in each dimension of $A \subset \mathbb{R}^L$, strictly convex in a for all $\epsilon > 0$, and satisfies $\lim_{\epsilon \to 0^+} f_\epsilon(a) = a$ for all $a \in A$. Observe that each of these functions is invertible. Given such a sequence of functions, define $\tilde{\pi}_\epsilon$ as the measure induced from π by $(a, \mu') \mapsto (f_\epsilon(a), \mu')$, and define the utility functions u_ϵ as satisfying

$$u_\epsilon(a', \bar{a}_\epsilon(s, r), s) = u(f_\epsilon^{-1}(a'), \bar{\alpha}(s, r), s) \, \forall s \in S, r \in R,$$

where $\bar{a}_\epsilon(s, r) = E^{\tilde{\pi}_\epsilon}[a | s, r] = E^{\pi}[f_\epsilon(a) | s, r]$. By construction, $(\bar{\alpha}_\epsilon, \tilde{\pi}_\epsilon)$ is an equilibrium of the game with the u_ϵ utility functions.

By assumption, π is not s-measurable, and therefore the higher moments of the distribution

of a given some s depend on r under π.[41] We can therefore find a sequence of functions such that $\bar{a}_\epsilon(s, r) \neq \bar{a}_\epsilon(s, r')$ for all $\epsilon > 0$ and $r \neq r'$.

We conclude that the case in which the equilibrium is not s-measurable but the economy nevertheless exhibits zero non-fundamental volatility is non-generic, in the sense that local perturbations to the utility function will generate non-fundamental volatility in this case.

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

$$a^*(\mu_0') = E^{\mu_0'(s,r)}[(s, r) \cdot \psi(\bar{\alpha}_s, \bar{\alpha}_r)].$$

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,

$$\Lambda^{-1} = \Sigma_0 - \frac{\tau_0 \Sigma_0 \lambda_0 \lambda_0' \Sigma_0}{1 + \tau_0 \lambda_0' \Sigma_0 \lambda_0}. \tag{24}$$

which yields

$$\mu_0'(\omega) = N(\omega \frac{\Sigma_0 \lambda_0}{\lambda_0' \Sigma_0 \lambda_0} \frac{\tau_0 \lambda_0' \Sigma_0 \lambda_0}{1 + \tau_0 \lambda_0' \Sigma_0 \lambda_0}, \Sigma_0 - \frac{\tau_0 \Sigma_0 \lambda_0 \lambda_0' \Sigma_0}{1 + \tau_0 \lambda_0' \Sigma_0 \lambda_0})$$

and therefore

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

---

[41]Because A is bounded, a moment-generating function must exist under each conditional distribution. If all moments do not depend on r, then the conditional distribution does not depend on r. If this holds for all (s, r), then π is s-measurable on a measure one set.

The linear-quadratic utility function in (12) can be written as

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

which yields

$$E^{\mu_0'(s,r)}[u(a^*(\mu_0'), \vec{\alpha} \cdot (s, r), s)] = -V^{\mu_0'(s,r)}[(s, r) \cdot \psi(\bar{\alpha}_s, \bar{\alpha}_r)] - E^{\mu_0'(s,r)}[\beta(1 - \beta)(s - \bar{\alpha}_s s - \bar{\alpha}_r r)^2]$$

where $V[\cdot]$ denotes the variance. We have by (24)

$$V^{\mu_0'(s,r)}[(s, r) \cdot \psi(\bar{\alpha}_s, \bar{\alpha}_r)] = \psi(\bar{\alpha}_s, \bar{\alpha}_r)' \Sigma_0 \psi(\bar{\alpha}_s, \bar{\alpha}_r) - \frac{\tau_0(\psi(\bar{\alpha}_s, \bar{\alpha}_r)' \Sigma_0 \lambda_0)^2}{1 + \tau_0 \lambda_0' \Sigma_0 \lambda_0}$$

The equation

$$E^{\pi_0(a,\mu_0')}[E^{\mu_0'(s,r)}[u(a, \bar{a}, s)]] = \frac{\tau_0(\psi(\bar{\alpha}_s, \bar{\alpha}_r)' \Sigma_0 \lambda_0)^2}{1 + \tau_0 \lambda_0' \Sigma_0 \lambda_0}$$
$$- \psi(\bar{\alpha}_s, \bar{\alpha}_r)' \Sigma_0 \psi(\bar{\alpha}_s, \bar{\alpha}_r) - \beta(1 - \beta)E^{\mu_0(s,r)}[(s - \bar{\alpha}_s s - \bar{\alpha}_r r)^2]$$

follows. Note that this derivation assumes Σ_0 is invertible. However, the resulting expression for expected utility is continuous in the limit as $\Sigma_0 \to \begin{bmatrix} \sigma_s^2 & 0 \\ 0 & 0 \end{bmatrix}$ 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

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

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,

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

where Λ is the posterior precision matrix, which is

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

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 Statement of Equilibrium Existence Results

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

Proposition 12. *(i) With the mutual information cost C_{MI}, there exists a linear equilibrium with zero non-fundamental volatility: $\bar{\alpha}_r = 0$. (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, E.14. □

D.4 Welfare in Equilibria with and without Information Acquisition

Assume $\theta = (1 - \beta)^2 \sigma_s^2$. Let us recall the first-order condition and mean-consistency condition:

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

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

We must have $\tau_0 \sigma_s^2 \geq 0$, and if this quantity is strictly positive, $sgn(\bar{\alpha}_s) = sgn(1 - \beta + \beta\bar{\alpha}_s)$. 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 either $\bar{\alpha}_s = 0$ and $\tau_0 = 0$ or

$$\bar{\alpha}_s = \frac{2\beta - 1}{\beta},$$

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

$$E^{\pi_0(a, \mu_0')}[E^{\mu_0'(s,r)}[u(a, \bar{a}, s)]] - C_{MI}(\lambda_0, \tau_0, \Sigma_0, \bar{\alpha}_s, \bar{\alpha}_r)$$
$$= \frac{\tau_0 \psi(\bar{\alpha}_s, \bar{\alpha}_r)' \Sigma_0 \lambda_0}{1 + \tau_0 \lambda_0' \Sigma_0 \lambda_0}(\psi(\bar{\alpha}_s, \bar{\alpha}_r)' \Sigma_0 \lambda_0)$$
$$-\psi(\bar{\alpha}_s, \bar{\alpha}_r)' \Sigma_0 \psi(\bar{\alpha}_s, \bar{\alpha}_r) - \beta(1 - \beta)E^{\mu_0(s,r)}[(s - \bar{\alpha}_s s - \bar{\alpha}_r r)^2]$$
$$-\theta \ln(1 + \tau_0 \lambda_0' \Sigma_0 \lambda_0).$$

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

$$J(\tau_0, \bar{\alpha}_s) = E^{\pi_0(a, \mu_0')}[E^{\mu_0'(s,r)}[u(a, \bar{a}, s)]] - C_{MI}(\lambda_0, \tau_0, \Sigma_0, \bar{\alpha}_s, \bar{\alpha}_r)$$

as

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

When $\bar{\alpha}_S = \tau_0 = 0$,

$$J(0, 0) = -(1 - \beta)\sigma_s^2.$$

9

In the other equilibrium,

$$J(\frac{2\beta - 1}{\beta}, \sigma_s^{-2}(\frac{\beta^2}{(1-\beta)^2} - 1)) = \beta^2\sigma_s^2\frac{\beta^2 - (1-\beta)^2}{\beta^2}$$
$$- \beta^2\sigma_s^2 - \frac{(1-\beta)}{\beta}(1-\beta)^2\sigma_s^2$$
$$- (1-\beta)^2\sigma_s^2 \ln(\frac{\beta^2}{(1-\beta)^2}).$$

which simplifies to

$$J(\frac{2\beta - 1}{\beta}, \sigma_s^{-2}(\frac{\beta^2}{(1-\beta)^2} - 1)) = -(1-\beta)^2\sigma_s^2(\beta^{-1} + \ln(\frac{\beta^2}{(1-\beta)^2})).$$

Defining $x = \frac{1-\beta}{\beta}$,

$$1 - \beta = 1 - \frac{1}{1+x} = \frac{x}{1+x},$$

$$f(x) = \frac{J(\frac{2\beta-1}{\beta}, \sigma_s^{-2}(\frac{\beta^2}{(1-\beta)^2} - 1)) - J(0,0)}{\beta(1-\beta)\sigma_s^2} = 1 - x^2 + 2x\ln(x).$$

For $x \in (0, 1]$, which corresponds to $\beta \in [\frac{1}{2}, 1)$, $f(1) = 0$ and

$$f'(x) = 2(1 - x + \ln(x)) \le 0.$$

Consequently,

$$J(\frac{2\beta - 1}{\beta}, \sigma_s^{-2}(\frac{\beta^2}{(1-\beta)^2} - 1)) \ge J(0,0).$$

E   Proofs

E.1   Additional Technical Lemmas

Lemma 8. *The set $\Pi_0(\mu_0)$ is non-empty, compact, and convex.*

Proof. See technical appendix, G.1. □

Lemma 9. *The function $\phi : \mathcal{U}_0 \times \bar{\mathcal{A}} \to \mathcal{U}$ is continuous.*

Proof. See technical appendix, G.2. □

Lemma 10. *The expected net utility function $J_0 : \Pi_0(\mu_0) \times \bar{\mathcal{A}} \to \mathbb{R}$ defined by $J_0(\pi_0, \bar{\alpha}) = E^{\pi_0(a, \mu'_0)}[V(a, \phi\{\mu'_0, \bar{\alpha}\}) - D(\phi\{\mu'_0, \bar{\alpha}\}||\phi\{\mu_0, \bar{\alpha}\})]$ is continuous, and concave in π.*

Proof. See technical appendix, G.3. □

Lemma 11. *The function $M : \Pi_0(\mu_0) \to \bar{\mathcal{A}}$ defined by $M\{\pi_0\}(s, r) = E^{\pi_0(a, \mu'_0)}[a|s, r]$ for all $(s, r) \in S \times R$ is continuous.*

Proof. See technical appendix, G.4. □

Lemma 12. *(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 10.* Π_0^* *is non-empty, compact-valued, convex-valued, and upper hemi-continuous.*

Proof. By Lemma 8, 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 10. Π_0^* is convex-valued by the linearity of J in π_0 and the convexity of $\Pi_0(\mu_0)$ (Lemma 8). □

E.2 Proof of Lemma 1

Differentiability on the interior follows immediately from the fact that all three functions are 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.

E.3 Proof of Lemma 2

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

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

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

$$D(\mu'||\mu) < D(\eta_R\{\mu'||\mu\}||\mu).$$

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

11

$S \times R$ and $\mu'_0 \ll \mu_0$, μ'' is non-negative for sufficiently small ϵ. By construction, this distribution satisfies Bayes' consistency. We have

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

and defining π'_ϵ as the measure induced from π_ϵ by $(a, \mu') \mapsto (a, \eta_R\{\mu', \mu\})$,

$$E^{\pi'_\epsilon(a,\mu')}[D(\mu'||\mu)] = \epsilon D(\eta_R\{\mu'||\mu\}||\mu) + (1-\epsilon)D(\eta_R\{\frac{\mu - \epsilon\mu'}{1-\epsilon}, \mu\}||\mu)$$

$$= \epsilon D(\eta_R\{\mu'||\mu\}||\mu) + (1-\epsilon)D(\frac{\mu - \epsilon\eta_R\{\mu', \mu\}}{1-\epsilon}||\mu),$$

where the last follows from the linearity of η_R in its first argument (see (7)). By assumption, $D(\phi\{\mu'_0, \bar{\alpha}\}||\phi\{\mu_0, \bar{\alpha}\})$ is Fréchet-differentiable in μ'_0, and by definition D is minimized when $\mu' = \mu$, implying that the derivative is zero at this point. We have

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

and likewise $\lim_{\epsilon \to 0^+} \frac{1-\epsilon}{\epsilon}D(\frac{\mu - \epsilon\eta_R\{\mu', \mu\}}{1-\epsilon}||\mu) = 0$. Therefore,

$$\lim_{\epsilon \to 0^+} \epsilon^{-1}(E^{\pi_\epsilon(a,\mu')}[D(\mu'||\mu)] - E^{\pi'_\epsilon(a,\mu')}[D(\mu'||\mu)]) =$$

$$D(\mu'||\mu) - D(\eta_R\{\mu'||\mu\}||\mu) < 0,$$

contradicting $E^{\pi'_\epsilon(a,\mu')}[D(\mu'||\mu)] \leq E^{\pi_\epsilon(a,\mu')}[D(\mu'||\mu)]$ 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 $\bar{A}$,

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

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 $\mu'' = \phi\{\mu_0, \bar{\alpha}'\}$, $\eta_{\bar{A}}\{\mu', \mu''\} = \phi\{\mu'_0, \bar{\alpha}'\}$. Consequently, if

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

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'' is non-negative for sufficiently small ϵ. By construction, this distribution satisfies Bayes' consistency. We have

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

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

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

$$= \epsilon D(\phi\{\mu_0', \bar\alpha'\}||\phi\{\mu_0, \bar\alpha'\}) + (1-\epsilon)D(\frac{\phi\{\mu_0, \bar\alpha'\} - \epsilon\phi\{\mu_0', \bar\alpha'\}}{1-\epsilon}||\phi\{\mu_0, \bar\alpha'\}),$$

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 Fréchet-differentiable in μ_0', and by definition D is minimized when $\mu' = \mu$, implying that the derivative is zero at this point. We have

$$\lim_{\epsilon \to 0^+} \frac{(1-\epsilon)}{\epsilon} D(\frac{\mu - \epsilon\mu'}{1-\epsilon}||\mu) = \frac{D(\mu + \frac{\epsilon}{1-\epsilon}(\mu' - \mu)||\mu) - D(\mu||\mu)}{\frac{\epsilon}{1-\epsilon}} = 0.$$

Therefore,

$$\lim_{\epsilon \to 0^+} \epsilon^{-1}(E^{\pi_\epsilon(a,\mu')}[D(\mu'||\phi\{\mu_0, \bar\alpha\})] - E^{\pi'_\epsilon(a,\mu')}[D(\mu'||\phi\{\mu_0, \bar\alpha'\})]) =$$
$$D(\phi\{\mu_0', \bar\alpha\}||\phi\{\mu_0, \bar\alpha\}) - D(\phi\{\mu_0', \bar\alpha\}||\phi\{\mu_0, \bar\alpha'\}) \neq 0,$$

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

E.4   Proof of Lemma 3

We prove non-monotonicity via a counter-example.

Assume an MLRP property: $\frac{\mu_0(s_i, r_j)}{\mu_0(s_{i-1}, r_j)}$ is strictly increasing in r_j. Consider the example structure described in the text:

$$\mu_0'(s, r; \epsilon) = \mu_0(s, r)g_1(s)\exp(\epsilon(r - \bar r(s))),$$

where g_1 is some strictly increasing function of s and $\bar r(s; \epsilon) = \epsilon^{-1}\ln(E^{\mu_0}[\exp(\epsilon r)|s])$. We also assume $\bar\alpha(s, r) = \bar a_0$ for some $\bar a_0 \in \bar A$ and all $(s, r) \in S \times R$. Consequently, for $\mu'(\epsilon) = \phi\{\mu_0'(\epsilon), \bar\alpha\}$ and $\mu = \phi\{\mu_0, \bar\alpha\}$,

$$d\mu'(s, r, \bar a; \epsilon) = d\mu(s, r, \bar a)g_1(s)\exp(\epsilon(r - \bar r(s; \epsilon))).$$

By construction,

$$d\gamma_R\{\mu'(\epsilon)\}(s, \bar a) = g_1(s)d\gamma_R\{\mu\}(s, \bar a).$$

Therefore,

$$d\eta_R\{\mu'(\epsilon), \mu\} = g_1(s)d\mu(s, r, \bar{a})$$
$$= d\phi\{\mu_0'', \bar{\alpha}\},$$

where $\mu_0''(s, r) = g_1(s)\mu_0(s, r)$. We can write this as $\eta_R\{\mu'(\epsilon), \mu\} = \mu'(0)$.

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

$$(\mu_0'(s, r_j; \epsilon) + \mu_0'(s, r_{j-1}; \epsilon))D_{KL,j,s}(\mu_0'||\mu_0) =$$
$$\sum_{r \in \{r_j, r_{j-1}\}} \mu_0'(s, r; \epsilon)[\ln(\exp(\epsilon r))) - \ln(\frac{\mu_0(s, r_j)\exp(\epsilon r_j) + \mu_0(s, r_{j-1})\exp(\epsilon r_{j-1}))}{\mu_0(s, r_j) + \mu_0(s, r_{j-1})})].$$

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

$$\frac{\partial}{\partial \epsilon}(\mu_0'(s, r_j; \epsilon) + \mu_0'(s, r_{j-1}; \epsilon))D_{KL,j,s}(\mu_0'(\epsilon)||\mu_0)|_{\epsilon=0} =$$
$$g_1(s) \sum_{r \in \{r_j, r_{j-1}\}} \frac{\mu_0(s, r)}{\mu_0(s, r_j) + \mu_0(s, r_{j-1})}[r - \frac{\mu_0(s, r_j)}{\mu_0(s, r_j) + \mu_0(s, r_{j-1})}r_j - \frac{\mu_0(s, r_{j-1})}{\mu_0(s, r_j) + \mu_0(s, r_{j-1})}r_{j-1}] = 0.$$

We also have

$$(\mu_0'(s_i, r_j; \epsilon) + \mu_0'(s_{i-1}, r; \epsilon))D_{KL,i,r}(\mu_0'||\mu_0) =$$
$$\sum_{s \in \{s_i, s_{i-1}\}} \mu_0(s, r)g_1(s)\frac{\exp(\epsilon r)}{E^{\mu_0}[\exp(\epsilon r')|s])}\{\ln(\frac{g(s)}{E^{\mu_0}[\exp(\epsilon r')|s])})$$
$$- \ln(\frac{\mu_0(s_i, r)\frac{g_1(s_i)}{E^{\mu_0}[\exp(\epsilon r')|s_i])} + \mu_0(s_{i-1}, r)\frac{g_1(s_{i-1})}{E^{\mu_0}[\exp(\epsilon r')|s_{i-1}])}}{\mu_0(s_i, r) + \mu_0(s_{i-1}, r)})\}.$$

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

$$\frac{\partial}{\partial \epsilon}(\mu_0'(s_i, r_j; \epsilon) + \mu_0'(s_{i-1}, r; \epsilon))D_{KL,i,r}(\mu_0'||\mu_0)|_{\epsilon=0} =$$
$$\sum_{s \in \{s_i, s_{i-1}\}} \mu_0(s, r)g_1(s)(r - E^{\mu_0}[r'|s])(\ln(g_1(s)) - \ln(\frac{\mu_0(s_i, r)g_1(s_i) + \mu_0(s_{i-1}, r)g_1(s_{i-1})}{\mu_0(s_i, r) + \mu_0(s_{i-1}, r)}))+$$
$$\sum_{s \in \{s_i, s_{i-1}\}} \mu_0(s, r)g_1(s)(-E^{\mu_0}[r'|s] + \frac{\mu_0(s_i, r)g_1(s_i)E^{\mu_0}[r'|s_i] + \mu_0(s_i, r)g_1(s_{i-1})E^{\mu_0}[r'|s_{i-1}]}{\mu_0(s_i, r)g_1(s_i) + \mu_0(s_{i-1}, r)g_1(s_{i-1})}).$$

The second line simplifies to zero.

It follows that

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

$$\sum_{i=2}^{|S|}\sum_{j=1}^{|R|}\sum_{s\in\{s_i,s_{i-1}\}} \mu_0(s,r)(r_j-E^{\mu_0}[r'|s])\frac{g_1(s)}{s_i-s_{i-1}}[\ln(g_1(s))-\ln(\frac{\mu_0(s_i,r_j)g_1(s_i)+\mu_0(s_{i-1},r_j)g_1(s_{i-1})}{\mu_0(s_i,r_j)+\mu_0(s_{i-1},r_j)})].$$

By the MLRP property,

$$\frac{\mu_0(s_i,r_j)}{\mu_0(s_{i-1},r_j)}$$

is strictly increasing in r_j, and consequently by $g_1(s_i) > g_1(s_{i-1})$,

$$\frac{g_1(s)}{s_i-s_{i-1}}[\ln(g_1(s))-\ln(\frac{\mu_0(s_i,r_j)g_1(s_i)+\mu_0(s_{i-1},r_j)g_1(s_{i-1})}{\mu_0(s_i,r_j)+\mu_0(s_{i-1},r_j)})]$$

is strictly decreasing in r_j. It follows by the covariance rule that

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

$$\epsilon\sum_{i=2}^{|S|}\sum_{s\in\{s_i,s_{i-1}\}} \mu_0(s)E^{\mu_0}[r-E^{\mu_0}[r'|s]|s]E^{\mu_0}[\frac{g_1(s)}{s_i-s_{i-1}}(\ln(g_1(s))-\ln(\frac{\mu_0(s_i,r)g_1(s_i)+\mu_0(s_{i-1},r)g_1(s_{i-1})}{\mu_0(s_i,r)+\mu_0(s_{i-1},r)}))|s]$$

which yields

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

implying that for sufficiently small $\epsilon > 0$,

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

E.5 Proof of Proposition 5

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

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

as in Lemma 10. We can invoke the theorem of the maximum, Lemma 12, 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 11, the function $M : \Pi_0(\mu_0) \to \bar{\mathcal{A}}$ 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^*, $M \circ \Pi_0^* : \Pi_0(\mu_0) \rightrightarrows \Pi_0(\mu_0)$, 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 a symmetric BNE under the alternative formulation (Definition 4).

Define π^* as the measure induced from π_0^* by $(a, \mu_0') \to (a, \phi\{\mu_0', \bar{\alpha}^*\})$. It follows that $(\pi^*, \bar{\alpha}^*)$ is a symmetric BNE under Definition 3.

E.6 Additional Lemmas

Lemma 13. *Let $\bar{\alpha} \in \bar{\mathcal{A}}$ 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 $\mu'' = \eta_R\{\mu'', \phi\{\mu_0, \bar{\alpha}\}\}$. 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 $a \in A$, $(a, \mu') \in supp(\pi_{\epsilon, \rho})$ for at most one $\mu' \in \mathcal{U}$.*

Proof. See the technical appendix, G.5. $\qquad\square$

Lemma 14. *Suppose the divergence D is not monotone in R on the set of s-measurable priors. Then there exists an s-measurable, mean-consistent $(\pi \in \Pi(\phi\{\mu_0, \bar{\alpha}\}), \bar{\alpha} \in \bar{\mathcal{A}})$ such that π has finite support, for each $a \in A$, $(a, \mu') \in supp(\pi)$ 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 $(a, \mu') \mapsto (a, \eta_R\{\mu', \mu\})$ exists with*

$$C(\hat{\pi}, \phi\{\mu_0, \bar{\alpha}\}) < C(\pi, \phi\{\mu_0, \bar{\alpha}\}).$$

Proof. See technical appendix, G.6. $\qquad\square$

Lemma 15. *Let D be a divergence monotone in R, and let $(\bar{\pi}, \bar{\alpha})$ be any s-measurable, mean-consistent strategy profile such that π has finite support and, for each $a \in A$, $(a, \mu') \in supp(\pi)$ for at most one $\mu' \in \mathcal{U}$. There exists, for any $\chi > 0$ and $r_0 \in R$, a mean-critical utility function*

$$u(a, \bar{a}, s) = g(a; s) - \frac{\chi}{2}|\bar{a} - \bar{\alpha}(s, r_0)|^2 - \chi(a - \bar{a}) \cdot (\bar{a} - \bar{\alpha}(s, r_0))$$

such that $(\pi, \bar{\alpha})$ is a symmetric BNE.

Proof. See the technical appendix, G.7. $\qquad\square$

E.7 Proof of Proposition 6

This proof refers several of the additional lemmas found in Technical Appendix Section E.1.

Sufficiency. First, we prove the "if" part of the proposition. This essentially reproduces our existence proof (Proposition 5) 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 $(a, \mu'_0) \in supp(\pi_0)$, there exists a function $g_{\mu'_0} : S \to \mathbb{R}$ such that $\frac{d\mu'_0(s,r)}{d\mu_0(s,r)} = g_{\mu'_0}(s)$ 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 s-measurable.

Define $\Pi_0^S(\mu_0) \subseteq \Pi_0(\mu_0)$ as the set of s-measurable strategies.

Lemma 16. *The set $\Pi_0^S(\mu_0)$ is non-empty, compact, and convex.*

Proof. See technical appendix, G.1. $\qquad\qquad\square$

Using this result and the continuity of the expected utility function $J_0(\pi_0, \bar{\alpha})$ defined in Lemma 10, we can invoke the theorem of the maximum. Define $\bar{\mathcal{A}}^S \subset \bar{\mathcal{A}}$ as the set of s-measurable aggregate action functions. Let $\Pi^* : \bar{\mathcal{A}}^S \Rightarrow \Pi^S(\mu_0)$ 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 hemi-continuous. By Lemma 11, the function $M : \Pi_0(\mu_0) \to \bar{\mathcal{A}}$ defined by $M\{\pi_0\}(s,r) = E^{\pi_0}[a|s,r]$ for all $(s,r) \in S \times R$ is continuous. Moreover, if $\pi_0 \in \Pi_0^S(\mu_0)$, then for all $s \in S$ and $r, r' \in R$,

$$E^{\pi_0(a,\mu'_0)}[a|s,r] = \int_{supp(\pi_0)} a g_{\mu'_0}(s) d\pi_0(a, \mu'_0)$$
$$= E^{\pi_0(a,\mu'_0)}[a|s,r']$$

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 $(a, \mu'_0) \to (a, \phi\{\mu'_0, \bar{\alpha}^*\})$ 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 $\mu' \ll \phi\{\mu_0, \bar{\alpha}^*\}$,

$$V(a, \mu') = V(a, \eta_R\{\mu', \mu\}).$$

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 $(a, \mu') \mapsto (a, \eta_R\{\mu', \phi\{\mu_0, \bar{\alpha}^*\}\})$ satisfies

$$E^{\hat{\pi}(a,\mu')}[V(a, \mu')] = E^{\pi(a,\mu')}[V(a, \mu')],$$

and by monotonicity in R for the prior $\phi\{\mu_0, \bar{\alpha}^*\}$, $E^{\hat{\pi}(a,\mu')}[D(\mu' || \phi\{\mu_0, \bar{\alpha}\})] \leq E^{\pi(a,\mu')}[D(\mu' || \phi\{\mu_0, \bar{\alpha}\})]$. 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 a symmetric BNE.

Necessity. Next, we prove the "only if" part of the proposition. Let $\mu = \phi\{\mu_0, \bar{\alpha}\}$, and suppose D is not monotone in R on for all s-measurable $\bar{\alpha}$. Then by Lemma 14, 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 s-measurable, and $\hat{\pi}$ is not (π is induced from $\hat{\pi}$ by $(a, \mu') \mapsto (a, \eta_R\{\mu', \mu\})$ and $\hat{\pi} \neq \pi$).

First, let us observe by continuity (Assumption 1) that if this inequality holds, it holds for some for some $\bar{\alpha} \in \bar{\mathcal{A}}$ whose image is contained in the relative interior of $\bar{A}$. Let us define, for all $\mu' \in \mathcal{U}$ with $\mu' \ll \mu$,

$$D_R(\mu' || \mu) = D(\eta_R\{\mu', \mu\} || \mu)$$

and observe by definition that D_R is an R-monotone divergence, and that by the s-measurability of π,

$$E^{\pi(a,\mu')}[D(\mu' || \mu)] = E^{\pi(a,\mu')}[D_R(\mu' || \mu)].$$

By Lemma 15, for the divergence D_R, there exists for any $\chi > 0$ a utility function u of the sort described in that lemma, with

$$u(a, \bar{a}, s) = g(a; s) - \frac{\chi}{2}|\bar{a} - \bar{\alpha}(s, r_0)|^2 - \chi(a - \bar{a}) \cdot (\bar{a} - \bar{\alpha}(s, r_0)).$$

By the s-measurability of $\bar{\alpha}$, for all $\mu' \in \mathcal{U}$ with $\mu' \ll \mu$,

$$V(a, \mu') = V(a, \eta_R\{\mu', \mu\}),$$

and consequently

$$E^{\hat{\pi}(a,\mu')}[V(a, \mu')] = E^{\pi(a,\mu')}[V(a, \mu')].$$

By $C(\hat{\pi}, \mu) < C(\pi, \mu)$, $(\pi, \bar{\alpha})$ cannot be an equilibrium with the divergence D. Moreover, because π is a best response under D_R to $\bar{\alpha}$, any other s-measurable $\pi' \in \Pi(\mu)$ must satisfy

$$E^{\pi'(a,\mu')}[V(a,\mu') - D(\mu'||\mu)] = E^{\pi'(a,\mu')}[V(a,\mu') - D_R(\mu'||\mu)]$$
$$\leq E^{\pi(a,\mu')}[V(a,\mu') - D_R(\mu'||\mu)].$$

Consequently,

$$E^{\pi'(a,\mu')}[V(a,\mu') - D(\mu'||\mu)] < E^{\hat{\pi}(a,\mu')}[V(a,\mu') - D(\mu'||\mu)],$$

and therefore no s-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

$$u(a, \tilde{\alpha}(s, r_0), s) = g(a; s) + \frac{\chi}{2}|\tilde{\alpha}(s, r_0) - \bar{\alpha}(s, r_0)|^2 - \chi(a - \bar{\alpha}(s, r_0)) \cdot (\tilde{\alpha}(s, r_0) - \bar{\alpha}(s, r_0)).$$

Thus, by mean-consistency,

$$E^{\tilde{\pi}(a,\mu')}[u(a, \tilde{\alpha}(s, r_0), s)|s, r] < E^{\tilde{\pi}(a,\mu')}[g(a; s)|s, r] - \frac{\chi}{2}\epsilon^2.$$

If we consider instead the policy $\pi^+ \in \Pi(\phi\{\mu_0, \tilde{\alpha}\})$ induced from π by $(a, \mu') \to (a, \eta_{\bar{A}}\{\mu', \phi\{\mu_0, \tilde{\alpha}\}\})$, we have by construction $E^{\pi^+(a,\mu')}[a|s, r] = E^{\pi(a,\mu')}[a|s, r] = \bar{\alpha}(s, r)$ and therefore

$$E^{\pi^+(a,\mu')}[u(a, \tilde{\alpha}(s, r_0), s)|s, r] = E^{\pi(a,\mu')}[g(a; s)|s, r] + \frac{\chi}{2}|\tilde{\alpha}(s, r_0) - \bar{\alpha}(s, r_0)|^2$$
$$> E^{\pi(a,\mu')}[g(a; s)|s, r] + \frac{\chi}{2}\epsilon^2.$$

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 $\lim_{n\to\infty} \chi_n = \infty$, 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, $\lim_{n\to\infty} \bar{\alpha}_n = \bar{\alpha}$. By the theorem of the maximum (Lemma 12), the optimal policy correspondence for the agents is upper hemi-continuous, and consequently $\lim_{n\to\infty} \pi_{0,n}$ must be a mean-consistent, s-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.8   Proof of Proposition 7

This proof refers several of the additional lemmas found in Technical Appendix Section E.1.

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 10 and $M\{\pi_0\}$ denotes the aggregate action function implied by π_0, as defined in Lemma 11. 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.

The set $\Pi_0(\mu_0)$ is convex and compact in the weak* topology, $M\{\pi_0\}$ is continuous, and $J(\pi_0, \bar{\alpha})$ is continuous (see Lemmas 8, 11, and 10). Consequently, a maximizer of the planner's problem exists.

Let $(\bar{\alpha}^+, \pi_0^+)$ be a maximum of the planner's problem. The function $J_0(\pi_0, M\{\pi_0\})$ is directionally differentiable in π_0 at π_0^+:

$$\lim_{\epsilon \to 0^+} \frac{1}{\epsilon}(J_0((1-\epsilon)\pi_0^+ + \epsilon\pi_0', M\{(1-\epsilon)\pi_0^+ + \epsilon\pi_0'\}) - J_0(\pi_0^+, M\{\pi_0^+\})) = J_0(\pi_0', M\{\pi_0^+\}) - J_0(\pi_0^+, M\{\pi_0^+\})$$

for all $\pi' \in \Pi_0(\mu_0)$. This result follows from the observation that

$$E^{\pi_0(a,\mu_0')}[V(a, \phi\{\mu_0', \bar{\alpha}\})]$$

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,

$$\lim_{\epsilon \to 0^+} \frac{1}{\epsilon}(M\{(1-\epsilon)\pi_0^+ + \epsilon\pi_0'\} - M\{\pi_0^+\}) = M(\pi_0') - M(\pi_0^+),$$

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

Because π_0^+ is maximal, we must have $J_0(\pi_0', M\{\pi_0^+\}) - J_0(\pi_0^+, M\{\pi_0^+\}) \leq 0$ for all $\pi_0' \in \Pi_0(\mu_0)$. Therefore, by $\bar{\alpha}^+ = M\{\pi_0^+\}$, π^+ is a mean-consistent best response in the agent's problem to $\bar{\alpha}^+$, and thus $(\bar{\alpha}^+, \pi^+)$ is a symmetric BNE.

E.9   Additional Lemma

Lemma 17. *Let* $(\pi_0^*, \bar{\alpha}^*) \in \Pi_0(\mu_0) \times \bar{A}$ *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* $\Omega_0 \subset supp(\pi_0^*)$ *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 function* $m_\epsilon : A \times \mathcal{U}_0 \to A \times \mathcal{U}_0$ *that is differentiable in* ϵ *and that satisfy, 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, G.8. □

E.10 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 17).

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 10 and $M\{\pi_0\}$ denotes the aggregate action function implied by π_0, as defined in Lemma 11. The agent's problem is $\max_{\pi_0 \in \Pi_0(\mu_0)} J_0(\pi_0, \bar{\alpha})$, taking $\bar{\alpha} \in \bar{\mathcal{A}}$ as given.

Suppose there exists a mean-critical utility function u such that $(\pi_0^*, \bar{\alpha}^*)$ is a constrained efficient strategy profile and an equilibrium, and for which

$$\nabla_{\bar{\alpha}} C_0(\pi_0^*, \mu_0, \bar{\alpha}^*) \neq \mathbf{0}, \tag{25}$$

Let us apply Lemma 17 defined above, and suppose that for some (s_0, r_0) satisfying

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

there exists a π_0^*-positive-measure set $\Omega_0 \subset supp(\pi_0^*)$ such that for all $(a, \mu_0') \in \Omega_0$, $a > \bar{\alpha}^*(s_0, r_0)$.

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 17, $m_\epsilon(a, \mu_0')$ 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 differentiability of D, this must equal

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

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 17 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 (25). 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 symmetric BNE satisfying the interiority assumptions, $\nabla_{\bar{\alpha}} C_0(\pi_0^*, \mu_0, \bar{\alpha}^*) = \vec{0}$.

E.11 Proof of Proposition 9

This proof refers several of the additional lemmas found in Technical Appendix Section E.1, and uses the perturbation lemma above (Lemma 17).

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 10 and $M\{\pi_0\}$ denotes the aggregate action function implied by π_0, as defined in Lemma 11. The agent's problem is $\max_{\pi_0 \in \Pi_0(\mu_0)} J_0(\pi_0, \bar{\alpha})$, taking $\bar{\alpha} \in \bar{\mathcal{A}}$ 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 symmetric BNE with $supp(\pi_0) \subset A \times int(\mathcal{U}_0)$ and, for some $(s_0, r_0) \in S \times R$,

$$\frac{\partial}{\partial \bar{a}} E^{\pi_0^*(a,\mu_0')}[u(a, \bar{a}, s)|s_0, r_0]|_{\bar{a}=\bar{\alpha}(s_0,r_0)} \neq 0.$$

Let us apply Lemma 17, and suppose that there exists a π_0^*-positive-measure set $\Omega_0 \subset supp(\pi_0^*)$ such that for all $(a, \mu_0') \in \Omega_0$, $a > \bar{\alpha}^*(s_0, r_0) = E^{\pi_0^*(a,\mu_0')}[a|s_0, r_0]$.

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 17, $m_\epsilon(a, \mu_0')$ 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

$$\frac{\partial}{\partial \epsilon} J(\pi_{0,\epsilon}, M\{\pi_{0,\epsilon}\})|_{\epsilon=0} = E^{\mu_0}[(\frac{\partial}{\partial \epsilon} M\{\pi_{0,\epsilon}\}(s,r)|_{\epsilon=0}) \cdot \frac{\partial}{\partial \bar{a}} E^{\pi_0^*(a,\mu_0')}[u(a, \bar{a}, s)|s, r]|_{\bar{a}=\bar{\alpha}(s,r)}]$$

$$= \mu_0(s_0, r_0) x \frac{\partial}{\partial \bar{a}} E^{\pi_0^*(a,\mu_0')}[u(a, \bar{a}, s)|s_0, r_0]|_{\bar{a}=\bar{\alpha}(s_0,r_0)},$$

where $x > 0$ is the constant defined in Lemma 17. 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 $\frac{\partial}{\partial \bar{a}} E^{\pi_0^*}[u(a, \bar{a}, s)|s_0, r_0]|_{\bar{a}=\bar{\alpha}(s_0,r_0)} \neq 0$, no set Ω_0 as described in Lemma 17 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

$$\frac{\partial}{\partial \bar{a}} E^{\mu_0(s,r,\bar{a})}[u(\bar{\alpha}(s_0, r_0), \bar{a}, s)]|_{\bar{a}=\bar{\alpha}(s_0,r_0)} \neq 0,$$

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 symmetric BNE satisfying the interiority assumptions, $\frac{\partial}{\partial \bar{a}} E^{\pi_0^*(a,\mu_0')}[u(a,\bar{a},s)|s_0,r_0]|_{\bar{a}=\bar{\alpha}(s_0,r_0)} = 0$.

E.12 Proof of Proposition 10

By the assumption of invariance,

$$D(\phi\{\mu_0',\bar{\alpha}\}||\phi\{\mu_0,\bar{\alpha}\}) = D(\phi\{\mu_0',\bar{\alpha}'\}||\phi\{\mu_0,\bar{\alpha}'\})$$

for any arbitrary $\bar{\alpha}' \in \bar{\mathcal{A}}$.

The planner's problem can be rewritten as $\max_{\pi_0 \in \Pi_0(\mu_0)} J_0(\pi_0, M(\pi_0))$, where

$$J(\bar{\alpha}, \pi_0) = E^{\pi_0(a,\mu_0')}[E^{\mu_0'}[u(a,\bar{\alpha}(s,r),s)] - D(\phi\{\mu_0',\bar{\alpha}'\}||\phi\{\mu_0,\bar{\alpha}'\})]$$

and $M(\pi_0)$ denotes the aggregate action function implied by π_0 (as in Lemma 11). Similarly, the agent's problem is $\max_{\pi_0 \in \Pi_0(\mu_0)} J_0(\pi_0, \bar{\alpha})$, taking $\bar{\alpha} \in \bar{\mathcal{A}}$ as given.

The set $\Pi_0(\mu_0)$ is convex and compact, $M(\pi_0)$ is continuous, and $J_0(\pi_0, \bar{\alpha})$ is continuous (see Lemmas 8, 10, and 11). Consequently, a maximizer of the planner's problem exists.

Plugging in the functional form of a mean-critical utility function (Lemma 4), for any mean-consistent $(\pi_0, \bar{\alpha})$,

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

If G is concave, J is concave over mean-consistent pairs: if $(\pi_{0,1}, \bar{\alpha}_1)$ and $(\pi_{0,2}, \bar{\alpha}_2)$ are mean-consistent pairs, then for $\lambda \in (0,1)$, $(\lambda\pi_{0,1} + (1-\lambda)\pi_{0,2}, \lambda\bar{\alpha}_1 + (1-\lambda)\bar{\alpha}_2) \in \Pi_0(\mu_0) \times \bar{\mathcal{A}}$ is mean-consistent and delivers weakly higher utility. If G is strictly concave, utility is strictly higher whenever $\bar{\alpha}_1 \neq \bar{\alpha}_2$.

We next show that, with a divergence D invariant in $\bar{A}$ and a mean-critical utility function, every symmetric BNE is a critical point of the planner's problem. Let $(\pi_0^*, \bar{\alpha}^*)$ be a symmetric BNE. By the definition of the agent's problem, π_0^* is a critical value (a maximum) of $J_0(\cdot, \bar{\alpha}^*) : \Pi_0(\mu_0) \to \mathbb{R}$. By the definition of mean-consistency and the mean-critical property of the utility function, for all $\bar{\alpha}'' \in \bar{\mathcal{A}}$,

$$\lim_{\epsilon \to 0^+} \frac{1}{\epsilon} J_0(\pi_0^*, \bar{\alpha}^* + \epsilon(\bar{\alpha}'' - \bar{\alpha}^*)) = 0$$

It follows that $(\pi_0^*, \bar{\alpha}^*)$ is a critical value of $J_0(\pi_0, \bar{\alpha})$, and hence that $(\pi_0^*, \bar{\alpha}^*)$ 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 symmetric BNE are constrained-efficient. If G is strictly concave, two constrained efficient equilibria with different aggregate action functions cannot exist, as otherwise a convex combination would generate strictly higher utility, contradicting the assumption of constrained efficiency.

Define

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

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

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

is strictly concave, and the solution to the planner's problem involves a unique $\bar{\alpha}$.

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 $\chi > 0$,

$$u(a,\bar{a},s) = g(a;s) - \frac{\chi}{2}|\bar{a} - \bar{\alpha}(s,r_0)|^2 - \chi(a - \bar{a}) \cdot (\bar{a} - \bar{\alpha}(s,r_0)).$$

Defining $G(\bar{a};s) = -\frac{\chi}{2}|\bar{a} - \bar{\alpha}(s,r_0)|^2$ yields

$$u(a,\bar{a},s) = g(a;s) + G(\bar{a};s) + (a - \bar{a}) \cdot \frac{\partial}{\partial\bar{a}}G(\bar{a},s),$$

and hence by Lemma 4 u is mean-critical. By Proposition 5, an equilibrium exists, and by the proof of Proposition 6, no s-measurable equilibrium exists.

By the strict concavity of G and Proposition E.12, all equilibria share a common aggregate action function and are constrained efficient.

E.14   Proof of Proposition 12

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 $\lambda_0'\Sigma_0\lambda = 1$. 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

$$\tau_0^* \in \arg\max_{\tau_0}(1 - \beta + \beta\bar{\alpha}_s)^2\sigma_s^2(1 - \frac{1}{1 + \tau_0}) + \theta\ln(\frac{1}{1 + \tau_0}),$$

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 $g : \mathbb{R}^2 \to \mathbb{R}^2$ 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 $K \subset R^2$ 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, $(\lambda_0^* = (\sigma_s^{-1}, 0)', \vec{\alpha}^* = (\bar{\alpha}_s^*, 0), \tau_0^*)$ is an s-measurable equilibrium.

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). Let us call $\hat{\lambda}_0, \hat{\tau}_0$ the optimal policies under the Hébert and Woodford [2021] normalization. Those authors show that if $\theta < |\Sigma_0\psi(\vec{\alpha})|^2$, agents will respond by gathering information, and otherwise agents will not gather information.

Consequently, if $\theta \geq ((1 - \beta)\sigma_s^2)^2$, an equilibrium with no information gathering exists. Suppose in what follows that $\theta < ((1 - \beta)\sigma_s^2)^2$. In this case, Hébert and Woodford [2021] show that

$$\hat{\lambda}_0 = (\theta\Sigma_0^{-1} + \hat{\tau}_0 I)^{-1}\psi(\vec{\alpha}),$$
$$\theta = |(\Sigma_0^{-1} + \theta^{-1}\hat{\tau}_0 I)^{-1}\psi(\vec{\alpha})|^2.$$

Observe that

$$\psi(\vec{\alpha})'\Sigma_0\hat{\lambda}_0 = \theta\hat{\lambda}_0'\hat{\lambda}_0 + \hat{\tau}_0\hat{\lambda}_0'\Sigma_0\hat{\lambda}_0 > 0,$$

which leads to the result in Hébert and Woodford [2021] that

$$E[(s,r)\cdot\psi(\vec{\alpha})|\omega] = \hat{\tau}_0\omega.$$

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

$$\hat{\lambda}_0 = \theta^{-1}(1-\beta)\Sigma_0(e_1 - \hat{\tau}_0\hat{\lambda}_0). \tag{26}$$

Using the normalization $\hat{\lambda}_0'\hat{\lambda}_0 = \theta^{-1}$ we can write

$$\hat{\tau}_0 = (1-\beta)\hat{\tau}_0\hat{\lambda}_0'\Sigma_0(e_1 - \hat{\tau}_0\hat{\lambda}_0).$$

Define

$$v = e_1 - \hat{\tau}_0\hat{\lambda}_0,$$

and we have

$$v = (I + \theta^{-1}\hat{\tau}_0(1-\beta)\Sigma_0)^{-1}e_1$$

and

$$\hat{\tau}_0 = (1-\beta)\hat{\tau}_0\hat{\lambda}_0'\Sigma_0 v.$$

Observe that $|v| \leq 1$ for all $\hat{\tau}_0 \geq 0$.

Define the function $g : \mathbb{R}^3 \to \mathbb{R}^3$ by

$$g(v, \hat{\tau}_0) = \begin{bmatrix} (I + \theta^{-1}\hat{\tau}_0(1-\beta)\Sigma_0)^{-1}e_1 \\ (1-\beta)\hat{\tau}_0\hat{\lambda}_0'\Sigma_0 v \end{bmatrix}.$$

Observe by assumption that $\theta^{\frac{1}{2}} < (1-\beta)\sigma_s^2 < ||\Sigma_0||$, where $||\Sigma_0||$ denotes the matrix operator norm.

Consequently, if $\hat{\tau}_0 \leq 1$, then

$$\hat{\tau}_0\hat{\lambda}_0'\Sigma_0 v < |\hat{\lambda}_0| \cdot ||\Sigma_0|| < 1.$$

It follows that the function g maps the unit sphere to the unit sphere, and by Brouwer's fixed point theorem a fixed point exists.

Let $(v^*, \hat{\tau}_0^*)$ be the fixed point. By $\theta^{\frac{1}{2}} < (1 - \beta)\sigma_s^2$, we must have $\hat{\tau}_0^* > 0$. We can recover $\hat{\lambda}_0^* = (\hat{\tau}_0^*)^{-1}(e_1 - v^*)$ and

$$(\vec{\alpha}^*)' = \hat{\tau}_0^* \hat{\lambda}_0^*.$$

Defining $x = (\hat{\lambda}_0^* \Sigma_0 \hat{\lambda}_0^*)^{-\frac{1}{2}}$, we can define $\lambda_0^* = x\hat{\lambda}_0$ and $\tau_0^* = x^{-2}\hat{\tau}_0^*$. Thus constructed, $(\tau_0^*, \lambda_0^*, \vec{\alpha}^*)$ is an equilibrium.

Observe, by the argument in the text or (26) 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.

E.15   Proof of Lemma 6

The "if" part is essentially immediate:

$$\nabla_{\bar{a}} u(a, \bar{a}, s) = (a - \bar{a}) \cdot H(\bar{a}; s),$$

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 ρ with full support on some $(a, s, r) \in A \times S \times R$. By the definition of a mean-critical utility function, we must have

$$\nabla_{\bar{a}} u(a, \bar{a}, s)|_{\bar{a}=a} = 0.$$

In case of concavity in $\bar{a}$, it follows that

$$a \in \arg\max_{\bar{a} \in \bar{A}} u(a, \bar{a}, s).$$

Define the function

$$f(a, \bar{a}, s) = u(a, a, s) - u(a, \bar{a}, s).$$

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 $\bar{a}$,

$$a \in \arg\min_{\bar{a} \in \bar{A}} u(a, \bar{a}, s),$$

and

$$f(a, \bar{a}, s) = u(a, \bar{a}, s) - u(a, a, s)$$

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 ρ consisting of two point masses, $\rho = \alpha\delta_{(a_1, s_1, r_1)} + (1 - \alpha)\delta_{(a_2, s_1, r_1)}$. Let $\kappa \in \Delta(A)$ be the conditional density given (s_1, r_1), $\kappa = \alpha\delta_{a_1} + (1-\alpha)\delta_{a_2}$. By mean-consistency,

$$\int_A a d\kappa(a) \in \arg\min_{\bar{a} \in A} \int_A f(a, \bar{a}; s) d\kappa(a).$$

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 $F : A \times S \to \mathbb{R}$ convex in A for all $s \in S$. In the concave in $\bar{a}$ case, defining

$$g(a; s) = u(a, a, s) + F(a; s)$$

and $G(\bar{a}; s) = F(\bar{a}; s)$ yields

$$\begin{aligned} u(a, \bar{a}, s) &= u(a, a, s) - f(a, \bar{a}; s) \\ &= g(a; s) + G(\bar{a}; s) + (a - \bar{a}) \cdot \nabla G(\bar{a}; s). \end{aligned}$$

In the convex in $\bar{a}$ case, defining $g(a; s) = u(a, a, s) - G(a; s)$ and $G(\bar{a}; s) = -F(\bar{a}; s)$ yields

$$\begin{aligned} u(a, \bar{a}, s) &= u(a, a, s) + f(a, \bar{a}; s) \\ &= g(a; s) + G(\bar{a}; s) - (a - \bar{a}) \cdot \nabla G(\bar{a}; s). \end{aligned}$$

F Proof for the Continuous State Case

This section provides proofs for the claims in Proposition 11. The proofs are essentially notational adaptions of the proofs of Propositions 5, 6, 7, and 10 in the main text.

F.1 Proof of Existence and Sufficiency of 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 18. *The function $\phi_F : \mathcal{F}_0 \to \mathcal{U}_0$ defined by $d\phi_F\{f_0\}(s, r) = f_0(s, r) d\mu_0(s, r)$ is continuous.*

Proof. If $f_{0,n} \to f_0$ in the sup-norm topology, then for all continuous $g : S \times R \to \mathbb{R}$,

$$\int_{S \times R} g(s, r) f_{0,n}(s, r) d\mu_0(s, r) \to \int_{S \times R} g(s, r) f_0(s, r) d\mu_0(s, r)$$

by Aliprantis and Border [2006] 15.7 and the continuity of the product of continuous functions. □

It is convenient to define a strategy $\bar{\pi}_0 \in \bar{\Pi}_0$ as s-measurable if, for all $(a, f_0) \in supp(\bar{\pi}_0)$, $f_0(s, r) = f_0(s, r')$ 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 s-measurable.

Define $\bar{\Pi}_0^S \subseteq \bar{\Pi}_0$ as the set of s-measurable strategies.

Lemma 19. *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, G.11 $\hfill\square$

Define the objective function

$$\bar{J}(\bar{\pi}_0, \bar{\alpha}) = E^{\bar{\pi}_0(a, f_0)}[V(a, \phi\{\phi_F\{f_0\}, \bar{\alpha}\}) - D^{\mathcal{F}}(f_0 || \mu_0; \bar{\alpha})]. \tag{27}$$

The following lemma shows that this function is continuous.

Lemma 20. *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, the function $V(a, \phi\{\phi_F\{f_0\}, \bar{\alpha}\}) = E^{\phi\{\phi_F\{f_0\}, \bar{\alpha}\}(s, r, \bar{a})}[u(a, \bar{a}, s)]$ 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 10 for details). $\hfill\square$

We next invoke the theorem of the maximum. Let $\bar{\Pi}^* : \bar{\mathcal{A}} \rightrightarrows \bar{\Pi}_0$ 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 $\bar{\Pi}_S^* : \bar{\mathcal{A}}^S \rightrightarrows \bar{\Pi}_0^S$ 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 21. *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 $\bar{\Pi}_0^S$. $\hfill\square$

Lemma 22. *The function $M : \bar{\Pi}_0 \to \bar{\mathcal{A}}^+$ defined by $M\{\bar{\pi}_0\}(s, r) = \int_{supp(\bar{\pi}_0)} a f_0(s, r) d\bar{\pi}_0(a, f_0)$ 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, G.12. $\hfill\square$

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 $\bar{\alpha}^* = M\{\pi_0^*\}$, $\bar{\pi}_0^*$ 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 a symmetric BNE, and hence that a symmetric BNE 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 π_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 $\mu' \ll \phi\{\mu_0, \bar{\alpha}^S\}$,

$$\hat{u}(a, \mu') = \hat{u}(a, \eta_R\{\mu', \mu\}).$$

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 $(a, \mu') \mapsto (a, \eta_R\{\mu', \phi\{\mu_0, \bar{\alpha}^S\}\})$ satisfies

$$E^{\hat{\pi}(a,\mu')}[\hat{u}(a, \mu')] = E^{\pi(a,\mu')}[\hat{u}(a, \mu')],$$

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 $\pi' \in \Pi(\phi\{\mu, \bar{\alpha}^S\})$ via this mapping, where π' is induced from π by the mapping $(a, \mu') \mapsto (a, \eta_R\{\mu', \phi\{\mu_0, \bar{\alpha}^S\})$. 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 a symmetric BNE.

F.2 Proof of Sufficiency of Invariance in $\bar{A}$

The planner's problem can be rewritten as $\max_{\bar{\pi}_0 \in \bar{\Pi}_0(\mu_0)} \bar{J}(M\{\bar{\pi}_0\}, \bar{\pi}_0)$, where $\bar{J}$ is defined in (27) above and $M(\bar{\pi}_0)$ is defined as in Lemma 22. The agent's problem is $\max_{\bar{\pi}_0 \in \bar{\Pi}_0(\mu_0)} \bar{J}(\bar{\alpha}, \bar{\pi}_0)$, 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 $\bar{J}(\bar{\alpha}, \bar{\pi}_0)$ is continuous (see Lemmas 19, 20, and 22 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 $\bar{J}(M\{\bar{\pi}_0\}, \bar{\pi}_0)$ is directionally differentiable in $\bar{\pi}_0$ at $\bar{\pi}_0^+$:

$$\lim_{\epsilon \to 0^+} \frac{1}{\epsilon}(\bar{J}(M\{(1-\epsilon)\bar{\pi}_0^+ + \epsilon\bar{\pi}_0'\}, (1-\epsilon)\bar{\pi}_0^+ + \epsilon\bar{\pi}_0') - \bar{J}(M\{\bar{\pi}_0^+\}, \bar{\pi}_0^+)) = \bar{J}(M\{\bar{\pi}_0^+\}, \bar{\pi}_0') - \bar{J}(M\{\bar{\pi}_0^+\}, \pi_0^+)$$

for all $\bar{\pi}_0' \in \bar{\Pi}_0$. This result follows from the observation that

$$E^{\bar{\pi}_0(a,f_0)}[V(a, \phi\{\phi_F\{f_0\}, \bar{\alpha}\})]$$

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,

$$\lim_{\epsilon \to 0^+} \frac{1}{\epsilon}(M\{(1-\epsilon)\bar{\pi}_0^+ + \epsilon\bar{\pi}_0'\} - M\{\bar{\pi}_0^+\}) = M\{\bar{\pi}_0'\} - M\{\bar{\pi}_0^+\},$$

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

Because $\bar{\pi}_0^+$ is maximal, we must have $\bar{J}(M\{\bar{\pi}_0^+\}, \bar{\pi}_0') - \bar{J}(M\{\bar{\pi}_0^+\}, \pi_0^+) \leq 0$ 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 $(\bar{\alpha}^+, \pi_0^+)$ is a symmetric BNE.

F.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 (27) above and $M\{\bar{\pi}_0\}$ is defined as in Lemma 22. 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 19, 20, and 22 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 (Lemma 4), for any mean-consistent $(\bar{\pi}_0, \bar{\alpha})$,

$$\bar{J}(\bar{\pi}_0, \bar{\alpha}) = E^{\bar{\pi}_0(a,f_0)}[g(a; s) - D(\bar{\phi}\{f_0, \bar{\alpha}'\}||\phi\{\mu_0, \bar{\alpha}'\})] + E^{\mu_0}[G(\bar{\alpha}(s, r); s)].$$

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 symmetric BNE is a critical point of the planner's problem. Let $(\bar{\pi}_0^*, \bar{\alpha}^*)$ be a symmetric BNE. By the definition of the agent's problem, $\bar{\pi}_0^*$ is a critical value (a maximum) of $\bar{J}(\cdot, \bar{\alpha}^*) : \bar{\Pi}_0(\mu_0) \to \mathbb{R}$. By the definition of mean-consistency and the mean-critical property of the utility

function, for all $\bar{\alpha}'' \in \bar{\mathcal{A}}^+$,

$$\lim_{\epsilon \to 0^+} \frac{1}{\epsilon} \bar{J}(\bar{\pi}_0^*, \bar{\alpha}^* + \epsilon(\bar{\alpha}'' - \bar{\alpha}^*)) = 0$$

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 symmetric BNE 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).

G   Additional Proofs

G.1   Proof of Lemma 8

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 $\mu_0 = \int_{supp(\pi_n)} \mu_0' d\pi_n(a, \mu_0') \to \int_{supp(\pi_0)} \mu_0' d\pi_0(a, \mu_0')$, 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 $\pi_1, \pi_2 \in \Pi_0(\mu_0)$, $\int_{supp(\pi_1)} \mu_0' d\pi_1(a, \mu_0') = \mu_0$ and $\int_{supp(\pi_2)} \mu_0' d\pi_2(a, \mu_0') = \mu_0$, 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 $\alpha \in (0, 1)$.

G.2 Proof of Lemma 9

Let $g : S \times R \times \bar{A} \to \mathbb{R}$ be a continuous and bounded function. Then for any $(\bar{\alpha}_n, \mu_{0,n}) \in \bar{\mathcal{A}} \times \mathcal{U}_0$,

$$\int_{S \times R \times \bar{A}} g(s, r, \bar{a}) d\phi\{\mu_{0,n}, \bar{\alpha}_n\}(s, r, \bar{a}) = \int_{S \times R} g(s, r, \bar{\alpha}_n(s, r)) d\mu_{0,n}(s, r).$$

Suppose $(\bar{\alpha}_n, \mu_{0,n}) \to (\bar{\alpha}, \mu_0')$. 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 $|\bar{a} - \bar{a}'| < \delta$,

$$|g(s, r, \bar{a}) - g(s, r, \bar{a}')| < \epsilon.$$

Consequently, for any $\epsilon > 0$, there exists an n_0 such that

$$\sup_{(s,r) \in S \times R} |g(s, r, \bar{\alpha}(s, r)) - g(s, r, \bar{\alpha}_n(s, r))| < \epsilon,$$

which is to say $g(s, r, \bar{\alpha}_n(s, r)) \to g(s, r, \bar{\alpha}(s, r))$ in the sup norm topology.

It follows by Aliprantis and Border [2006] 15.7 that

$$\int_{S \times R} g(s, r, \bar{\alpha}_n(s, r)) d\mu_{0,n}(s, r) \to \int_{S \times R} g(s, r, \bar{\alpha}(s, r)) d\mu_0'(s, r),$$

and hence that

$$\int_{S \times R \times \bar{A}} g(s, r, \bar{a}) d\phi\{\mu_{0,n}, \bar{\alpha}_n\}(s, r, \bar{a}) \to \int_{S \times R \times \bar{A}} g(s, r, \bar{a}) d\phi\{\mu_0', \bar{\alpha}\}(s, r, \bar{a}).$$

G.3 Proof of Lemma 10

Concavity is immediate by linearity.

The function

$$V(a, \mu') = E^{\mu'(s, r, \bar{a})}[u(a, \bar{a}, s)]$$

is continuous on $A \times \mathcal{U}$ by the continuity of u 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 9, 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 $(\bar{\alpha}_n, \pi_{0,n}) \to (\bar{\alpha}, \pi_0)$. 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, $h(\cdot, \bar{\alpha}_n(s,r)) \to h(\cdot, \bar{\alpha}(s,r))$ in the sup norm topology.

It follows by Aliprantis and Border [2006] 15.7 that

$$E^{\pi_{0,n}(a,\mu_0')}[h(a, \mu_0', \bar{\alpha}_n)] \to E^{\pi_0(a,\mu_0')}[h(a, \mu_0', \bar{\alpha})]$$

and hence that $J(\cdot)$ is continuous.

G.4   Proof of Lemma 11

Define the function $g_{s,r} : A \times \mathcal{U}_0 \to \mathbb{R}$ by

$$g_{s,r}\{a, \mu_0'\} = a \frac{\mu_0'(s,r)}{\mu_0(s,r)}.$$

By the finiteness of $S \times R$, $\mu_{0,n}' \to \mu_0'$ implies point-wise convergence. Consequently, if $(a_n, \mu_{0,n}') \to (a, \mu_0')$, then $g_{s,r}(a_n, \mu_{0,n}') \to g_{s,r}(a, \mu_0')$ for all $(s,r) \in S \times R$. Thus, $g_{s,r}$ 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 $\pi_{0,n} \to \pi_0$, then $M\{\pi_{0,n}\}(s,r) \to M\{\pi_0\}(s,r)$ 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.

G.5   Proof of Lemma 13

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

$$a_0 = E^{\mu_0(s,r)}[\bar{\alpha}(s,r)].$$

Define, for some $\lambda > 1$, $\bar{A}_\lambda = \{a \in \mathbb{R}^L : \exists (s,r) \in S \times R \; s.t. \; a = \lambda\bar{\alpha}(s,r) + (1-\lambda)a_0\}$. 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 $\bar{A}_\lambda$. Define $a_1 \in A$ as an arbitrary element of the relative interior of $\bar{A}_\lambda$ 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 $\mu'' = \eta_R\{\mu', \mu\}$, f_0 does not depend on r. Now define a shifted $\bar{\alpha}$ function,

$$\tilde{\alpha}(s,r) = \bar{\alpha}(s,r) + \frac{\epsilon\rho f_0(s)}{1 - \epsilon\rho f_0(s) - \epsilon(1-\rho)}(\bar{\alpha}(s,r) - a_1) + \frac{\epsilon(1-\rho)}{1 - \epsilon\rho f_0(s) - \epsilon(1-\rho)}(\bar{\alpha}(s,r) - a_0)$$
$$= \frac{\bar{\alpha}(s,r) - a_1\epsilon\rho f_0(s)a_1 - \epsilon(1-\rho)a_0}{1 - \epsilon\rho f_0(s) - \epsilon(1-\rho)}$$

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 $\bar{A}_\lambda$. 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 $\bar{A}_\lambda$. Consequently, for each $s \in S$, there is a probability measure $\nu_s \in \Delta(\bar{A}_\lambda)$ satisfying

$$E^{\nu_s(a)}[a] = \tilde{\alpha}(s,r)$$

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

$$\nu_{s,\epsilon,\rho}(a) = (1 - \epsilon\rho f_0(s) - \epsilon(1-\rho))\nu_s(a)$$
$$+ \epsilon\rho f_0(s)\delta_{a_1}(a) + \epsilon(1-\rho)\delta_{a_0}(a),$$

where $\delta_{a_1}(a)$ and $\delta_{a_0}(a)$ are the point masses on a_0 and a_1.

By construction, for all $r \in R$,

$$E^{\nu_{s,\epsilon,\rho}(a)}[a] = (1 - \epsilon\rho f_0(s) - \epsilon(1-\rho))E^{\nu_s(a)}[a]$$
$$+ \epsilon\rho f_0(s)a_1 + \epsilon(1-\rho)a_0$$
$$= \bar{\alpha}(s,r).$$

Define $A^*_{\epsilon,\rho} = \{a \in \bar{A}_\lambda \cup \{a_0, a_1\} : \exists s \in S \text{ s.t. } \nu_{s,\epsilon,\rho}(a) > 0\}$. Define, for each $a \in A^*$, the posterior measure $\mu_{0,a} \in \mathcal{U}_0$ by

$$\frac{d\mu_{0,a}(s,r)}{d\mu_0(s,r)} = \frac{\nu_{s,\epsilon,\rho}(a)}{E^{\mu_0(s',r')}[\nu_{s',\epsilon,\rho}(a)]},$$

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 $\pi_0 \in \Pi_0(\mu_0)$ that is a finite sum of point masses,

$$d\pi_0(a, \mu_0') = \sum_{a' \in A^*} E^{\mu_0(s,r)}[\nu_{s,\epsilon,\rho}(a')]\delta_{(a',\mu_{0,a'})}(a, \mu_0').$$

This induces a $\pi_{\epsilon,\rho} \in \Pi(\mu)$ by $(a, \mu_0') \mapsto (a, \phi\{\mu_0', \bar{\alpha}\})$. By construction, $\pi_\epsilon(a_1, \mu'') = E^{\mu_0(s,r)}[\nu_{s,\epsilon,\rho}(a_1)] = \rho$ and $\pi_\epsilon(a_0, \mu) = E^{\mu_0(s,r)}[\nu_{s,\epsilon,\rho}(a_0)] = (1 - \rho)$, and for each $a \in A^*$, $(a, \mu') \in supp(\pi_{\epsilon,\rho})$ for at most one μ'.

G.6 Proof of Lemma 14

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

$$D(\mu'||\phi_{\bar{A}}\{\mu_0, \bar{\alpha}\}) < D(\eta_R\{\mu', \phi\{\mu_0, \bar{\alpha}\}\}||\phi\{\mu_0, \bar{\alpha}\}).$$

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 $\bar{A}$ and that is not a constant function.

Define $\mu'' = \eta_R\{\mu', \mu\}$, 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 13, which defines a $\pi_{\epsilon,\rho}$ that is mean-consistent with $\bar{\alpha}$. Let $\pi_0 \in \Pi_0(\mu_0)$ be the measure induced from $\pi_{\epsilon,\rho}$ by $(a, \mu''') \mapsto (a, \gamma_{\bar{A}}\{\mu'''\})$.

Let us now define a family $\pi_{0,\tau} \in \Pi_0(\mu_0)$ induced from π_0 by the mapping

$$(a, \mu_0') \mapsto \begin{cases} (a, \mu_0') & a \notin \{a_0, a_1\} \cup \mu_0' \notin \{\mu_0'', \mu_0\}, \\ (a, \mu_0'' + \tau(\gamma_{\bar{A}}\{\mu'\} - \mu_0'')) & (a, \mu_0') = (a_1, \mu_0''), \\ (a, \mu_0 - \frac{\rho}{1-\rho}\tau(\gamma_{\bar{A}}\{\mu'\} - \mu_0'')) & (a, \mu_0') = (a_0, \mu_0) \end{cases}$$

for some $\tau \geq 0$. 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 $\mu' = \mu$,

$$\frac{\partial}{\partial\tau} C(\pi_\tau, \mu)|_{\tau=0^+} = \epsilon\rho \frac{\partial}{\partial\tau} D(\phi\{\mu_0'' + \tau(\gamma_{\bar{A}}\{\mu'\} - \mu_0''), \bar{\alpha}\}|\mu))|_{\tau=0}.$$

By the convexity of the divergence,

$$\frac{\partial}{\partial\tau} D(\phi\{\mu_0'' + \tau(\gamma_{\bar{A}}\{\mu'\} - \mu_0''), \bar{\alpha}\}|\mu))|_{\tau=0} + D(\mu''||\mu) \leq D(\mu'||\mu),$$

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,

$$\eta_R\{\phi\{\mu_0'' + \tau(\gamma_{\bar{A}}\{\mu'\} - \mu_0''), \bar{\alpha}\}, \mu\} = \eta_R\{\mu'' + \tau(\mu' - \mu''), \mu\},$$
$$= (1 - \tau)\eta_R\{\mu'', \mu\} + \tau\eta_R\{\mu', \mu\}$$
$$= \mu''$$

and likewise

$$\eta_R\{\phi\{\mu_0 - \frac{\rho}{1-\rho}\tau(\gamma_{\bar{A}}\{\mu'\} - \mu_0''), \bar{\alpha}\}, \mu\} = \mu.$$

Therefore, π is induced from π_τ by $(a, \mu') \to (a, \eta_R\{\mu', \mu\})$, proving the claim for $\hat{\pi} = \pi_\tau$.

G.7   Proof of Lemma 15

Define $\mu = \phi\{\mu_0, \bar{\alpha}\}$. Define the function $D_S : \Delta(S) \to [0, \infty]$ by

$$D_S(\mu_s; \mu) = D(\mu\{\mu_s; \mu\}||\mu)$$

where

$$d\mu(\mu_s; \mu) = \frac{d\mu_s}{d\gamma_{R\bar{A}}\{\mu\}}d\mu$$

and $\gamma_{R\bar{A}} : \bar{\mathcal{U}} \to \Delta(S)$ computes the marginal distribution on 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 $h : S \to \mathbb{R}$ such that

$$D_S(\mu'_s; \mu) \geq D_S(\mu_s; \mu) + E^{\mu'_s(s)}[h(s; \mu_s)] - E^{\mu_s(s)}[h(s; \mu_s)].$$

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 $E^{\mu_s(s)}[h(s; \mu_s)] = 0$, as $\hat{h}(s; \mu_s) = h(s; \mu_s) - E^{\mu_s(s)}[h(s; \mu_s)]$ 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 $\mu' \ll \mu$,

$$D(\mu'||\mu) \geq D_S(\mu_s; \mu) + E^{\mu'_s(s)}[h(s; \mu_s)].$$

Define $A^* \subset A$ as the set of actions with $(a, \mu'_a) \in supp(\pi)$ for some $\mu'_a \in \mathcal{U}$, and let μ_a be the unique corresponding measure.

Define, for all $a \in A^*$, $s \in S$, and an arbitrary $r_0 \in R$,

$$\tilde{u}(a, \bar{\alpha}(s, r_0), s) = D(\mu_a||\mu) + h(s; \gamma_{R\bar{A}}\{\mu_a\}).$$

We extend the utility function $\tilde{u}$ to other values of $\bar{a}$ using, for all $a \in A^*$,

$$\tilde{u}(a, \bar{a}, s) = \tilde{u}(a, \bar{\alpha}(s, r_0), s) - \chi(a - \bar{\alpha}(s, r_0))^T \cdot (\bar{a} - \bar{\alpha}(s, r_0)) + \frac{\chi}{2}|\bar{a} - \bar{\alpha}(s, r_0)|^2.$$

We extend this utility function to other actions $a \in A \setminus A^*$ by defining

$$\tilde{u}(a, \bar{a}, s) = \min_{a' \in A^*} \tilde{u}(a', \bar{a}, s) - |a' - a|^2.$$

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 $\tilde{u}$ is not differentiable at certain points. However, we can define a mollified version of it,

$$u_\delta(a, \bar{a}, s) = \int_A \omega_\delta(a'' - a)\tilde{u}(a'', \bar{a}, s)da'',$$

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,

$$u_\delta(a, \bar{\alpha}(s, r_0), s) = \tilde{u}(a, \bar{\alpha}(s, r_0), s) - c_\delta$$

for all $a \in A^*$, and that u_δ is continuously differentiable, while preserving the property that actions not in A^* are dominated.

Defining $u(a, \bar{a}, s) = u_\delta(a, \bar{\alpha}(s, r_0), s) + c_\delta$ results in the utility function

$$u(a, \bar{a}, s) = g(a; s) - \frac{\chi}{2}|\bar{a} - \bar{\alpha}(s, r_0)|^2 - \chi(a - \bar{a})^T \cdot (\bar{a} - \bar{\alpha}(s, r_0)),$$

where $g(a; s) = u_\delta(a, \bar{\alpha}(s, r_0), s) + c_\delta$.

Let us first show that $(\pi, \bar{\alpha})$ is an equilibrium. Consider the relaxed agent's problem,

$$\max_{\pi_0 \in \Delta(A \times \mathcal{U}_0)} E^{\pi_0(a, \mu_0')}[E^{\mu_0'(s,r)}[u(a, \bar{\alpha}(s, r_0), s)] - D(\phi\{\mu_0', \bar{\alpha}\}||\mu)],$$

which is the agent's problem without the Bayes-consistency constraint. It is immediate a necessary and sufficient condition for optimality is that, for any $(a, \mu_0') \in supp(\pi_0^*)$,

$$E^{\mu_0'(s,r)}[u(a, \bar{\alpha}(s, r_0), s)] - D(\phi\{\mu_0', \bar{\alpha}\}||\mu) \geq E^{\mu_0''(s,r)}[u(a', \bar{\alpha}(s, r_0), s)] - D(\phi\{\mu_0'', \bar{\alpha}\}||\mu)$$

for all $(a', \mu_0'') \in A \times \mathcal{U}_0$.

Because the actions in $A \setminus A^*$ are dominated by actions in A^*, no such action is optimal in this problem. Choosing $a' = a$ for some $a \in a^*$ and $\mu_0' = \gamma_{\bar{A}}\{\mu_a\}$ yields,

$$0 \geq D(\mu_a||\mu) - D(\mu'||\mu) + E^{\mu'(s,r,\bar{a})}[h(s; \gamma_{R\bar{A}}\{\mu_a\})],$$

which holds by the definition of the sub-gradient. By construction, for all $a \in a^*$,

$$E^{\mu_a(s,r,\bar{a})}[u(a, \bar{a}, s)] - D(\mu_a||\mu)] = 0,$$

and consequently π is a maximizer of the actual agent's problem. It follows that $(\pi, \bar{\alpha})$ is an equilibrium.

G.8 Proof of Lemma 17

Let us think of $\omega = (a, \mu_0') \in supp(\pi_0^*)$ 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 $\epsilon \in \mathbb{R}$,

$$d\nu_\epsilon(\omega|s, r) = \begin{cases} d\nu(\omega|s, r) & (s, r) \neq (s_0, r_0), \\ (1 - \epsilon + \frac{\epsilon}{\nu(\Omega_0|s, r)}) d\nu(\omega|s, r) & (s, r) = (s_0, r_0), \omega \in \Omega_0, \\ (1 - \epsilon) d\nu(\omega|s, r) & otherwise. \end{cases}$$

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

$$\frac{\mu_{0,\epsilon}'(s, r; \omega)}{\mu_0(s, r)} = \begin{cases} \frac{\mu_0'(s, r; \omega)}{1 + (\frac{\epsilon}{\nu(\Omega_0|s_0, r_0)} - \epsilon)\mu_0'(s_0, r_0; \omega)} & (s, r) \neq (s_0, r_0), \omega \in \Omega_0 \\ \frac{(1 - \epsilon + \frac{\epsilon}{\nu(\Omega_0|s_0, r_0)})\mu_0'(s, r; \omega)}{1 + (\frac{\epsilon}{\nu(\Omega_0|s_0, r_0)} - \epsilon)\mu_0'(s_0, r_0; \omega)} & (s, r) = (s_0, r_0), \omega \in \Omega_0 \\ \frac{\mu_0'(s, r; \omega)}{1 - \epsilon\mu_0'(s_0, r_0; \omega)} & (s, r) \neq (s_0, r_0), \omega \in supp(\pi_0^*) \setminus \Omega_0 \\ \frac{(1 - \epsilon)\mu_0'(s, r; \omega)}{1 - \epsilon\mu_0'(s_0, r_0; \omega)} & (s, r) = (s_0, r_0), \omega \in supp(\pi_0^*) \setminus \Omega_0 \end{cases}$$

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,

$$E^{\pi_{0,\epsilon}(a,\mu_0')}[a|s, r] = \begin{cases} \bar{\alpha}^*(s, r) & (s, r) \neq (s_0, r_0) \\ (1 - \epsilon)\bar{\alpha}^*(s, r) + \epsilon\frac{\int_{\Omega_0} a(\omega) d\nu(\omega|s_0, r_0)}{\nu(\Omega_0|s_0, r_0)} & otherwise, \end{cases}$$

and by the definition of Ω_0,

$$x = \frac{\int_{\Omega_0} a(\omega) d\nu(\omega|s_0, r_0)}{\nu(\Omega_0|s_0, r_0)} > 0.$$

G.9 Proof of Lemma 5

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,

$$E^{\pi_0(a,\mu')}[\mu_0'] = \mu_0.$$

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 $h : A \times \mathcal{U}_0 \to \Omega$. By the Radon-Nikodym theorem, there exists a set of functions $g_\omega : S \times R \to \mathbb{R}_+$ for all $\omega \in supp(\bar{\nu})$ such that

$$\int_\Omega g_\omega(s,r) d\bar{\nu}(\omega) = 1, \ \mu_0\text{-a.e.}$$

Define

$$\bar{g}(s,r) = 1 - \int_\Omega g_\omega(s,r) d\bar{\nu}(\omega)$$

and

$$f_\omega(s,r) = \begin{cases} 1 & \bar{g}(s,r) \neq 0 \\ g_\omega(s,r) & \bar{g}(s,r) = 0. \end{cases}$$

We have $f_\omega(s,r) = g_\omega(s,r)$ 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 $\omega \in supp(\bar{\nu})$, and as zero otherwise. We have, for all $\omega \in supp(\bar{\nu})$,

$$\bar{f}\{\nu,\mu\}(\omega) = \int_{S \times R} f_\omega(s,r) d\mu_0(s,r) = 1$$

and

$$f^\omega\{\nu,\mu\}(s,r,\bar{a}) = f_\omega(s,r),$$

and therefore $\hat{\pi}\{\bar{\nu},\nu,\mu\} = \pi$.

G.10   Proof of Lemma 7

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,

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

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 $h : A \times \mathcal{F}_0^+ \to \Omega$. Define $f_{h(a,f_0)} = f_0$ and define

$$\nu(\omega|s,r,\bar{a}) = f_\omega(s,r)$$

for all $\omega \in supp(\bar{\nu})$, and as zero otherwise. We have, for all $\omega \in supp(\bar{\nu})$,

$$\bar{f}\{\nu, \mu\}(\omega) = \int_{S \times R} f_\omega(s, r) d\mu_0(s, r) = 1$$

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

$$g(s) = \frac{\bar{\alpha}_s(s, r)}{1 + \bar{\alpha}_s(s, r)^2} \frac{\partial}{\partial s'} \nu(\omega|s', r, \bar{\alpha}(s, r))|_{s'=1}.$$

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 $g(s)$ is bounded. Define

$$\nu_1(\omega|s, r, \bar{a}) = \min(\max(\bar{a} - \bar{\alpha}(s, r), -\delta), \delta) * g(s)$$
$$+ \nu(\omega|s, r, \bar{\alpha}(s, r)).$$

For sufficiently small δ, by the boundedness of $g(s)$, this is non-negative and the optimality condition (22).

We can verify that

$$\int_{supp(\bar{\nu})} \nu_1(\omega|s, r, \bar{a}) d\bar{\nu}(\omega) = 1$$

by observing that

$$\frac{\bar{\alpha}_s(s, r)}{1 + \bar{\alpha}_s(s, r)^2} \int_{supp(\bar{\nu})} \frac{\partial}{\partial s'} \nu(\omega|s', r, \bar{a})|_{s'=s} d\bar{\nu}(\omega) = 0$$

everywhere.

G.11 Proof of Lemma 19

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, $g(a, f_0) = f_0$ 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 $1 = \int_{supp(\bar{\pi}_n)} f_0 d\bar{\pi}_n(a, f_0) \to \int_{supp(\bar{\pi}_0)} f_0 d\bar{\pi}_0(a, f_0)$, 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 $\bar{\pi}_1, \bar{\pi}_2 \in \bar{\Pi}_0$, $\int_{supp(\bar{\pi}_1)} f_0 d\pi_1(a, f_0) = 1$ and $\int_{supp(\bar{\pi}_2)} f_0 d\bar{\pi}_2(a, f_0) = 1$, 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.

G.12   Proof of Lemma 22

Define the function $g_{s,r} : A \times \mathcal{F}_0^+ \to \mathbb{R}$ by

$$g_{s,r}\{a, f_0\} = af_0(s, r).$$

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

$$\begin{aligned} |a_n f_{0,n}(s, r) - af_0(s, r)| &\leq |a_n - a| \sup_{(s',r')\in S\times R} |f_{0,n}(s', r')| \\ &+ |a| \sup_{(s',r')\in S\times R} |f_{0,n}(s', r') - f_0(s', r')| \end{aligned}$$

and consequently, by the compactness of $\mathcal{F}_0^+$ and A, if $(a_n, f_{0,n}) \to (a, f_0)$, then $g_{s,r}(a_n, f_{0,n}) \to g_{s,r}(a, f_0)$ for all $(s, r) \in S \times R$. Consequently, $g_{s,r}$ 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 $\bar{\pi}_{0,n} \to \bar{\pi}_0$, then $M\{\bar{\pi}_{0,n}\}(s, r) \to M\{\bar{\pi}_0\}(s, r)$ 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 $M\{\bar{\pi}_0\}$, $\bar{A}^+$, 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, and hence 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 $(a, f_0) \in supp(\bar{\pi}_0)$, $f_0(s, r) = f_0(s, r')$ 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 $(r, r') \in R$.