Unrestricted information acquisition*

Tommaso Denti[†]

July 2022[‡]

Abstract

In many strategic environments, information acquisition is a central component of the game that is played. Being uncertain about a payoff-relevant state, a player in a game has a two-fold incentive to acquire information: learning the state and learning what others know. We develop a model of information acquisition in games that accounts for players' incentive to learn what others know. In applications to rational inattention and global games, we prove the power of this incentive. When information acquisition is "independent"—that is, players can acquire information only about the state—severe coordination problems emerge among rationally inattentive players. When information acquisition is "unrestricted"—that is, players can acquire information about the state and each other's information in a flexible way—we show that rational inattention admits a sharp logit characterization, and we provide a new rationale for selecting risk dominant equilibria in coordination games.

1 Introduction

Strategic interaction often is not based on fixed prior information: information acquisition is a central component of the game that is played. Facing uncertainty about a

*This paper is based on my dissertation, submitted to the MIT Economics Department in June 2016. I am grateful to my advisors Muhamet Yildiz and Juuso Toikka for their continuing guidance and support. I thank the editor, Marina Halac, and the referees for their comments and suggestions. I also thank Marco Battaglini, Dirk Bergemann, Larry Blume, Alessandro Bonatti, Giulia Brancaccio, Glenn Ellison, Tibor Heumann, Mihai Manea, Massimo Marinacci, Alessandro Pavan, Ming Yang, Alex Wolitzky, and, especially, Stephen Morris for very helpful discussions.

[†]Cornell University. Email: tjd237@cornell.edu.

[‡]First version: November 2015.

payoff-relevant state, a player has a *two-fold* incentive to acquire information: learning the state *and* learning what others know. Learning the state is of course necessary to optimally choose what action to take, but it is not sufficient. A player realizes that the actions opponents take depend on what information they have; this generates an incentive to learn what others know.

In this paper, we propose a model of information acquisition in games that accounts for players' incentive to learn what others know. We fix a game where the players are uncertain about some feature of the environment that affects payoffs; such uncertainty is measured by a common prior over possible states. The traditional approach is to also fix an information structure and analyze the Bayes Nash equilibria. An *information structure* details how the joint probability distribution of the players' signals depends on the state; in a *Bayes Nash equilibrium*, each player observes her signal, uses Bayes' rule to compute beliefs, and takes an action that is optimal in expectation.

Here, we make the information structure endogenous. For each player, we fix a feasible set of costly experiments. An *experiment* details how the probability distribution of the player's own signal depends on the state and the opponents' signals; different experiments represent different levels of information about the state and what others know; the associated costs incorporate limitations in acquiring information. Every information structure induces an experiment for each player. We say that a probability distribution over states and actions is *robust to information acquisition* if it corresponds to a Bayes Nash equilibrium for some information structure such that for each player, the induced experiment maximizes the value minus the cost of information.

Robustness to information acquisition is the new concept we introduce in this paper: it reflects the idea that no player has an incentive to change what she knows about the state and what others know. It can been seen as a refinement of *Bayes correlated equilibrium* (Bergemann and Morris, 2016).[1] As Bergemann and Morris show, a probability distribution over states and actions is a Bayes correlated equilibrium if it corresponds to a Bayes Nash equilibrium for some information structure. Robustness to information acquisition adds the condition that the experiments induced by the information structure are individually optimal in balancing the value and the cost of

[1]Bayes correlated equilibrium has gained attention, especially for applications to *information design* (Bergemann and Morris, 2019).

information.

Whether or not there exists an equilibrium that is robust to information acquisition depends on what experiments are feasible and what properties the cost of information satisfies. When there is no prior uncertainty about the state, in which case a Bayes correlated equilibrium coincides with an Aumann correlated equilibrium, a robust equilibrium exists under broad conditions on the cost of information—any Nash equilibrium is robust to information acquisition, provided that the players can privately randomize over actions. In general, it is possible that no equilibrium is robust to information acquisition; intuitively, it could be that the players do not "agree" on any information structure. The concept is, nevertheless, useful in a variety of circumstances, as we show in applications.

We present applications to *rational inattention* (Sims, 2003) and *global games* (Carlsson and van Damme, 1993). In the applications, we focus on two cases of particular interest. We say that information acquisition is *unrestricted* if all experiments are feasible: the players can acquire information about the state and each other's information in a flexible way. We say that information acquisition is *independent* if all feasible experiments are independent of others' signals: the players can acquire information only about the state. Most previous works have focused on independent information acquisition (e.g., Persico, 2000; Bergemann and Valimaki, 2002; Yang, 2015); our model nests these contributions. By comparing unrestricted and independent information acquisition, we are able to assess the effect of players' incentive to learn what others know.

In the first application, we study games between rationally inattentive players. Rational inattention is an influential model of costly information acquisition in economics.[2] According to this model, agents pay attention "as if" they optimally acquire costly information; in a widespread specification, the cost of information is proportional to the expected reduction in the entropy of beliefs (Matejka and McKay, 2015). One challenge has been understanding the implications of rational inattention in settings—such as games—where noise is endogenously correlated across agents.[3]

[2]See Mackowiak, Matejka and Wiederholt (Forthcoming) for a survey of the literature. See also Alaoui and Penta (2022) for a broader approach to rational inattention.

[3]As Sims (2010, p. 171) puts it: "Another issue that arises in bringing rational inattention to equilibrium models is that the rational inattention models of individual behavior have nothing to say about properties of information processing error other than its conditional distribution given decision choices. [...] Information-processing noise that is independent across agents will average out in macroeconomic behavior, whereas highly correlated information processing noise will become

Focusing on *potential games* (Monderer and Shapley, 1996), we provide a simple and transparent characterization of rational inattention in games. We show that a probability distribution over states and actions is robust to *unrestricted* information acquisition if and only if the players follow a *joint logit rule*. The result can be seen as a many player extension of Matejka and McKay's *individual logit rule* for rationally inattentive decision makers, one of the cornerstones of the literature.

Independent information acquisition generates starkly different predictions. In a thought-provoking work, Yang (2015) studies a coordination game between rationally inattentive agents who can acquire information *only* about the state; he finds a large multiplicity of equilibria. Yang's indeterminacy result casts doubt on whether a meaningful analysis of rational inattention in games is possible. We revisit his findings in the broader context of potential games and highlight the significance of players' incentive to learn what others know. When information acquisition is *unrestricted*, the behavior of rationally inattentive players admits a sharp logit characterization; in particular, as the cost of information vanishes, a *unique* equilibrium arises in which the players coordinate on maximizing the potential function of the game.

We expand the study of information acquisition and equilibrium selection in a second application where we focus on global games. Global games is a prominent approach to equilibrium selection in coordination games.[4] According to this approach, an equilibrium is selected if it survives a small perturbation of common knowledge. One limitation of global games is that the selection rule crucially depends on the nature of the perturbation (Kajii and Morris, 1997; Weinstein and Yildiz, 2007). Studying the case of the negligible costs of information, we provide a micro-foundation to perturbations of common knowledge and revisit selection rules from global games. Our analysis provides a new rationale for selecting *risk dominant equilibria* (Harsanyi and Selten, 1988).

A few recent papers have also relaxed the independence assumption on information acquisition in games (e.g., Hellwig and Veldkamp, 2009; Myatt and Wallace, 2012). These papers mostly focus on specific environments—for example, beauty contests where uncertainty is normally distributed. Our model, which applies to general games and arbitrary spaces of uncertainty, follows in their footsteps in accounting for players' incentive to learn what others know.

an additional source of macroeconomic randomness."

[4]See Morris and Shin (2003) for a survey of the literature.

2 Robustness to information acquisition

2.1 Definition

A finite set of individuals $N = \{1, \ldots, n\}$, with typical element i, play a simultaneous game. The players are uncertain about a payoff-relevant state θ, which is drawn from a finite set Θ according to a full-support probability distribution $\pi \in \Delta(\Theta)$. A player i has a finite set of actions A_i and a utility function $u_i : \Theta \times A \to \mathbb{R}$, where we adopt the usual notation $A_{-i} = \prod_{j \neq i} A_j$ and $A = A_i \times A_{-i}$.

Let $\Delta_\pi(\Theta \times A)$ be the set of joint probability distributions over states and actions whose marginal distribution over states is π:

$$\Delta_\pi(\Theta \times A) = \left\{ \rho \in \Delta(\Theta \times A) : \sum_{a \in A} \rho(\theta, a) = \pi(\theta) \text{ for all } \theta \in \Theta \right\}.$$

We take the view that correlation in actions is the result of players acting on information about the state and what others know. We want to formalize the idea that $\rho \in \Delta_\pi(\Theta \times A)$ is "robust to information acquisition" if no player has an incentive to change what she knows about the state and what others know.

We start with the notion of correlated equilibrium for games of incomplete information as in Bergemann and Morris (2016).[5]

Definition 1 (Bergemann and Morris, 2016)**.** A probability distribution $\rho \in \Delta_\pi(\Theta \times A)$ is a *Bayes correlated equilibrium (BCE)* if for all $i \in N$ and $a_i, a_i' \in A_i$,

$$\sum_{\theta \in \Theta} \sum_{a_{-i} \in A_{-i}} \left(u_i(\theta, a_i, a_{-i}) - u_i(\theta, a_i', a_{-i}) \right) \rho(\theta, a_i, a_{-i}) \geq 0.$$

As Bergemann and Morris show, BCE characterizes all Bayes Nash equilibria that can arise across information structures. For each player i, let X_i be a standard Borel space of signals—we adopt the notation $X_{-i} = \prod_{j \neq i} X_j$ and $X = X_i \times X_{-i}$.[6] An *information structure* $\mu \in \Delta_\pi(\Theta \times X)$ is a joint probability distribution over states and signals whose marginal distribution over states is π. For each player i, a *mixed-action rule* is a measurable function $\sigma_i : X_i \to \Delta(A_i)$; we denote by $\sigma_i(a_i|x_i) \in [0, 1]$

[5] See also the *Bayesian solution* of Forges (1993).

[6] Throughout, Cartesian products of measurable spaces are endowed with the product σ-algebras.

the probability that player i takes an action a_i after having observed a signal x_i.[7] We denote by Σ_i the set of i's mixed-action rules—we adopt the notation $\Sigma_{-i} = \prod_{j \neq i} \Sigma_j$ and $\Sigma = \Sigma_i \times \Sigma_{-i}$. With a slight abuse of notation, we define

$$u_i(\theta, \sigma(x)) = \sum_{a \in A} u_i(\theta, a) \prod_{j \in N} \sigma_j(a_j|x_j).$$

Proposition 1 (Bergemann and Morris, 2016)**.** *Assume that for every player i, X_i contains more elements than A_i. A probability distribution $\rho \in \Delta_\pi(\Theta \times A)$ is a BCE if and only if there are an information structure $\mu \in \Delta_\pi(\Theta \times X)$ and a profile of mixed-action rules $\sigma \in \Sigma$ that satisfy the following conditions:*

(i) For all $\theta \in \Theta$ and $a \in A$,

$$\rho(a|\theta) = \int_X \prod_{i \in N} \sigma_i(a_i|x_i) \, \mathrm{d}\mu(x|\theta).$$

(ii) For all $i \in N$, σ_i is an optimal solution of

$$\max_{\sigma_i' \in \Sigma_i} \int_{\Theta \times X} u_i(\theta, \sigma_i'(x_i), \sigma_{-i}(x_{-i})) \, \mathrm{d}\mu(\theta, x).$$

Condition (i) states that ρ is the probability distribution over states and actions induced by μ and σ. Condition (ii) states that σ is a Bayes Nash equilibrium when μ is the players' information structure.

The concept of BCE takes the conservative view that all information structures are possible. However, intuition suggests that some information structures are more likely to occur than others. For an extreme example, suppose that the utility of a player is constant in the state and the opponents' actions. Intuition suggests that such a player has little incentive to acquire information: if acquiring information is costly, then an information structure where the player knows a lot may be viewed as not "robust."

Next we formalize a notion of robustness to information acquisition. For each player i, an *experiment* is a measurable function $P_i : \Theta \times X_{-i} \to \Delta(X_i)$ that details how the distribution of i's signal depends on the state and the opponents' signals; we denote by $P_i(B_i|\theta, x_{-i}) \in [0, 1]$ the probability that x_i belongs to a Borel set $B_i \subseteq X_i$

[7] A function $\sigma_i : X_i \to \Delta(A_i)$ is measurable if for every $a_i \in A_i$, $\sigma_i(a_i|x_i)$ is measurable in x_i.

conditional on θ being the state and x_{-i} being the signals observed by the opponents.[8] We fix a set of feasible experiments $\mathcal{P}_i$—we adopt the notation $\mathcal{P}_{-i} = \prod_{j \neq i} \mathcal{P}_j$ and $\mathcal{P} = \mathcal{P}_i \times \mathcal{P}_{-i}$. A function $C_i : \mathcal{P}_i \times \Delta_\pi(\Theta \times X_{-i}) \to \mathbb{R}_+$ represents i's *cost of acquiring information*, which we allow to depend both on i's experiment and on i's ex ante beliefs about the state and the opponents' signals.[9] For every information structure $\mu \in \Delta_\pi(\Theta \times X)$, we denote by $\mu_{-i} \in \Delta_\pi(\Theta \times X_{-i})$ the marginal distribution over states and the opponents' signals.

Definition 2. A probability distribution $\rho \in \Delta_\pi(\Theta \times A)$ is *robust to information acquisition* if there is an information structure $\mu \in \Delta_\pi(\Theta \times X)$, a profile of experiments $P \in \mathcal{P}$, and a profile of mixed-action rules $\sigma \in \Sigma$ such that the following conditions are satisfied:

(i) For all $\theta \in \Theta$ and $a \in A$,

$$\rho(a|\theta) = \int_X \prod_{i \in N} \sigma_i(a_i|x_i) \, \mathrm{d}\mu(x|\theta).$$

(ii) For all $\theta \in \Theta$, $B_i \subseteq X_i$, and $B_{-i} \subseteq X_{-i}$,

$$\mu(B_i \times B_{-i}|\theta) = \int_{B_{-i}} P_i(B_i|\theta, x_{-i}) \, \mathrm{d}\mu_{-i}(x_{-i}|\theta)$$

(iii) For all $i \in N$, (P_i, σ_i) is an optimal solution of

$$\max_{P_i' \in \mathcal{P}_i, \sigma_i' \in \Sigma_i} \int_{\Theta \times X} u_i(\theta, \sigma_i'(x_i), \sigma_{-i}(x_{-i})) \, \mathrm{d}P_i'(x_i|\theta, x_{-i}) \, \mathrm{d}\mu_{-i}(\theta, x_{-i}) - C_i(P_i', \mu_{-i}).$$

Condition (i) states that ρ is the probability distribution over states and actions induced by μ and σ. Condition (ii) states that P_i is, with respect to μ, a (regular) conditional distribution of x_i given θ and x_{-i}. Condition (iii) states that (P_i, σ_i) is an optimal solution of an information acquisition problem where player i maximizes the value of information (i.e., expected utility) minus the cost of information.

[8] A function $P_i : \Theta \times X_{-i} \to \Delta(X_i)$ is measurable if for every Borel set $B_i \subseteq X_i$, $P_i(B_i|\theta, x_{-i})$ is measurable in θ and x_{-i}.

[9] The dependence of C_i on i's ex ante beliefs is important for applications to rational inattention but can otherwise be dispensed with. See Denti, Marinacci and Rustichini (2002) for an analysis of the relation between the cost of information and prior beliefs.

Definition 2 is the central concept we introduce in this paper: it reflects the idea that ρ is robust to information acquisition when no player has an incentive to change what she knows about the state and what others know. It is easy to see that if ρ is robust to information acquisition, then ρ is a BCE. In the rest of the paper, we will present many instances in which the opposite is not true and the notion of robustness we propose has bite.

In most studies of information acquisition in games, players can acquire information only about the state (e.g., Persico, 2000; Bergemann and Valimaki, 2002; Yang, 2015). Definition 2 nests these contributions: it is enough to assume that for every player i and every feasible experiment $P_i \in \mathcal{P}_i$, the distribution of x_i depends only on θ and not on x_{-i}. In that case, we say that information acquisition is *independent.*

In applications, we will compare independent information acquisition to the case in which *all* experiments are feasible—that is, for every player i, $\mathcal{P}_i$ is the set of *all* measurable functions from $\Theta \times X_{-i}$ to $\Delta(X_i)$. In that case, we say that information acquisition is *unrestricted.* Comparing independent and unrestricted information acquisition will allow us to assess the power of the incentive to learn what others know.

We conclude the section by discussing a few technical and conceptual issues concerning Definition 2:

Remark 1. Definition 2 treats the information structure μ and the profile of experiments P as separate objects, instead of deriving P from μ via condition (ii). The reason is that (ii) determines each P_i only μ_{-i}-almost surely, while C_i could assign different costs to different versions of the conditional distribution of x_i given θ and x_{-i} (think of the case in which i's cost of information is independent of her ex-ante beliefs about the state of the the signals of others).

Remark 2. As in Bergemann and Morris (2016), Definition 2 maintain the hypothesis that the players share a common prior over the entire space of uncertainty (states and signals). Non-common priors, however, naturally arise in the context of information acquisition. For example, there is no shortage of situations in which individuals agree to disagree on the quality of information sources (e.g., newspapers of different slants).

Remark 3. Definition 2 assumes the existence of a real ex-ante stage where players evaluates experiments before nature draws the signals. This is different from Harsanyi's interim approach to incomplete information in which the ex-ante stage is merely a mathematical artifact to represent the players' hierarchies of beliefs.

Remark 4. Bergemann and Morris (2016) also consider a more general version of Definition 1 where a default information structure μ^* is given, and BCE characterizes all equilibria that can arise across information structures μ that provide "more information" than μ^*.[10] Analogously, we could extend Definition 2 by requiring that for each player i, each feasible experiment provides "more information" to player i than μ_i^*.[11]

2.2 A revelation principle

We have defined robustness to information acquisition for arbitrary signal spaces. Our initial finding is a version of the revelation principle: we may focus on experiments whose outcomes are mixed-action recommendations.[12] The result holds under broad assumptions on the cost of information.

To formalize assumptions and results, we first need to define two rankings of experiments. The first ranking is standard:

Definition 3 (Blackwell, 1951)**.** An experiment P_i is a *garbling* of an experiment P_i', denoted by $P_i \preceq P_i'$, if there is a measurable function $K_i : X_i \to \Delta(X_i)$ such that for all $\theta \in \Theta$, $B_i \subseteq X_i$, and $x_{-i} \in X_{-i}$,

$$P_i(B_i|\theta, x_{-i}) = \int_{X_i} K_i(B_i|x_i') \, \mathrm{d}P_i'(x_i'|\theta, x_{-i}).$$

The second ranking of experiments is—to the best of our knowledge—new:

Definition 4. An experiment P_i is a *shrinking* of an experiment P_i', denoted by $P_i \trianglelefteq P_i'$, if there is a measurable function $K_{-i} : \Theta \times X_{-i} \to \Delta(X_{-i})$ such that for all $\theta \in \Theta$, $B_i \subseteq X_i$, and $x_{-i} \in X_{-i}$,

$$P_i(B_i|\theta, x_{-i}) = \int_{X_{-i}} P_i'(B_i|\theta, x_{-i}') \, \mathrm{d}K_{-i}(x_{-i}'|\theta, x_{-i}).$$

Garbling and shrinking capture two complementary yet different notions of informativeness. In both cases, one obtains P_i by adding noise to P_i', thus making P_i "less informative" than P_i'. The difference is in the *timing* with which the noise is added.

[10] We refer the interested reader to their paper for the appropriate definition of "more information."

[11] Formally, every $P_i \in \mathcal{P}_i$ is a garbling of a conditional μ_i^*-distribution of x_i given (θ, x_{-i}).

[12] The term "revelation principle" usually applies to situations where players *send* signals. Here, as in Bergemann and Morris (2016), we use the same term in situations where players *receive* signals.

In the case of a garbling, one obtains P_i by *first* performing P_i' and *then* adding noise to its outcome x_i'—noise represented by K_i. In the case of a shrinking, one obtains P_i by *first* adding noise to x_{-i}—noise represented by K_{-i}—and *then* performing P_i' on the realized x_{-i}'. The following example describes a garbling that is not a shrinking.

Example 1. Let P_i and P_i' be a pair of completely uninformative experiments: under P_i and under P_i', the distribution of x_i is independent of θ and x_{-i}. As is well known, P_i is a garbling of P_i' (and vice versa). However, as readily checked, P_i is a shrinking of P_i' if and only if $P_i = P_i'$, that is, if and only if the distribution of x_i is the same under P_i and under P_i'. ▲

Next is an example of a shrinking that is not a garbling.

Example 2. Let P_i and P_i' be a pair of experiments that coincide up to a relabelling of X_{-i}: for all $\theta \in \Theta$, $B_i \subseteq X_i$, and $x_{-i} \in X_{-i}$,

$$P_i(B_i|\theta, x_{-i}) = P_i'(B_i|\theta, f(x_{-i}))$$

where $f : X_{-i} \to X_{-i}$ is a bijective function such that both f and f^{-1} are measurable. It is easy to see that P_i is a shrinking of P_i' (and vice versa). On the other hand, in general, P_i need not be a garbling of P_i'. ▲

Our revelation principle holds when the cost of information is monotone with respect to garbling and shrinking. Next is monotonicity with respect to garbling:

Assumption 1. *If $P_i \preceq P_i'$ and $P_i' \in \mathcal{P}_i$, then $P_i \in \mathcal{P}_i$ and $C_i(P_i, \mu_{-i}) \leq C_i(P_i', \mu_{-i})$.*

In other terms, if P_i is a garbling of P_i' and P_i' is feasible, then P_i is feasible and less expensive than P_i'.

To define monotonicity with respect to shrinking, we need to account for the possible dependence of the cost of information on prior beliefs. Thus, we extend the notion of shrinking to *pairs* of experiments and priors: we define $(P_i, \mu_{-i}) \trianglelefteq^* (P_i', \mu_{-i}')$ if there is a measurable function $K_{-i} : \Theta \times X_{-i} \to \Delta(X_{-i})$ such that for all $\theta \in \Theta$, $B_i \subseteq X_i$, $x_{-i} \in X_{-i}$, and $B_{-i} \subseteq X_{-i}$,

$$P_i(B_i|\theta, x_{-i}) = \int_{X_{-i}} P_i'(B_i|\theta, x_{-i}')\, \mathrm{d}K_{-i}(x_{-i}'|\theta, x_{-i}),$$

$$\mu_{-i}'(B_{-i}|\theta) = \int_{X_{-i}} K_{-i}(B_{-i}|\theta, x_{-i}')\, \mathrm{d}\mu_{-i}(x_{-i}'|\theta).$$

Assumption 2. *If* $(P_i, \mu_{-i}) \trianglelefteq^* (P_i', \mu_{-i}')$ *and* $P_i' \in \mathcal{P}_i$*, then* $P_i \in \mathcal{P}_i$ *and* $C_i(P_i, \mu_{-i}) \leq C_i(P_i', \mu_{-i}')$*.*

If the cost of information is independent of prior beliefs, then the assumption has the following equivalent, simpler formulation: if $P_i \trianglelefteq P_i'$ and $P_i' \in \mathcal{P}_i$, then $P_i \in \mathcal{P}_i$ and $C_i(P_i) \leq C_i(P_i')$.[13] In other terms, if P_i is a shrinking of P_i' and P_i' is feasible, then P_i is feasible and less expensive than P_i'.

Finally, we assume that there are more signals than mixed actions:

Assumption 3. *For every player* i*,* $\Delta(A_i)$ *is a Borel subset of* X_i*.*

We are now ready to state the revelation principle. In short, we denote by $\mathcal{A}_i$ the set of i's mixed actions—we adopt the notation $\mathcal{A}_{-i} = \prod_{j\neq i} \mathcal{A}_j$ and $\mathcal{A} = \mathcal{A}_i \times \mathcal{A}_{-i}$. Let $\mathcal{P}_{\mathcal{A}_i}$ be the set of i's feasible experiments $P_i \in \mathcal{P}_i$ whose outcomes are mixed-action recommendations:

$$\mathcal{P}_{\mathcal{A}_i} = \{P_i \in \mathcal{P}_i : \text{for all } \theta \text{ and } x_{-i}, P_i(\mathcal{A}_i|\theta, x_{-i}) = 1\}.$$

We adopt the notation $\mathcal{P}_{\mathcal{A}_{-i}} = \prod_{j\neq i} \mathcal{P}_{\mathcal{A}_j}$ and $\mathcal{P}_{\mathcal{A}} = \mathcal{P}_{\mathcal{A}_i} \times \mathcal{P}_{\mathcal{A}_{-i}}$.

Proposition 2. *Under Assumptions 1-3, a probability distribution* $\rho \in \Delta_\pi(\Theta \times A)$ *is robust to information acquisition if and only if there are* $\mu \in \Delta_\pi(\Theta \times \mathcal{A})$ *and* $P \in \mathcal{P}_{\mathcal{A}}$ *such that the following conditions are satisfied:*

(i) For all $\theta \in \Theta$ *and* $a \in A$*,*

$$\rho(a|\theta) = \int_X \prod_{i\in N} \alpha_i(a_i)\, \mathrm{d}\mu(\alpha|\theta).$$

(ii) For all $\theta \in \Theta$*,* $B_i \subseteq \mathcal{A}_i$*, and* $B_{-i} \subseteq \mathcal{A}_{-i}$*,*

$$\mu(B_i \times B_{-i}|\theta) = \int_{B_{-i}} P_i(B_i|\theta, \alpha_{-i})\, \mathrm{d}\mu_{-i}(\alpha_{-i}|\theta)$$

(iii) For all $i \in N$*,* P_i *is an optimal solution of*

$$\max_{P_i'\in\mathcal{P}_{\mathcal{A}_i}} \int_{\Theta\times\mathcal{A}} u_i(\theta, \alpha)\, \mathrm{d}P_i'(\alpha_i|\theta, \alpha_{-i})\, \mathrm{d}\mu_{-i}(\theta, \alpha_{-i}) - C_i(P_i', \mu_{-i}).$$

[13] If $(P_i, \mu_{-i}) \trianglelefteq^* (P_i', \mu_{-i}')$, then $P_i \trianglelefteq P_i'$. Conversely, if $P_i \trianglelefteq P_i'$, then for every $\mu_{-i} \in \Delta_\pi(\Theta \times X_{-i})$, there exists $\mu_{-i}' \in \Delta_\pi(\Theta \times X_{-i})$ such that $(P_i, \mu_{-i}) \trianglelefteq^* (P_i', \mu_{-i}')$.

In sum, without loss of generality, we can focus on experiments whose outcomes are mixed-action recommendations. The proposition is reminiscent of existing versions of the revelation principle for single-agent information acquisition problems (e.g., Corollary 1 in Matejka and McKay, 2015) and for games with independent information acquisition (e.g., Proposition 1 in Yang, 2015). Endogenously correlated information introduces two new features:

- When there is a single agent or information acquisition is independent, we may focus on experiments whose outcomes are *pure-action* recommendations. In the case of correlated information acquisition, we look at *mixed-action* recommendations. This accounts for the fact that players can *privately* randomize over actions: players can acquire information about others' signals, but not about others' realized actions.

- When there is a single agent or information acquisition is independent, a revelation principle holds under Assumption 1 and a weaker version of Assumption 3, namely, that A_i is a subset of X_i. In the case of correlated information acquisition, Assumption 2 is also necessary. When information acquisition is independent, Assumption 2 is always satisfied: an experiment $P_i : \Theta \to \Delta(X_i)$ is a shrinking of an experiment $P_i' : \Theta \to \Delta(X_i)$ if and only if $P_i = P_i'$.

To get some intuition for the role of Assumptions 1 and 2 in the revelation principle, let $\rho \in \Delta_\pi(\Theta \times A)$ be robust to information acquisition; take also $\mu \in \Delta_\pi(\Theta \times X)$, $P \in \mathcal{P}$, and $\sigma \in \Sigma$ that satisfy (i)-(iii) of Definition 2. The idea behind the revelation principle is to combine μ and σ into a single information structure $\nu \in \Delta_\pi(\Theta \times \mathcal{A})$ such that signals are mixed-action recommendations. The derived information structure ν generates, for every player i, an experiment $Q_i \in \mathcal{P}_{\mathcal{A}_i}$ that is (a version of) the conditional distribution of α_i given θ and α_{-i}. One can show that, under Assumptions 1 and 2, Q_i is optimal for player i. The key observation is that Q_i can be obtained by garbling *and* shrinking P_i. Thus, under Assumptions 1 and 2, Q_i is less expensive than P_i, and the optimality of Q_i (in the sense of Proposition 2-(iii)) follows from the optimality of P_i (in the sense of Definition 2-(iii)).

Many cost functions used in practice satisfy Assumptions 1 and 2. Next we provide a notable example from rational inattention that we will use later on in applications.

Example 3. For simplicity, suppose that X is finite. Let $I(P_i, \mu_{-i})$ be the *mutual information* of x_i and (θ, x_{-i}):

$$I(P_i, \mu_{-i}) = \sum_{\theta,x} P_i(x_i|\theta, x_{-i})\mu_{-i}(\theta, x_{-i}) \ln\left(\frac{P_i(x_i|\theta, x_{-i})\mu_{-i}(\theta, x_{-i})}{P_{\mu_{-i}}(x_i)\mu_{-i}(\theta, x_{-i})}\right)$$

where $P_{\mu_{-i}} \in \Delta(X_i)$ is the marginal distribution of i's signal. Mutual information is a measure of dependence between random variables, widely used in information theory (Cover and Thomas, 2006, Chapter 2). In many applications of rational inattention (e.g., Matejka and McKay, 2015), the cost of information is based on mutual information.

It is known that mutual information is monotone with respect to garbling.[14] Mutual information is also monotone with respect to shrinking. To see this, take (P_i, μ_{-i}) and (P'_i, μ'_{-i}) such that $(P_i, \mu_{-i}) \trianglelefteq (P'_i, \mu'_{-i})$. Simple algebra shows that $P_{\mu_{-i}} = P'_{\mu'_{-i}}$. Let $Q_i : X_i \to \Delta(\Theta \times X_{-i})$ and $Q'_i : X_i \to \Delta(\Theta \times X_{-i})$ be the conditional distributions of (θ, x_{-i}) given x_i induced by (P_i, μ_{-i}) and (P'_i, μ'_{-i}). We can view Q_i and Q'_i as experiments with input x_i and output (θ, x_{-i}). Simple algebra shows that Q_i is a garbling of Q'_i. Thus $I(P_{\mu_{-i}}, Q_i) \leq I(P_{\mu_{-i}}, Q'_i) = I(P'_{\mu'_{-i}}, Q'_i)$. Since mutual information is symmetric, $I(P_i, \mu_{-i}) = I(P_{\mu_{-i}}, Q_i)$ and $I(P'_i, \mu'_{-i}) = I(P'_{\mu'_{-i}}, Q'_i)$. We conclude that $I(P_i, \mu_{-i}) \leq I(P'_i, \mu'_{-i})$. ▲

Popular generalizations of mutual information also satisfy Assumptions 1 and 2:

Example 3 (Continued)**.** A popular generalization of mutual information is based on the notion of *f-divergence* (Csiszar, 1963, 1967; Ali and Silvey, 1966). Let $f : [0, \infty) \to \mathbb{R}$ be a continuous convex function. The quantity

$$I_f(P_i, \mu_{-i}) = \sum_{\theta,x} P_{\mu_{-i}}(x_i)\mu_{-i}(\theta, x_{-i}) f\left(\frac{P_i(x_i|\theta, x_{-i})\mu_{-i}(\theta, x_{-i})}{P_{\mu_{-i}}(x_i)\mu_{-i}(\theta, x_{-i})}\right)$$

is the *f-divergence* between the joint distribution of (θ, x) and the product of the marginal distributions of x_i and (θ, x_{-i}).[15] Mutual information corresponds to the case in which $f(t) = t \ln t$; in such a case, the f-divergence is called *relative entropy* (Cover and Thomas, 2006, Chapter 2).

[14]It follows from the so-called *data-processing inequality* (Cover and Thomas, 2006, Section 2.8).
[15]We adopt the convention $0f(0/0) = 0$.

It is known that I_f is monotone with respect to garbling.[16] Reasoning as above, we can check that I_f is also monotone with respect to shrinking. We will use f-divergences later on in an application to global games. ▲

Next is an example, also inspired by information theory, where the cost of information is independent of prior beliefs and satisfies Assumptions 1 and 2.

Example 3 (Continued)**.** Let the cost of information be the *channel capacity* of P_i:

$$I^*(P_i) = \max_{\mu_{-i}} I(P_i, \mu_{-i}).$$

Channel capacity is the maximum mutual information across all priors (Cover and Thomas, 2006, Chapter 7). For applications of channel capacity in economics, see Woodford (2012) and Nimark and Sundaresan (2019).

It is easy to see that channel capacity inherits monotonicity with respect to garbling from mutual information. Interestingly, channel capacity also inherits monotonicity with respect to shrinking. To see this, take P_i and P_i' such that $P_i \trianglelefteq P_i'$, with corresponding $K_{-i} : \Theta \times X_{-i} \to \Delta(X_{-i})$. For every prior μ_{-i}, we can define

$$\mu'_{-i}(\theta, x'_{-i}) = \sum_{x_{-i}} K_{-i}(x'_{-i}|\theta, x_{-i})\mu_{-i}(\theta, x_{-i}).$$

By construction, $(P_i, \mu_{-i}) \trianglelefteq (P_i', \mu'_{-i})$. Since mutual information is monotone with respect to shrinking, $I(P_i, \mu_{-i}) \leq I(P_i', \mu'_{-i}) \leq I^*(P_i')$. Since the choice of μ_{-i} is arbitrary, $I^*(P_i) \leq I(P_i')$.

The same argument applies to a generalization of channel capacity based on f-divergences:

$$I_f^*(P_i) = \max_{\mu_{-i}} I_f(P_i, \mu_{-i}).$$

Also I_f^* is monotone with respect to garbling and shrinking. ▲

As Example 2 highlights, monotonicity with respect to shrinking implies a form of invariance to the labelling of the space of uncertainty. Thus, cost functions that are sensitive to such a labelling (e.g., Pomatto, Strack and Tamuz, 2020; Hebert and Woodford, 2021) are, in general, not monotone with respect to shrinking.

While quite natural, the notion of shrinking of experiments is—to the best of our knowledge—new. The vast majority of cost functions used in applications are

[16] It follows from a so-called *data-processing lemma* (Liese and Vajda, 2006, Theorem 14).

monotone with respect to garbling. Versions of Assumption 2 occur in rational inattention. For example, the notions of invariance from Caplin, Dean and Leahy (2022, Definition 4) and Hebert and La'O (2021, Lemma 2) can be seen as special cases of monotonicity with respect to shrinking.

2.3 Learning about others' actions

We have assumed that players can *privately* randomize over actions; that is, players can randomize over actions without worrying that opponents will spy on the outcomes of the randomizations. The revelation principle, formalized by Proposition 2, highlights an implication of this restriction: without loss of generality, we can focus on mixed-action recommendations, but not on pure-action recommendations.

Of course, players' primary incentive is to acquire information about others' *realized actions*, not about others' mixed actions. To incorporate this incentive into the analysis, we propose a stronger notion of robustness to information acquisition.

For each player i, a *pure-action rule* is a measurable function $s_i : X_i \to A_i$. We denote by S_i the set of i's pure-action rules—we adopt the notation $S_{-i} = \prod_{j \neq i} S_j$ and $S = S_i \times S_{-i}$.

Definition 5. A probability distribution $\rho \in \Delta_\pi(\Theta \times A)$ is *strongly robust to information acquisition* if there is an information structure $\mu \in \Delta_\pi(\Theta \times X)$, a profile of experiments $P \in \mathcal{P}$, and a profile of pure-action rules $s \in S$ such that the following conditions are satisfied:

(i) For all $\theta \in \Theta$ and $a \in A$,

$$\rho(a|\theta) = \mu(\{x : s(x) = a\}|\theta).$$

(ii) For all $\theta \in \Theta$, $B_i \subseteq X_i$, and $B_{-i} \in X_{-i}$,

$$\mu(B_i \times B_{-i}|\theta) = \int_{B_{-i}} P_i(B_i|\theta, x_{-i}) \, \mathrm{d}\mu_{-i}(x_{-i}|\theta)$$

(iii) For all $i \in N$, (P_i, s_i) is an optimal solution of

$$\max_{P_i' \in \mathcal{P}_i, s_i' \in S_i} \int_{\Theta \times X} u_i(\theta, s_i'(x_i), s_{-i}(x_{-i})) \, \mathrm{d}P_i'(x_i|\theta, x_{-i}) \, \mathrm{d}\mu_{-i}(\theta, x_{-i}) - C_i(P_i', \mu_{-i}).$$

Definition 5 is a special case of Definition 2: if ρ is strongly robust to information acquisition, then it is robust. The difference is that in Definition 5, players must take pure actions conditional on all signal realizations. Players can still randomize over actions, but they must do it by adding noise to the experiments they perform, noise that others can acquire information about.

For strong robustness, a strong version of the revelation principle holds: without loss of generality, we can focus on experiments whose outcomes are pure-action recommendations. To state the result, let $\mathcal{P}_{A_i}$ be the set of i's feasible experiments $P_i \in \mathcal{P}_i$ whose outcomes are pure-action recommendations: for all $\theta \in \Theta$ and $x_{-i} \in X_{-i}$, $P_i(A_i|\theta, x_{-i}) = 1$. We adopt the notation $\mathcal{P}_{A_{-i}} = \prod_{j \neq i} \mathcal{P}_{A_j}$ and $\mathcal{P}_A = \mathcal{P}_{A_i} \times \mathcal{P}_{A_{-i}}$.

Proposition 3. *Under Assumptions 1-3, a probability distribution $\rho \in \Delta_\pi(\Theta \times A)$ is strongly robust to information acquisition if and only if there is $P \in \mathcal{P}_A$ such that the following conditions are satisfied:*

(i) For all $\theta \in \Theta$ and $a \in A$,

$$\rho(a|\theta) = P_i(a_i|\theta, a_{-i})\rho_{-i}(a_{-i}|\theta)$$

(ii) For all $i \in N$, P_i is an optimal solution of

$$\max_{P_i' \in \mathcal{P}_{A_i}} \sum_{\theta, a} u_i(\theta, a) P_i'(a_i|\theta, a_{-i})\rho_{-i}(\theta, a_{-i}) - C_i(P_i', \rho_{-i}).$$

The next example makes concrete the difference between robustness (Definition 2) and strong robustness (Definition 5) to information acquisition.

Example 4. Let Θ be a singleton, and let $\alpha \in \mathcal{A}$ be a Nash equilibrium of the corresponding complete-information game. For every $a \in A$, define

$$\rho(a) = \prod_{i \in N} \alpha_i(a_i).$$

We claim that, under Assumption 1, ρ is robust to information acquisition. Let P be a profile of uninformative experiments, and let $\mu \in \Delta_\pi(\Theta \times X)$ be the product distribution induced by P. For every $i \in N$ and $x_i \in X_i$, define

$$\sigma_i(a_i|x_i) = \alpha_i(a_i).$$

It is easy to see that μ, P, and σ satisfy (i)-(iii) of Definition 5. In particular, uninformative experiments are optimal: all experiments have zero value (since Θ is a singleton and a_{-i} is independent of x_{-i}) and uninformative experiments are the least costly (because of Assumption 1). Overall, we conclude that ρ is robust to information acquisition. ▲

The example shows that under broad conditions on costs, every Nash equilibrium of a complete-information game is robust to information acquisition. But it may not be *strongly* robust:

Example 4 (Continued)**.** Absent uncertainty about θ, an experiment $P_i : X_{-i} \to \Delta(X_i)$ is *fully informative* if it reduces all uncertainty about θ and x_{-i}. Assume that for some player i, there exists a fully informative experiment P_i such that

$$C_i(P_i, \rho_{-i}) < \sum_{a_{-i}} \rho_{-i}(a_{-i}) \max_{a_i} u_i(a_i, a_{-i}) - \sum_{a} u_i(a)\rho(a) \tag{1}$$

where $\rho \in \Delta_\pi(\Theta \times A)$ is, as above, the action distribution induced by a Nash equilibrium $\alpha \in \mathcal{A}$.

We claim that ρ is *not* strongly robust to information acquisition. Since i's opponents cannot privately randomize, a fully informative experiment reduces all uncertainty not only about x_{-i}, but also about a_{-i}. Thus, by (1), player i prefers to acquire full information and best respond to the opponents' realized actions, rather than to acquire no information and best respond to the opponents' mixed actions. This shows that ρ, where actions are independent, is not strongly robust to information acquisition. ▲

As the example highlights, strong robustness to information acquisition can be a demanding notion. One may wonder whether there is any interesting application of such a notion. In Section 3, we will show that there is an important class of games—potential games—where unrestricted information acquisition has a neat characterization for rationally inattentive players.

The literature on incomplete-information games highlights two kinds of uncertainty that players may face: *fundamental uncertainty* (i.e., uncertainty about the state) and *strategic uncertainty* (i.e., uncertainty about others' actions). The interaction between fundamental and strategic uncertainty is crucial in many settings (global games, robust mechanism design, etc.).

Both fundamental and strategic uncertainty provide an incentive to acquire information. In Definition 5, players can reduce both fundamental and strategic uncertainty. In Definition 2, players can reduce strategic uncertainty only insofar as it is determined by first-order and higher-order fundamental uncertainty.

2.4 Existence

When Θ is a singleton, a robust equilibrium exists under broad conditions on the cost of information; see Example 4. When Θ is not a singleton, the existence of a robust equilibrium depends on the set of feasible experiments and on the properties of the cost of information. Next is an example of non-existence.

Example 5. The state could be either high or low: $\Theta = \{h, l\}$ with $\pi(h) = \pi(l) = 1/2$. There are two players: player 1 wants to match the state, and player 2 wants to match player 1's action. Given $A_1 = A_2 = \{h, l\}$, define u_1 and u_2 by

$$u_1(\theta, a_1, a_2) = \begin{cases} 1 & \text{if } a_1 = \theta \\ 0 & \text{otherwise,} \end{cases} \quad \text{and} \quad u_2(\theta, a_1, a_2) = \begin{cases} 1 & \text{if } a_1 = a_2 \\ 0 & \text{otherwise.} \end{cases}$$

As in applications of rational inattention, we assume that the cost of information is based on mutual information. For $P_i \in \mathcal{P}_i$ and $\mu_{-i} \in \Delta(\Theta \times X_{-i})$, we define

$$C_i(P_i, \mu_{-i}) = \lambda I(P_i, \mu_{-i})$$

where $\lambda > 0$ is a scale factor that parameterizes the marginal cost of information, and $I(P_i, \mu_{-i})$ is the mutual information of x_i and (θ, x_{-i}); see Example 3 for the definition of mutual information. We also assume that all experiments are feasible and that signal spaces are rich: for every player i, $\mathcal{P}_i$ is the set of all measurable functions $P_i : \Theta \times X_{-i} \to \Delta(X_i)$, and $\mathcal{A}_i$ is a Borel subset of X_i.

We claim that for λ sufficiently small, there is no $\rho \in \Delta_\pi(\Theta \times A)$ that is robust to information acquisition. Using tools from rational inattention (e.g., Matejka and McKay, 2015), one can verify that for ρ to be robust, it must be that (i) a_1 is conditionally independent of a_2 given θ, and (ii) a_2 is conditionally independent of θ given a_1. Condition (i) reflects the fact that player 1 only cares about the state; so she disregards any additional costly information about player 2. Condition (ii) reflects the fact that player 2 only cares about player 1; so she disregards any additional costly information about the state.

It is easy to verify that (i) and (ii) coexist only if a_2 is independent of θ and a_1. If λ is sufficiently small, this cannot happen: when information is sufficiently cheap, player 1 makes her action dependent on the state, and player 2 makes his action dependent on player 1's action. This shows that when λ is sufficiently small, there is no equilibrium that is robust to information acquisition. ▲

The example shows that an equilibrium that is robust to information acquisition may not exist. The concept is, nevertheless, useful in a variety of circumstances, as we show next in applications. General conditions for the existence and non-existence of equilibria that are robust to information acquisition are largely an open question.

3 Application to potential games

We present an application to rational inattention in games. We consider *potential games* (Monderer and Shapley, 1996). In a potential game, a function $v : \Theta \times A \to \mathbb{R}$, called *potential*, summarizes players' incentives to take actions: for all $\theta \in \Theta$, $i \in N$, $a_i, a_i' \in A_i$, and $a_{-i} \in A_{-i}$,

$$u_i(\theta, a_i, a_{-i}) - u_i(\theta, a_i', a_{-i}) = v(\theta, a_i, a_{-i}) - v(\theta, a_i', a_{-i}).$$

The class of potential games, although restrictive, includes many games that are relevant for information economics—among others, common interest games, congestion games, games on networks, and linear-quadratic games.

As is standard in rational inattention (e.g., Matejka and McKay, 2015), we assume that the cost of information is based on mutual information. For $P_i \in \mathcal{P}_i$ and $\mu_{-i} \in \Delta(\Theta \times X_{-i})$, we define

$$C_i(P_i, \mu_{-i}) = \lambda I(P_i, \mu_{-i})$$

where $\lambda > 0$ is a scale factor that parametrizes the marginal cost of information, and $I(P_i, \mu_{-i})$ is the mutual information of x_i and (θ, x_{-i}); see Example 3 for the definition of mutual information. We also assume that signal spaces are rich: for every player i, $\mathcal{A}_i = \Delta(A_i)$ is a Borel subset of X_i.

Previous work on rational inattention in games has focused on *independent* information acquisition: for every player i and every feasible experiment $P_i \in \mathcal{P}_i$, the

	Invest	*Don't*
Invest	θ, θ	$\theta - 1, 0$
Don't	$0, \theta - 1$	$0, 0$

Figure 1: An investment game.

distribution of x_i depends only on θ and not on x_{-i}. In this literature, a central paper is Yang (2015), who shows that severe coordination problems may arise among rationally inattentive players.

Yang considers the investment game depicted in Figure 1, which is a potential game. He assumes that information acquisition is independent and shows that a large multiplicity of equilibria arises. In particular, as λ goes to zero, coordination may occur on *any* equilibrium of the complete-information game. Yang's results are thought-provoking and cast doubt on whether a meaningful analysis of rational inattention in games is possible.

An element that is missing from Yang's analysis is the possibility for the players to acquire *correlated* information—that is, information about each other's information. Intuition suggests that correlated information may play an important role in coordination games, and, more broadly, in strategic settings. Our model allows us to incorporate this ingredient into the analysis.

Here, instead of assuming that information acquisition is independent, we assume that all experiments are feasible: for every player i, $\mathcal{P}_i$ is the set of *all* measurable functions from $\Theta \times X_{-i}$ to $\Delta(X_i)$. Thus, information acquisition is *unrestricted*: the players can acquire information about the state and each other's information in a flexible way.

In stark contrast with Yang's indeterminacy result, unrestricted information acquisition admits a sharp characterization:

Proposition 4. *Fix a potential function v and a scale factor λ. A probability distribution $\rho \in \Delta_\pi(\Theta \times A)$ is robust (Definition 2) to unrestricted information acquisition if and only if there are $\mu_i \in \Delta(\mathcal{A}_i)$, with $i \in N$, that satisfy the following conditions:*

(i) For all $\theta \in \Theta$ and $a \in A$,

$$\rho(a|\theta) = \frac{\int_{\mathcal{A}} \prod_{i=1}^n \alpha_i(a_i) e^{v(\theta,\alpha)/\lambda} \prod_{i=1}^n \mathrm{d}\mu_i(\alpha_i)}{\int_{\mathcal{A}} e^{v(\theta,\alpha)/\lambda} \prod_{i=1}^n \mathrm{d}\mu_i(\alpha_i)}.$$

(ii) $(\mu_1, \ldots, \mu_n)$ *is a pure Nash equilibrium of a common-interest game where players' payoff function* $V : \prod_{i=1}^n \Delta(\mathcal{A}_i) \to \mathbb{R}$ *is given by*

$$V(\nu_1, \ldots, \nu_n) = \sum_\theta \pi(\theta) \log \left(\int_\mathcal{A} e^{v(\theta,\alpha)/\lambda} \prod_{i=1}^n \mathrm{d}\nu_i(\alpha_i) \right).$$

Moreover, for all $i \in N$ *and* $a_i \in A_i$*,*

$$\sum_{\theta, a_{-i}} \rho(\theta, a_i, a_{-i}) = \int_{\mathcal{A}_i} \alpha_i(a_i)\, \mathrm{d}\mu_i(\alpha_i).$$

Thus, rationally inattentive players follow a *joint logit rule*. By (i), the distribution over actions $\rho \in \Delta_\pi(\Theta \times A)$ is generated by a distribution over mixed actions $\mu \in \Delta_\pi(\Theta \times \mathcal{A})$ given by

$$\mu(B|\theta) = \frac{\int_B e^{v(\theta,\alpha)/\lambda} \prod_{i=1}^n \mathrm{d}\mu_i(\alpha_i)}{\int_\mathcal{A} e^{v(\theta,\alpha)/\lambda} \prod_{i=1}^n \mathrm{d}\mu_i(\alpha_i)}.$$

The log-likelihood ratio of mixed-action profiles α and α' is

$$\ln \frac{\mathrm{d}\mu(\alpha|\theta)}{\mathrm{d}\mu(\alpha'|\theta)} = \frac{v(\theta, \alpha) - v(\theta, \alpha')}{\lambda},$$

which is a logistic regression.

The joint logit rule is pinned down by v and λ, which are primitives of the model, and by $\mu_1, \ldots, \mu_n$, which are endogenous objects. For every player i, μ_i is the marginal distribution over mixed actions. By (ii), the profile $(\mu_1, \ldots, \mu_n)$ is a "pure" Nash equilibrium of an auxiliary potential game where the players have a common payoff function V. Choices in the auxiliary game are *independent*; yet, remarkably, they determine the correlation structure in the original game.

As a corollary of Proposition 4, we obtain a characterization of strong robustness to information acquisition.

Corollary 1. *Fix a potential function* v *and a scale factor* λ*. A probability distribution* $\rho \in \Delta_\pi(\Theta \times A)$ *is strongly robust (Definition 5) to unrestricted information acquisition if and only if there are* $\alpha_i \in \mathcal{A}_i$*, with* $i \in N$*, that satisfy the following conditions:*

(i) For all $\theta \in \Theta$ and $a \in A$,

$$\rho(a|\theta) = \frac{e^{v(\theta,a)/\lambda} \prod_{i \in N} \alpha_i(a_i)}{\sum_{a'} e^{v(\theta,a')/\lambda} \prod_{i \in N} \alpha_i(a_i)}.$$

(ii) $(\alpha_1, \ldots, \alpha_n)$ is a pure Nash equilibrium of a common-interest game where players' payoff function $V : \mathcal{A} \to \mathbb{R}$ is given by

$$V(\alpha'_1, \ldots, \alpha'_n) = \sum_{\theta} \pi(\theta) \log \left(\sum_{a} e^{v(\theta,a)/\lambda} \prod_{i \in N} \alpha'_i(a_i) \right).$$

Moreover, for all $i \in N$ and $a_i \in A_i$,

$$\sum_{\theta, a_{-i}} \rho(\theta, a_i, a_{-i}) = \alpha_i(a_i).$$

Corollary 1 follows more or less immediately from Proposition 4. For every player i, take μ_i as in Proposition 4. Strong robustness is the special case in which μ_i puts positive probability only on pure actions:

$$\mu_i(\{\delta_{a_i} : a_i \in A_i\}) = 1$$

where $\delta_{a_i} \in \mathcal{A}_i$ is the Dirac measure concentrated on a_i. Then, it is clear that (i) and (ii) in Corollary 1 correspond to (i) and (ii) in Proposition 4.

We also obtain the result that *there exists* an equilibrium that is strongly robust (hence, robust) to information acquisition:

Corollary 2. *For every potential function v and scale factor λ, there exists a probability distribution $\rho \in \Delta_\pi(\Theta \times A)$ that is strongly robust (Definition 5) to unrestricted information acquisition.*

The result follows from the existence of an equilibrium in the auxiliary potential game in Corollary 1-(ii). As Example 4 highlights, the existence of a strongly robust equilibrium is demanding. Here, one can prove a partial converse to Corollary 2: for every utility profile $(u_i)_{i \in N}$ and scale factor λ, if there exists a probability distribution $\rho \in \Delta_\pi(\Theta \times A)$ that is strongly robust to unrestricted information acquisition, then the game admits a potential function on the support of ρ.[17]

[17] A formal proof of the result was included in a previous draft of the paper and is available from

A special case of interest is the negligible cost of information—that is, $\lambda \to 0$. Under independent information acquisition, Yang (2015) finds that anything goes. When information acquisition is unrestricted, we obtain the opposite result, limit uniqueness:

Corollary 3. *Fix a potential function v. Let $a^* \in A^*$ and $\theta^* \in \Theta^*$ satisfy the following conditions:*

(i) The quantity $v(\theta^*, a)$ is uniquely maximized at $a = a^*$.

(ii) For each player i, a_i^* is the unique equilibrium action in some state θ.

For every $\lambda > 0$, let $\rho^\lambda \in \Delta_\pi(\Theta \times A)$ be robust (Definition 2) to unrestricted information acquisition. Then

$$\lim_{\lambda \to 0} \rho^\lambda(a^*|\theta^*) = 1.$$

Thus, when information acquisition is correlated and the cost of information is negligible, a *unique* outcome arises: coordination occurs on the action profile a^* that maximizes the potential. The result holds under a richness condition: for every player i, a_i^* is the unique equilibrium action in some state—Yang (2015) makes a similar richness assumption. Corollary 3 suggests an application of our model to equilibrium selection, which we carry out in the next section.

Proposition 4 proposes a many player extension of the individual logit rule for rationally inattentive agents as in Matejka and McKay (2015). Any single-agent decision problem trivially is a potential game; if N is a singleton, then Proposition 4 is a restatement of Matejka and McKay's main result. Matejka and McKay provide essential tools for applications of rational inattention; we adapt and extend their results to games.

For the proof of Proposition 4, we start from player-by-player *individual* logit rules à la Matejka and McKay, and show that they generate a *joint* logit rule, leveraging on the potential structure of the game. The interested reader can find detailed arguments in the appendix.

Logit rules and potential games show up in evolutionary game theory. To illustrate the relationship to our results, take $(\alpha_1, \ldots, \alpha_n)$ as in Corollary 1. Suppose that Θ is

the author upon request.

	Left	*Right*
Left	θ_{ll}, θ_{ll}	θ_{lr}, θ_{rl}
Right	θ_{rl}, θ_{lr}	θ_{rr}, θ_{rr}

Figure 2: Symmetric 2×2 games.

a singleton and all α_i are uniform distributions. Then

$$\rho(a) = \frac{e^{v(a)/\lambda} \prod_{i=1}^{n} \alpha_i(a_i)}{\sum_{a'} e^{v(a')/\lambda} \prod_{i=1}^{n} \alpha_i(a_i)} = \frac{e^{v(a)/\lambda}}{\sum_{a'} e^{v(a')/\lambda}}.$$

One can view such $\rho \in \Delta(A)$ as the stationary distribution of a *logit best-response dynamic* (Blume, 1993), which is prominent in evolutionary game theory. Here, the α_i are endogenous objects that, except for knife-edge cases, are not uniform.

4 Application to global games

We present an application to equilibrium selection in coordination games: we revisit global games and provide a new rationale for risk dominant equilibria.

We consider symmetric 2×2 games as depicted in Figure 2. When $\theta = (\theta_{ll}, \theta_{rl}, \theta_{lr}, \theta_{rr})$ is common knowledge, there are generically one or three equilibria; in the latter case, there is one weak equilibrium in mixed actions and two strict equilibria in pure actions. Thus, without loss of generality, we may focus on the case in which $(Left, Left)$ or $(Right, Right)$ or both are strict equilibria: we assume that for every $\theta \in \Theta$,

$$\max \{\theta_{ll} - \theta_{rl}, \theta_{rr} - \theta_{lr}\} > 0 \quad \text{and} \quad \min \{\theta_{ll} - \theta_{rl}, \theta_{rr} - \theta_{lr}\} \neq 0.$$

As is standard in global games, we assume that there are states in which *Left* and *Right* are dominant actions: there are $\theta, \theta' \in \Theta$ such that

$$\min \{\theta_{ll} - \theta_{rl}, \theta_{lr} - \theta_{rr}\} > 0 \quad \text{and} \quad \min \{\theta'_{rl} - \theta'_{ll}, \theta'_{rr} - \theta'_{lr}\} > 0.$$

Such states may have arbitrarily small probability. Thus, we may view dominance regions as richness conditions on the space of uncertainty.

For the cost of information, we consider a popular generalization of mutual information based on f-divergences. Let $f : [0,\infty) \to \mathbb{R}$ be a continuous function that is strictly convex and continuously differentiable on $(0,\infty)$. For $P_i \in \mathcal{P}_i$ and $\mu_{-i} \in \Delta(\Theta \times X_{-i})$, we define

$$C_i(P_i, \mu_{-i}) = \lambda\left(I_f(P_i, \mu_{-i}) - f(1)\right)$$

where $\lambda > 0$ is a scale factor that parameterizes the marginal cost of information, and $I_f(P_i, \mu_{-i})$ is the f-divergence between the joint distribution of (θ, x) and the product of the marginal distributions of x_i and (θ, x_{-i}); see Example 3 for the definition of I_f.[18] We also assume that signal spaces are rich: for every player i, $\mathcal{A}_i = \Delta(A_i)$ is a Borel subset of X_i.

Let all experiments be feasible: for every player i, $\mathcal{P}_i$ is the set of *all* measurable functions from $\Theta \times X_{-i}$ to $\Delta(X_i)$. Thus, information acquisition is *unrestricted*: the players can acquire information about the state and each other's information in a flexible way.

For every scale factor $\lambda > 0$, let $\rho^\lambda \in \Delta_\pi(\Theta \times A)$ be robust to unrestricted information acquisition. We analyze the limit case of the *negligible* cost of information: we ask what happens to ρ^λ as λ converges to zero. We interpret the case of $\lambda \to 0$ as a perturbation of common knowledge. The nature of the perturbation is *endogenous*: it depends on players' incentives to acquire information. The next existence result makes sure that the problem we set out to solve is well-defined.

Lemma 1. *For every scale factor $\lambda > 0$, there exists $\rho^\lambda \in \Delta_\pi(\Theta \times A)$ that is robust (Definition 2) to unrestricted information acquisition.*

The notion of *risk dominance* will be central to our analysis. According to Harsanyi and Selten (1988), $Left$ risk dominates *Right* at θ if

$$\theta_{ll} - \theta_{rl} > \theta_{rr} - \theta_{lr}.$$

Right risk dominates $Left$ if the reverse strict inequality holds.

For an intuition behind risk dominance, take the perspective of a player who is unsure about the opponent's action. It is easy to verify that $\theta_{ll} - \theta_{rl} > \theta_{rr} - \theta_{lr}$ if and only if $Left$ is a best response to a broader range of beliefs than *Right*; one could deem playing $Left$ "less risky" than playing *Right*.

[18] Subtracting $f(1)$ guarantees that $C_i(P_i, \mu_{-i}) \geq 0$ (see, e.g., Liese and Vajda, 2006, Theorem 5).

The notion of risk dominance is standard in the literature. Next we introduce a related comparative notion. We say that $Left$ is *less risky* at θ than at θ' if

$$\theta_{ll} - \theta_{rl} > \theta'_{ll} - \theta'_{rl} \qquad \text{and} \qquad \theta_{rr} - \theta_{lr} < \theta'_{rr} - \theta'_{lr}.$$

Harsanyi and Selten (1988) argue that when a game has multiple equilibria, the players "should" coordinate on the risk dominant action. Carlsson and van Damme (1993) formalize their intuition by identifying a class of perturbations of common knowledge that select the risk dominant equilibrium. A number of subsequent works have revisited Carlsson and van Damme's results, both from an ex-ante perspective (e.g., Kajii and Morris, 1997) and an interim perspective (e.g., Weinstein and Yildiz, 2007). A common theme is that minute details of the perturbation of common knowledge determine whether or not the risk dominant equilibrium is selected. This raises the question of what perturbations of common knowledge are more likely to occur. The next proposition provides an answer.

Proposition 5. *Assume* $\lim_{t\to 0} f'(t) = -\infty$. *For every* $\lambda > 0$, *let* $\rho^{\lambda} \in \Delta_{\pi}(\Theta \times A)$ *be robust (Definition 2) to unrestricted information acquisition. For* $\rho = \lim_{\lambda\to 0} \rho^{\lambda}$, *the following conditions hold:*

(i) If $Left$ *is risk dominant at* θ, *then*

$$\rho(Left, Left|\theta) \geq \rho(Left, Left).$$

If $Right$ *is risk dominant at* θ, *then the reverse inequality holds.*

(ii) If $Left$ *is less risky at* θ *than at* θ', *then*

$$\rho(Left, Left|\theta) \geq \rho(Left, Left|\theta').$$

Figure 3 provides a graphical representation of the content of the proposition. The horizontal axis orders the states according to the comparative notion of riskiness introduced above: θ precedes θ' if $Left$ is less risky at θ than at θ'.[19] The state $\hat{\theta}$ is the risk dominance cutoff: if $\theta < \hat{\theta}$, then $Left$ is risk dominant; if $\theta > \hat{\theta}$, then $Right$

[19] In general, such ordering is not complete: there could be two states θ and θ' such that $Left$ is neither less nor more risky at θ than at θ'. For ease of exposition, Figure 3 considers the special case in which the ordering is complete.

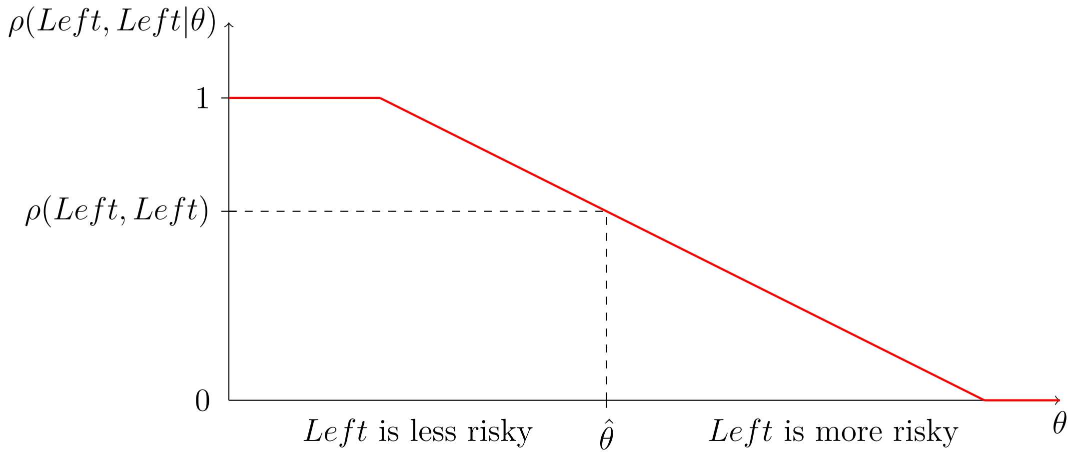

Figure 3: A graphical representation of Proposition 5.

is risk dominant. The vertical axis plots the conditional probability that the players coordinate on $Left$, as a function of the state.

By (ii), $\rho(Left, Left|\theta)$ is decreasing in θ: the more risky $Left$ is, the more likely the players coordinate on *Right*. The quantity $\rho(Left, Left)$ is the marginal probability that the players coordinate on $Left$. By (i), $\rho(Left, Left|\theta)$ is larger than $\rho(Left, Left)$ when $Left$ is risk dominant at θ, and smaller when *Right* is risk dominant at θ. Extreme values of θ correspond to dominance regions where $\rho(Left, Left|\theta)$ is either zero or one.

These results hold under the hypothesis that f is infinitely steep at zero: as $t \to 0$, $f'(t) \to -\infty$. This guarantees that the marginal cost of acquiring full information is infinite. Otherwise, acquiring full information would be optimal for λ sufficiently small, trivially leading to complete information and multiple equilibria.

Overall, Proposition 5 suggests that when information is endogenous and cheap, it is more likely that players coordinate on the risk dominant action. There is a simple intuition behind this result. Suppose that $Left$ is the risk dominant action. Acquiring information, a player wants to minimize the probability of miscoordination. The probability of miscoordination is determined by the probability of choosing *Right* when the opponent chooses $Left$, and by the probability of choosing $Left$ when the opponent chooses *Right*. By analogy with statistical decision theory, we can call these two mistakes a *type I error* and a *type II error*. So the question is: choosing an experiment, what is the player's priority: minimizing the type I error or the type

II error? Intuitively, the priority is to minimize the type I error, because erroneously choosing $Left$ is less risky than erroneously choosing *Right*, $Left$ being the risk dominant action. This pushes the players toward coordination on the risk dominant action.

We see Proposition 5 as supporting the risk dominance selection rule from global games. We emphasize, however, that Proposition 5 *does not* state that the risk dominant action is played with probability one. As λ goes to zero, the players could coordinate on a correlated equilibrium where with some probability, they take the risk dominant action, and with the remaining probability, they take the risk dominated action. A key feature of Proposition 5 is that the probability with which the players take the risk dominant action is "increasing" in the state—increasing according to the comparative notion of riskiness described above. In applications, lack of monotonicity in the state is one of the most compelling critiques of multiple equilibrium narratives such as the theory of sunspots (Morris and Shin, 2001); our results reiterate such critiques.[20]

When I_f is mutual information, the risk dominant action *is* played with probability one in the limit. The result follows from Corollary 3 in the previous section. It is easy to verify that the games represented in Figure 2 are potential games, and the risk dominant action is the maximizer of the potential function.

We have used a model of costly information acquisition to generate an endogenous perturbation of common knowledge; we borrowed the approach from Yang (2015). As discussed in the previous section, Yang (2015) studies a coordination game where players decide whether or not to invest in a risky project; his investment game belongs to the class of symmetric games we have analyzed in this section (cfr. Figures 1 and 2). Assuming that information acquisition is *independent* and costs are based on mutual information, Yang (2015) shows that multiple equilibria arise when information is cheap.

The basic intuition behind Yang's multiplicity result is that a coordination problem in taking actions may translate into a coordination problem in acquiring information. Our analysis shows that this intuition, although compelling at first glance, crucially relies on information acquisition being independent. When players can learn

[20] As Morris and Shin (2001, p. 139) put it, "[the theory of sunspots] runs counter to our intuition that bad fundamentals are somehow 'more likely' to trigger a financial crisis, or to tip the economy into recession. In other words, sunspot explanations do not provide a basis for exploring the correlation between the underlying fundamentals and the resultant economic outcomes."

what others know, other considerations come into play that may lead to equilibrium uniqueness. For example, the trade-off in minimizing the probability of miscoordination that we discussed above.

Morris and Yang (Forthcoming) revisit the results of Yang (2015) from a different perspective. They maintain the hypothesis that information acquisition is independent, but consider a class of cost functions that do not nest mutual information. In particular, they assume that it is costly to distinguish nearby states. They show that when information is cheap, a unique outcome arises: as in global games, players coordinate on the risk dominant action. Their approach is compelling in applications where the state space has an obvious topology. An advantage of our methods is that they apply to arbitrary state spaces.

In a follow-up to a previous version of this paper, Hoshino (2018) adopts the model of unrestricted information acquisition to study equilibrium selection in coordination games. He finds that *any* selection of complete-information equilibria can be sustained by some ad hoc cost functions. One limitation of this interesting result is that the cost functions Hoshino constructs are somewhat arbitrary. By contrast, the cost functions we consider in this section have widespread applications in economics, statistics, and information theory. Our takeaway from Hoshino (2018) is that the relationship between costly information acquisition and coordination problems deserves further investigation.

5 Conclusion

We have developed a model of correlated information acquisition in games. The model has allowed us to study players' incentive to learn what others know. In applications to rational inattention and global games, we have shown that the model is tractable and delivers interesting predictions: rationally inattentive players follow a joint logit rule; players endogenously coordinate on the risk dominant action when information is cheap. Both results crucially rely on information acquisition being correlated. When information acquisition is independent, the behavior of rationally inattentive players is undetermined and severe coordination problems arise in information acquisition. This shows the power of players' incentive to learn what others know.

A Proofs

Proof of Proposition 2. "If." Let $\mu \in \Delta_\pi(\Theta \times \mathcal{A})$ and $P \in \mathcal{P}_\mathcal{A}$ satisfy (i)-(iii) of Proposition 2. For every player i, let $\sigma_i \in \Sigma_i$ be an action rule such that for all $\alpha_i \in \mathcal{A}_i$, $\sigma_i(a_i|\alpha_i) = \alpha_i(a_i)$.

Since μ and P satisfy (i)-(ii) of Proposition 2, we see that μ, P, and σ satisfy (i)-(ii) of Definition 2. To prove (iii) of Definition 2, take an experiment $\tilde{P}_i \in \mathcal{P}_i$ and an action rule $\tilde{\sigma}_i \in \Sigma_i$. Define the experiment P_i' by

$$P_i'(B_i|\theta, x_{-i}) = \tilde{P}_i\left(\tilde{\sigma}_i^{-1}(B_i)|\theta, x_{-i}\right).$$

The experiment P_i' is a garbling of P_i. Thus, by Assumption 3, $P_i' \in \mathcal{P}_i$ and $C_i(P_i', \mu_{-i}) \leq C_i(\tilde{P}_i, \mu_{-i})$. Moreover, for all $\theta \in \Theta$ and $x_{-i} \in X_{-i}$, $P_i'(\mathcal{A}_i|\theta, x_{-i}) = 1$. Therefore, $P_i' \in \mathcal{P}_{\mathcal{A}_i}$. Using (iii) of Proposition 2, we obtain that

$$\begin{aligned}
&\int_{\Theta\times X} U_i(\theta, \tilde{\sigma}_i(x_i), \sigma_{-i}(x_{-i}))\,\mathrm{d}\tilde{P}_i(x_i|\theta, x_{-i})\,\mathrm{d}\mu_{-i}(\theta, x_{-i}) - C_i(\tilde{P}_i, \mu_{-i})\\
&\leq \int_{\Theta\times\mathcal{A}_i\times X_{-i}} U_i(\theta, \alpha_i, \sigma_{-i}(x_{-i}))\,\mathrm{d}P_i'(\alpha_i|\theta, x_{-i})\,\mathrm{d}\mu_{-i}(\theta, x_{-i}) - C_i(P_i', \mu_{-i})\\
&= \int_{\Theta\times\mathcal{A}} U_i(\theta, \alpha)\,\mathrm{d}P_i'(\alpha_i|\theta, \alpha_{-i})\,\mathrm{d}\mu_{-i}(\theta, \alpha_{-i}) - C_i(P_i', \mu_{-i})\\
&\leq \int_{\Theta\times\mathcal{A}} U_i(\theta, \alpha)\,\mathrm{d}P_i(\alpha_i|\theta, \alpha_{-i})\,\mathrm{d}\mu_{-i}(\theta, \alpha_{-i}) - C_i(P_i, \mu_{-i})\\
&= \int_{\Theta\times X} U_i(\theta, \sigma(x))\,\mathrm{d}P_i(x_i|\theta, x_{-i})\,\mathrm{d}\mu_{-i}(\theta, x_{-i}) - C_i(P_i, \mu_{-i}).
\end{aligned}$$

It follows that (iii) of Definition 2 holds.

"Only if." Let $\rho \in \Delta_\pi(\Theta \times A)$ be robust to information acquisition. Let $\mu \in \Delta_\pi(\Theta \times X)$, $P \in \mathcal{P}$, and $\sigma \in \Sigma$ as in Definition 2. Let $\nu \in \Delta_\pi(\Theta \times \mathcal{A})$ be the push-forward of μ under the function $x \mapsto \sigma(x)$. For every player i, let $K_{-i} : \Theta \times X_{-i} \to \Delta(X_{-i})$ be a version of the (regular) conditional μ_{-i}-probability of x_{-i} given θ and $\sigma_{-i}(x_{-i})$. Define the experiment Q_i by

$$Q_i(B_i|\theta, x_{-i}) = \int_{X_{-i}} P_i(\sigma_i^{-1}(B_i)|\theta, x'_{-i})\,\mathrm{d}K_{-i}(x'_{-i}|\theta, x_{-i}).$$

Claim 1. $Q_i \in \mathcal{P}_{\mathcal{A}_i}$ and $C_i(Q_i, \nu_{-i}) \leq C_i(P_i, \mu_{-i})$.

Proof of the claim. Define the experiment Q_i° by

$$Q_i^\circ(B_i|\theta, x_{-i}) = \int_{X_{-i}} P_i(B_i|\theta, x'_{-i})\, \mathrm{d}K_{-i}(x'_{-i}|\theta, x_{-i}).$$

For all $\theta \in \Theta$ and $B_{-i} \subseteq X_{-i}$,

$$\mu_{-i}(B_{-i}|\theta) = \int_{X_{-i}} K_{-i}(B_{-i}|\theta, x_{-i})\, \mathrm{d}\nu_{-i}(x_{-i}|\theta).$$

Thus $(Q_i^\circ, \nu_{-i}) \trianglelefteq (Q_i, \mu_{-i})$. In addition, $Q_i \preceq Q_i^\circ$ and for every $\theta \in \Theta$ and $x_{-i} \in X_{-i}$, $Q_i(\mathcal{A}_i|\theta, x_{-i}) = 1$. It follows from Assumptions 3 and 4 that $Q_i^\circ \in \mathcal{P}_i$, $Q_i \in \mathcal{P}_{\mathcal{A}_i}$, and

$$C_i(Q_i, \nu_{-i}) \leq C_i(Q_i^\circ, \nu_{-i}) \leq C_i(P_i, \mu_{-i}),$$

as desired. □

Since μ, P, and σ satisfy (i)-(ii) of Definition 2, we see that ν and Q satisfy (i)-(ii) of Proposition 2. To prove (iii) of Proposition 2, note that for every $Q'_i \in \mathcal{P}_{\mathcal{A}_i}$,

$$\begin{aligned}
&\int_{\Theta\times\mathcal{A}} U_i(\theta, \alpha)\, \mathrm{d}Q'_i(\alpha_i|\theta, \alpha_{-i})\, \mathrm{d}\nu_{-i}(\theta, \alpha_{-i}) - C_i(Q'_i, \nu_{-i}) \\
&= \int_{\Theta\times\mathcal{A}_i\times X_{-i}} U_i(\theta, \alpha_i, \sigma_{-i}(x_{-i}))\, \mathrm{d}Q'_i(\alpha_i|\theta, \sigma_{-i}(x_{-i}))\, \mathrm{d}\mu_{-i}(\theta, x_{-i}) - C_i(Q'_i, \nu_{-i}) \\
&\leq \max_{\sigma'_i} \int_{\Theta\times X} U_i(\theta, \sigma'_i(x_i), \sigma_{-i}(x_{-i}))\, \mathrm{d}Q'_i(x_i|\theta, \sigma_{-i}(x_{-i}))\, \mathrm{d}\mu_{-i}(\theta, x_{-i}) - C_i(Q'_i, \nu_{-i}).
\end{aligned}$$

Define the experiment P'_i by

$$P'_i(B_i|\theta, x_{-i}) = Q'_i(B_i|\theta, \sigma_{-i}(x_{-i})).$$

Claim 2. $(P'_i, \mu_{-i}) \trianglelefteq (Q'_i, \nu_{-i})$.

Proof of the claim. Define $K'_{-i} : \Theta \times X_{-i} \to \Delta(X_{-i})$ by

$$K_{-i}(B_{-i}|\theta, x_{-i}) = 1_{B_{-i}}\left(\sigma_{-i}(x_{-i})\right).$$

Then

$$P_i'(B_i|\theta, x_{-i}) = \int_{X_{-i}} Q_i'(B_i|\theta, x_{-i}') \,\mathrm{d}K_{-i}(x_{-i}'|\theta, x_{-i})$$
$$\nu_{-i}(B_{-i}|\theta) = \int_{X_{-i}} K_{-i}(B_{-i}|\theta, x_{-i}) \,\mathrm{d}\mu_{-i}(\theta, x_{-i}).$$

The desired result follows. □

It follows from Assumption 4 that $C_i(P_i', \mu_{-i}) \leq C_i(Q_i', \nu_{-i})$. Thus,

$$\begin{aligned}
&\max_{\sigma_i'} \int_{\Theta\times X} U_i(\theta, \sigma_i'(x_i), \sigma_{-i}(x_{-i})) \,\mathrm{d}Q_i'(x_i|\theta, \sigma_{-i}(x_{-i})) \,\mathrm{d}\mu_{-i}(\theta, x_{-i}) - C_i(Q_i', \nu_{-i}) \\
\leq &\max_{\sigma_i'} \int_{\Theta\times X} U_i(\theta, \sigma_i'(x_i), \sigma_{-i}(x_{-i})) \,\mathrm{d}P_i'(x_i|\theta, x_{-i}) \,\mathrm{d}\mu_{-i}(\theta, x_{-i}) - C_i(P_i', \mu_{-i}) \\
\leq &\int_{\Theta\times X} U_i(\theta, \sigma(x)) \,\mathrm{d}P_i(x_i|\theta, x_{-i}) \,\mathrm{d}\mu_{-i}(\theta, x_{-i}) - C_i(P_i, \mu_{-i})
\end{aligned}$$

where the last inequality follows from (iii) of Definition 2. As shown above, $C_i(Q_i, \nu_{-i}) \leq C_i(P_i, \mu_{-i})$. We obtain that

$$\begin{aligned}
&\int_{\Theta\times X} U_i(\theta, \sigma(x)) \,\mathrm{d}P_i(x_i|\theta, x_{-i}) \,\mathrm{d}\mu_{-i}(\theta, x_{-i}) - C_i(P_i, \mu_{-i}) \\
\leq &\int_{\Theta\times \mathcal{A}} U_i(\theta, \alpha) \,\mathrm{d}Q_i(\alpha_i|\theta, \alpha_{-i}) \,\mathrm{d}\nu_{-i}(\theta, \alpha_{-i}) - C_i(Q_i, \nu_{-i}).
\end{aligned}$$

We conclude that (iii) of Proposition 2 holds. ■

Proof of Proposition 3. The proof is analogous to the proof of Proposition 2; we omit the details. ■

Lemma 2. *Suppose that all experiments are feasible and signal spaces are rich enough to contain the sets of mixed actions. For every $i \in N$, $P_i \in \mathcal{P}_{\mathcal{A}_i}$, and $\mu_{-i} \in \Delta_\pi(\Theta \times \mathcal{A}_{-i})$, the following statements are equivalent:*

(i) P_i is an optimal solution of

$$\max_{P_i' \in \mathcal{P}_{\mathcal{A}_i}} \int_{\Theta\times \mathcal{A}} u_i(\theta, \alpha) \,\mathrm{d}P_i'(\alpha_i|\theta, \alpha_{-i}) \,\mathrm{d}\mu_{-i}(\theta, \alpha_{-i}) - \lambda I(P_i', \mu_{-i}).$$

(ii) There is $\mu_i \in \Delta(\mathcal{A}_i)$ that satisfies the following conditions:

(a) *For all* $B_i \subseteq \mathcal{A}_i$ *and* μ_{-i}*-almost all* θ *and* α_{-i}*,*

$$P_i(B_i|\theta, \alpha_{-i}) = \frac{\int_{B_i} e^{u_i(\theta,\alpha)/\lambda} \mathrm{d}\mu_i(\alpha_i)}{\int_{\mathcal{A}_i} e^{u_i(\theta,\alpha)/\lambda} \mathrm{d}\mu_i(\alpha_i)}.$$

(b) *For all* $\alpha_i \in \mathcal{A}_i$*,*

$$\int_{\Theta \times \mathcal{A}_{-i}} \frac{e^{u_i(\theta,\alpha)/\lambda}}{\int_{\mathcal{A}_i} e^{u_i(\theta,\alpha_i',\alpha_{-i})/\lambda} \mathrm{d}\mu_i(\alpha_i')} \mathrm{d}\mu_{-i}(\theta, \alpha_{-i}) \leq 1.$$

Moreover, for all $B_i \subseteq \mathcal{A}_i$*,*

$$\mu_i(B_i) = \int_{\Theta \times \mathcal{A}_{-i}} P_i(B_i|\theta, \alpha_{-i}) \, \mathrm{d}\mu_{-i}(\theta, \alpha_{-i}).$$

Proof. See Denti, Marinacci and Montrucchio (2020), who extend Matejka and McKay (2015) to continuous variables. □

Proof of Proposition 4. "If." Pick $\rho \in \Delta_\pi(\Theta \times A)$ and $\mu_i \in \Delta(\mathcal{A}_i)$, $i \in N$, that satisfy (i) and (ii) of Proposition 4. We claim that ρ is robust to unrestricted information acquisition.

To prove the claim, we use Proposition 2. Define $\mu \in \Delta_\pi(\Theta \times \mathcal{A})$ by

$$\mu(B|\theta) = \frac{\int_B e^{v(\theta,\alpha)/\lambda} \prod_{i=1}^n \mathrm{d}\mu_i(\alpha_i)}{\int_{\mathcal{A}} e^{v(\theta,\alpha)/\lambda} \prod_{i=1}^n \mathrm{d}\mu_i(\alpha_i)}. \tag{2}$$

For every i, define $P_i \in \mathcal{P}_{\mathcal{A}_i}$ by

$$P_i(B_i|\theta, \alpha_{-i}) = \frac{\int_{B_i} e^{v(\theta,\alpha)/\lambda} \mathrm{d}\mu_i(\alpha_i)}{\int_{\mathcal{A}_i} e^{v(\theta,\alpha)/\lambda} \mathrm{d}\mu_i(\alpha_i)}.$$

By construction, (i) and (ii) of Proposition 2 are satisfied. To prove that ρ is robust to unrestricted information acquisition, it remains to show that (iii) of Proposition 2 is satisfied.

By Monderer and Shapley (1996, Lemma 2.10), we have that

$$u_i(\theta, \alpha) - u_i(\theta, \alpha_i', \alpha_{-i}) = v(\theta, \alpha) - v(\theta, \alpha_i', \alpha_{-i}).$$

We deduce that

$$P_i(B_i|\theta,\alpha_{-i}) = \frac{\int_{B_i} e^{u_i(\theta,\alpha)/\lambda} \mathrm{d}\mu_i(\alpha_i)}{\int_{\mathcal{A}_i} e^{u_i(\theta,\alpha)/\lambda} \mathrm{d}\mu_i(\alpha_i)}. \tag{3}$$

In addition, because μ_i is a best response to $\bar{\mu}_{-i} = (\mu_j)_{j\neq i}$ in the auxiliary common-interest game V, we have that

$$\sum_\theta \pi(\theta) \frac{\int_{\mathcal{A}} e^{v(\theta,\alpha_i,\alpha_{-i})/\lambda} \prod_{j\neq i} \mathrm{d}\mu_j(\alpha_j)}{\int_{\mathcal{A}} e^{v(\theta,\alpha')/\lambda} \prod_{i=1}^n \mathrm{d}\mu_i(\alpha_i')} \leq 1, \qquad \alpha_i \in \mathcal{A}_i,$$

which is the first-order condition of the optimization problem

$$\max_{\nu_i} V(\nu_i, \bar{\mu}_{-i}).$$

It follows from (2) that

$$\int_{\Theta\times\mathcal{A}_{-i}} \frac{e^{v(\theta,\alpha)/\lambda}}{\int_{\mathcal{A}_i} e^{v(\theta,\alpha_i',\alpha_{-i})/\lambda} \mathrm{d}\mu_i(\alpha_i')} \mathrm{d}\mu_{-i}(\theta,\alpha_{-i}) \leq 1, \qquad \alpha_i \in \mathcal{A}_i.$$

Using again Monderer and Shapley (1996, Lemma 2.10), we conclude that

$$\int_{\Theta\times\mathcal{A}_{-i}} \frac{e^{u_i(\theta,\alpha)/\lambda}}{\int_{\mathcal{A}_i} e^{u_i(\theta,\alpha_i',\alpha_{-i})/\lambda} \mathrm{d}\mu_i(\alpha_i')} \mathrm{d}\mu_{-i}(\theta,\alpha_{-i}) \leq 1, \qquad \alpha_i \in \mathcal{A}_i. \tag{4}$$

Conditions (3) and (4) correspond to (ii.a) and (ii.b) of Lemma 2. We deduce that (iii) of Proposition 2 holds, as desired.

"Only if." Let $\rho \in \Delta_\pi(\Theta \times A)$ be robust to unrestricted information acquisition. We claim that (i) and (ii) of Proposition 4 hold.

To prove the claim, we use Proposition 2. Take $\mu \in \Delta_\pi(\Theta \times \mathcal{A})$ and $P \in \mathcal{P}_{\mathcal{A}}$ as in Proposition 2. For every player i, let $\mu_i \in \Delta(\mathcal{A}_i)$ be the marginal distribution of μ over i's mixed actions.

By (i) of Lemma 2, we have that

$$P_i(B_i|\theta,\alpha_{-i}) = \frac{\int_{B_i} e^{u_i(\theta,\alpha)/\lambda} \mathrm{d}\mu_i(\alpha_i)}{\int_{\mathcal{A}_i} e^{u_i(\theta,\alpha)/\lambda} \mathrm{d}\mu_i(\alpha_i)}.$$

In addition, by Monderer and Shapley (1996, Lemma 2.10), we have that

$$u_i(\theta,\alpha) - u_i(\theta,\alpha_i',\alpha_{-i}) = v(\theta,\alpha) - v(\theta,\alpha_i',\alpha_{-i}).$$

It follows that

$$P_i(B_i|\theta,\alpha_{-i}) = \frac{\int_{B_i} e^{v(\theta,\alpha)/\lambda} \mathrm{d}\mu_i(\alpha_i)}{\int_{\mathcal{A}_i} e^{v(\theta,\alpha)/\lambda} \mathrm{d}\mu_i(\alpha_i)}.$$

Aggregating across players, simple algebra shows that

$$\frac{\mathrm{d}\mu(\alpha|\theta)}{\prod_{i=1}^n \mathrm{d}\mu_i(\alpha_i)} = \frac{e^{v(\theta,\alpha)/\lambda}}{\int_{\mathcal{A}} e^{v(\theta,\alpha')/\lambda} \prod_{i=1}^n \mathrm{d}\mu_i(\alpha_i')}. \tag{5}$$

We deduce that (i) of Proposition 4 holds.

By (ii) of Lemma 2, we have that

$$\int_{\Theta\times\mathcal{A}_{-i}} \frac{e^{u_i(\theta,\alpha)/\lambda}}{\int_{\mathcal{A}_i} e^{u_i(\theta,\alpha_i',\alpha_{-i})/\lambda} \mathrm{d}\mu_i(\alpha_i')} \mathrm{d}\mu_{-i}(\theta,\alpha_{-i}) \le 1, \qquad \alpha_i \in \mathcal{A}_i.$$

Using again Monderer and Shapley (1996, Lemma 2.10), we obtain that

$$\int_{\Theta\times\mathcal{A}_{-i}} \frac{e^{v(\theta,\alpha)/\lambda}}{\int_{\mathcal{A}_i} e^{v(\theta,\alpha_i',\alpha_{-i})/\lambda} \mathrm{d}\mu_i(\alpha_i')} \mathrm{d}\mu_{-i}(\theta,\alpha_{-i}) \le 1, \qquad \alpha_i \in \mathcal{A}_i$$

In addition, it follows from (5) that

$$\frac{\mathrm{d}\mu_{-i}(\alpha_{-i}|\theta)}{\prod_{j\neq i} \mathrm{d}\mu_j(\alpha_i)} = \frac{\int_{\mathcal{A}_i} e^{v(\theta,\alpha_i',\alpha_{-i})/\lambda} \mathrm{d}\mu_i(\alpha_i')}{\int_{\mathcal{A}} e^{v(\theta,\alpha')/\lambda} \prod_{i=1}^n \mathrm{d}\mu_i(\alpha_i')}.$$

We deduce that

$$\int_{\Theta\times\mathcal{A}_{-i}} \frac{e^{v(\theta,\alpha)/\lambda}}{\int_{\mathcal{A}_i} e^{v(\theta,\alpha_i',\alpha_{-i})/\lambda} \mathrm{d}\mu_i(\alpha_i')} \mathrm{d}\mu_{-i}(\theta,\alpha_{-i}) \le 1, \qquad \alpha_i \in \mathcal{A}_i. \tag{6}$$

Simple calculus shows that (6) is the (necessary and sufficient) first-order condition of the optimization problem

$$\max_{\nu_i} V(\nu_i, \bar{\mu}_{-i})$$

where $\bar{\mu}_{-i} = (\mu_j)_{j\neq i}$. Since this is true for all players i, we conclude that (ii) of Proposition 4 holds. ■

Proof of Corollary 1. "If." Pick $\rho \in \Delta_\pi(\Theta \times A)$ and $\alpha \in \mathcal{A}$ that satisfy (i) and (ii) of Corollary 1. For every player i, choose $\mu_i \in \Delta(\mathcal{A}_i)$ such that

$$\mu_i\left(\delta_{a_i}\right) = \alpha_i(a_i) \tag{7}$$

where $\delta_{a_i} \in \mathcal{A}_i$ is the Dirac measure concentrated on $a_i \in A_i$. Thus, μ_i puts positive probability only on pure actions.

It is easy to see that (i) of Corollary 1 implies (i) of Proposition 4. By (ii) of Corollary 1, α_i is an optimal solution of

$$\max_{\alpha_i'} V(\alpha_i', \alpha_{-i}).$$

The first-order condition of this optimization problem is

$$\sum_\theta \pi(\theta) \frac{\sum_{a'_{-i}} e^{v(\theta,a_i,a'_{-i})/\lambda} \prod_{i\neq j} \alpha_j(a'_j)}{\sum_{a'} e^{v(\theta,a')/\lambda} \prod_{j\in N} \alpha_j(a'_j)} \leq 1, \qquad a_i \in A_i.$$

It follows from (7) that

$$\sum_\theta \pi(\theta) \frac{\int_{\mathcal{A}_{-i}} e^{v(\theta,a_i,\alpha'_{-i})/\lambda} \prod_{i\neq j} \mathrm{d}\mu_j(a'_j)}{\int_{\mathcal{A}} e^{v(\theta,\alpha')/\lambda} \prod_{j\in N} \mathrm{d}\mu_j(\alpha'_j)} \leq 1, \qquad a_i \in A_i.$$

In addition, by convexity of the exponential function, we have that

$$e^{v(\theta,\alpha_i,\alpha'_{-i})/\lambda} = \sum_{a_i} \alpha_i(a_i) e^{v(\theta,a_i,\alpha'_{-i})/\lambda}, \qquad \alpha_i \in \mathcal{A}_i.$$

Thus, we deduce that

$$\sum_\theta \pi(\theta) \frac{\int_{\mathcal{A}_{-i}} e^{v(\theta,\alpha_i,\alpha'_{-i})/\lambda} \prod_{i\neq j} \mathrm{d}\mu_j(a'_j)}{\int_{\mathcal{A}} e^{v(\theta,\alpha')/\lambda} \prod_{j\in N} \mathrm{d}\mu_j(\alpha'_j)} \leq 1, \qquad \alpha_i \in \mathcal{A}_i,$$

which is the first-order condition of the optimization problem

$$\max_{\nu_i \in \Delta(\mathcal{A}_i)} V(\nu_i, \bar{\mu}_{-i})$$

where $\bar{\mu}_{-i} = (\mu_j)_{j\neq i}$. Hence, (ii) of Proposition 4 holds. Overall, we have shown that

ρ is robust to information acquisition. Given that for every player i, μ_i puts positive probability only on pure actions, we conclude that ρ is *strongly* robust to unrestricted information acquisition.

"Only if." Let $\rho \in \Delta_\pi(\Theta \times A)$ be *strongly* robust to unrestricted information acquisition. In particular, ρ is robust to information acquisition, in the sense of Definition 2. Thus, there are $\mu_i \in \Delta(\mathcal{A}_i)$, $i \in N$, that satisfy (i) and (ii) of Proposition 4. The additional "strongly" qualification implies that for every player i, μ_i puts positive probability only on pure actions. Define $\alpha_i \in \mathcal{A}_i$ by

$$\alpha_i(a_i) = \mu_i(\delta_{a_i})$$

where $\delta_{a_i} \in \mathcal{A}_i$ is the Dirac measure concentrated on $a_i \in A_i$. It is easy to see that (i) and (ii) of Corollary 1 hold. In addition, for all $i \in N$ and $a_i \in A_i$,

$$\sum_{\theta, a_{-i}} \rho(\theta, a_i, a_{-i}) = \int_{\mathcal{A}_i} \alpha_i'(a_i)\mu_i(\alpha_i') = \alpha_i(a_i).$$

■

Proof of Corollary 2. The result follows from the existence of a maximizer of V in Corollary 1(ii). ■

Proof of Corollary 3. For every $\lambda > 0$ and $i \in N$, take μ_i^λ as in Proposition 4. For every $a \neq a^*$,

$$\begin{aligned}
\rho^\lambda(a|\theta^*) &= \frac{\int_{\mathcal{A}} \prod_{i=1}^n \alpha_i(a_i) e^{v(\theta^*,\alpha)/\lambda} \prod_{i=1}^n \mathrm{d}\mu_i^\lambda(\alpha_i)}{\int_{\mathcal{A}} e^{v(\theta^*,\alpha)/\lambda} \prod_{i=1}^n \mathrm{d}\mu_i^\lambda(\alpha_i)} \\
&\leq \frac{\int_{\mathcal{A}} \prod_{i=1}^n \alpha_i(a_i) e^{v(\theta,\alpha)/\lambda} \prod_{i=1}^n \mathrm{d}\mu_i^\lambda(\alpha_i)}{e^{v(\theta,a^*)/\lambda} \prod_{i=1}^n \mu_i^\lambda(\delta_{a_i^*})} \\
&= \frac{\int_{\mathcal{A}} \prod_{i=1}^n \alpha_i(a_i) e^{(v(\theta,\alpha)-v(\theta,a^*))/\lambda} \prod_{i=1}^n \mathrm{d}\mu_i^\lambda(\alpha_i)}{\prod_{i=1}^n \mu_i^\lambda(\delta_{a_i^*})}.
\end{aligned}$$

By (i), we have that

$$\limsup_{\lambda \to 0} \int_{\mathcal{A}} \prod_{i=1}^n \alpha_i(a_i) e^{(v(\theta,\alpha)-v(\theta,a^*))/\lambda} \prod_{i=1}^n \mathrm{d}\mu_i^\lambda(\alpha_i) = 0.$$

By (ii), we have that

$$\liminf_{\lambda \to 0} \mu_i^\lambda \left(\delta_{a_i^*}\right) > 0.$$

We deduce that $\rho^\lambda(a|\theta^*) \to 0$ as $\lambda \to 0$. Since the choice of $a \neq a^*$ was arbitrary, we conclude that $\rho^\lambda(a^*|\theta^*) \to 1$ as $\lambda \to 0$. ■

Proof of Lemma 1. Fix $\lambda > 0$; we will prove the stronger statement that there exists $\rho^\lambda \in \Delta_\pi(\Theta \times A)$ that is *strongly* robust to unrestricted information acquisition. To ease notation, we drop the superscript λ, and write ρ instead of ρ^λ.

Denote the players by i and j. For $\rho_j \in \Delta_\pi(\Theta \times A_j)$, define $BR_i(\rho_j) \subseteq \mathcal{P}_{A_i}$ by

$$BR_i(\rho_j) = \arg\max_{P_i \in \mathcal{P}_{A_i}} \sum_{\theta, a} u_i(\theta, a) P_i(a_i|\theta, a_j)\rho_j(\theta, a_j) - \lambda I_f(P_i, \rho_j).$$

Since $I_f(P_i, \mu_j)$ is jointly continuous in P_i and ρ_j, the correspondence

$$BR_i : \Delta_\pi(\Theta \times A_j) \rightrightarrows \mathcal{P}_{A_i}$$

has nonempty values and closed graph (by Berge's maximum theorem). In addition, $I_f(P_i, \rho_j)$ is convex in P_i.[21] Thus, BR_i also has convex values.

For $\rho_j \in \Delta_\pi(\Theta \times A_j)$, let $F_i(\rho_j) \subseteq \Delta_\pi(\Theta \times A_i)$ be the set of all $\rho_i \in \Delta(\Theta \times A_i)$ for which there exists $P_i \in \mathcal{P}_{A_i}$ such that for all θ and a_i,

$$\rho_i(\theta, a_i) = \sum_{a_j} P_i(a_i|\theta, a_j)\rho_j(\theta, a_j).$$

Since BR_i has closed graph and nonempty convex values, the correspondence

$$F_i : \Delta_\pi(\Theta \times A_j) \rightrightarrows \Delta_\pi(\Theta \times A_i)$$

has closed graph and nonempty convex values. Given that $A_j = A_i$, we deduce that F_i has a fixed point (by Kakutani's fixed-point theorem).

Let $\rho_j \in \Delta(\Theta \times A_i)$ be a fixed point of F_i; take a corresponding $P_i \in BR_i(\rho_j)$. Define $\rho \in \Delta(\Theta \times A)$ by

$$\rho(\theta, a) = P_i(a_i|\theta, a_j)\rho_j(\theta, a_j).$$

[21]This follows from the *joint* convexity of the mapping $(p, q) \to qf(p/q)$ on $\mathbb{R}^2_{++}$; see, e.g., Liese and Vajda (2006, p. 4398).

We claim that ρ is strongly robust to unrestricted information acquisition. To verify the claim, we use Proposition 3.

Let $P_j \in P_{A_j}$ be the mirror image of P_i: for all $\theta \in \Theta$ and $a \in A$,

$$P_j(a_j|\theta, a_i) = P_i(a'_i|\theta, a'_j) \quad \text{with} \quad a'_i = a_j \text{ and } a'_j = a_i.$$

Let $\rho_i \in \Delta(\Theta \times A_i)$ be the marginal distribution of ρ over states and i's actions. Since ρ_j is a fixed point of F_i, we have that for all $\theta \in \Theta$ and $a_i \in A_i$,

$$\rho_i(\theta, a_i) = \rho_j(\theta, a'_j) \quad \text{with} \quad a'_j = a_i.$$

It follows that (i) of Proposition 3 holds. In addition, by symmetry of the game, $P_i \in BR_i(\rho_j)$ implies $P_j \in BR_j(\rho_i)$. It also follows that (ii) of Proposition 3 holds. We conclude that ρ is strongly robust (hence, robust) to unrestricted information acquisition. ■

Lemma 3. *For λ sufficiently small, if ρ^λ is robust to unrestricted information acquisition, then it is strongly robust.*

Proof. Suppose that ρ^λ is robust to unrestricted information acquisition. Take $\mu^\lambda \in \Delta_\pi(\Theta \times \mathcal{A})$ and $P^\lambda \in \mathcal{P}_\mathcal{A}$ as in Proposition 2. To show that ρ^λ is also strongly robust, it is enough to show that for all players i and states θ,

$$\mu^\lambda\left(\{\delta_{Left}, \delta_{Right}\} \times \mathcal{A}_{-i}|\theta\right) = 1.$$

By contradiction, suppose that

$$\mu^\lambda(\{\delta_{Left}, \delta_{Right}\} \times \mathcal{A}_{-i}|\theta) < 1. \tag{8}$$

We distinguish between two cases. Case (i):

$$\int_{\Theta \times \mathcal{A}_{-i}} \left(u_i\left(\theta, Left, \alpha_{-i}\right) - u_i\left(\theta, Right, \alpha_{-i}\right)\right) P_i^\lambda\left(\delta_{Left}|\theta, \alpha_{-i}\right) \mathrm{d}\mu_{-i}^\lambda(\theta, \alpha_{-i}) = 0.$$

Thus, either player i never receives a mixed-action recommendation $\alpha_i = \delta_{Left}$, or when she receives it, she is indifferent between playing $Left$ and $Right$. When player i receives a mixed-action recommendation $\alpha_i \neq \delta_{Left}$, she must be willing to play

Right:

$$\int_{\Theta\times\mathcal{A}_{-i}} \left(u_i\left(\theta, Left, \alpha_{-i}\right) - u_i\left(\theta, Right, \alpha_{-i}\right)\right) P_i^\lambda\left(\alpha_i|\theta,\alpha_{-i}\right) \mathrm{d}\mu_{-i}^\lambda(\theta,\alpha_{-i}) \leq 0.$$

Therefore, player i would get the same expected utility by deviating to a direct experiment $Q_i^\lambda \in \mathcal{P}_{\mathcal{A}_i}$ such that for all $\theta \in \Theta$ and $\alpha_{-i} \in \mathcal{A}_{-i}$,

$$Q_i^\lambda(Right|\theta,\alpha_{-i}) = 1.$$

The experiment Q_i is uninformative, which implies that $Q_i \preceq P_i$. For λ sufficiently small, dominance regions guarantee that there are $\theta, \theta' \in \Theta$ such that

$$\rho^\lambda(Left, Left|\theta) > 1/2 \quad \text{and} \quad \rho^\lambda(Right, Right|\theta') > 1/2.$$

This means that P_i cannot be uninformative. We deduce that $Q_i^\lambda \prec P_i^\lambda$. Since f is *strictly* convex, $I_f(P_i^\lambda, \mu_{-i}^\lambda) > I_f(Q_i^\lambda, \mu_{-i}^\lambda)$. Thus, P_i^λ cannot be a best response to μ_{-i}^λ: contradiction.

Case (ii):

$$\int_{\Theta\times\mathcal{A}_{-i}} \left(u_i\left(\theta, Left, \alpha_{-i}\right) - u_i\left(\theta, Right, \alpha_{-i}\right)\right) P_i^\lambda\left(\delta_{Left}|\theta,\alpha_{-i}\right) \mathrm{d}\mu_{-i}^\lambda(\theta,\alpha_{-i}) > 0. \quad (9)$$

Thus, player i receives the mixed-action recommendation $\alpha_i = \delta_{Left}$ with positive probability, and when she receives it, she strictly prefers to play $Left$ rather than *Right*. When player i receives a mixed-action recommendation $\alpha_i \neq \{\delta_{Left}, \delta_{Right}\}$, she must be indifferent between $Left$ and *Right*:

$$\int_{\Theta\times\mathcal{A}_{-i}} \left(u_i\left(\theta, Left, \alpha_{-i}\right) - u_i\left(\theta, Right, \alpha_{-i}\right)\right) P_i^\lambda\left(\alpha_i|\theta,\alpha_{-i}\right) \mathrm{d}\mu_{-i}^\lambda(\theta,\alpha_{-i}) = 0. \quad (10)$$

Therefore, player i would get the same expected utility by deviating to a direct experiment $Q_i \in \mathcal{P}_{\mathcal{A}_i}$ such that for all $\theta \in \Theta$ and $\alpha_{-i} \in \mathcal{A}_{-i}$,

$$Q_i^\lambda\left(\delta_{Left}|\theta,\alpha_{-i}\right) = P_i^\lambda\left(\mathcal{A}_i \setminus \{\delta_{Right}\}|\theta,\alpha_{-i}\right),$$
$$Q_i^\lambda\left(\delta_{Right}|\theta,\alpha_{-i}\right) = P_i^\lambda\left(\delta_{Right}|\theta,\alpha_{-i}\right).$$

It is obvious that $Q_i^\lambda \preceq P_i^\lambda$. But further inspection reveals that $Q_i^\lambda \prec P_i^\lambda$: it follows from (8)-(10). Since f is *strictly* convex, $I_f(P_i^\lambda, \mu_{-i}^\lambda) > I_f(Q_i^\lambda, \mu_{-i}^\lambda)$. Thus, P_i^λ cannot be a best response to μ_{-i}^λ: contradiction. □

Proof of Proposition 5. By Lemma 3, we can assume that ρ^λ is *strongly* robust to information acquisition. Take $P^\lambda \in \mathcal{P}_\mathcal{A}$ as in Proposition 3. For every player i, denote by $\alpha_i^\lambda \in \mathcal{A}_i$ the marginal distribution over i's actions:

$$\alpha_i^\lambda(a_i) = \sum_{\theta, a_{-i}} \rho^\lambda(\theta, a_i, a_{-i}).$$

Define also $\alpha_i(a_i) = \lim_{\lambda \to 0} \alpha_i^\lambda(a_i)$. It is clear that

$$\alpha_i(a_i) = \sum_{\theta, a_{-i}} \rho(\theta, a_i, a_{-i}).$$

Dominance regions guarantee that there are $\theta, \theta' \in \Theta$ such that

$$\rho(Left, Left|\theta) = 1 = \rho(Right, Right|\theta').$$

This implies that for both players i,

$$\alpha_i(Left) \in (0, 1).$$

Thus, without loss of generality, we can assume that for all λ,

$$\alpha_i^\lambda(Left) \in (0, 1). \tag{11}$$

In short, for every $\theta \in \Theta$ and $a_{-i} \in A_{-i}$, we define

$$\Delta u_i(\theta, a_{-i}) = u_i(\theta, Left, a_{-i}) - u_i(\theta, Right, a_{-i}).$$

We prove the proposition in five claims:

Claim 3. If $\rho_{-i}^\lambda(\theta, a_{-i}) > 0$, then $P_i^\lambda(a_i|\theta, a_{-i}) \in (0, 1)$.

Proof of the claim. For ease of exposition, we drop the superscript λ. By contradiction, suppose that $\rho_{-i}(\theta, a_{-i}) > 0$ and $P_i(a_i|\theta, a_{-i}) \in \{0, 1\}$; without loss of generality, assume that $P_i(Left|\theta, a_{-i}) = 0$.

For every $\epsilon \in (0,1)$, consider the deviation $P_i^\epsilon \in \mathcal{P}_{\mathcal{A}_i}$ given by

$$P_i^\epsilon(Left|\theta', a'_{-i}) = \begin{cases} \epsilon & \text{if } (\theta', a'_{-i}) = (\theta, a_{-i}), \\ P_i(Left|\theta', a'_{-i}) & \text{if } (\theta', a'_{-i}) \neq (\theta, a_{-i}). \end{cases}$$

The quantity $I_f(P_i^\epsilon, \mu_{-i})$ is convex in P_i^ϵ; thus, it is convex in ϵ. Therefore, we have the inequality

$$I_f(P_i, \mu_{-i}) - I_f(P_i^\epsilon, \mu_{-i}) \geq -\frac{\partial I_f(P_i^\epsilon, \mu_{-i})}{\partial \epsilon}\epsilon. \tag{12}$$

Since $\alpha_i(Left) \in (0,1)$—see (11)—and $\lim_{t\to 0} f'(t) = -\infty$, basic calculus shows that

$$\lim_{\epsilon \to 0} \frac{\partial I_f(P_i^\epsilon, \mu_{-i})}{\partial \epsilon} = -\infty. \tag{13}$$

In addition, simple algebra shows that

$$\begin{aligned} &\sum_{\theta', a'} u_i(\theta', a') P_i(a'_i|\theta', a'_{-i}) \rho_{-i}(\theta', a'_{-i}) - \sum_{\theta', a'} u_i(\theta', a') P_i^\epsilon(a'_i|\theta', a'_{-i}) \rho_{-i}(\theta', a'_{-i}) \\ =& -\epsilon \Delta u_i(\theta, a_{-i}) \rho_{-i}(\theta, a_{-i}). \end{aligned} \tag{14}$$

By (13), we can choose ϵ small enough such that

$$\Delta u_i(\theta, a_{-i}) \rho_{-i}(\theta, a_{-i}) > \frac{\partial I_f(P_i^\epsilon, \mu_{-i})}{\partial \epsilon}.$$

It follows from (12) and (14) that P_i^ϵ is a profitable deviation: contradiction. □

Claim 4. For all $\theta \in \Theta$ and $a_{-i} \in A_{-i}$,

$$\Delta u_i(\theta, a_{-i}) = \lim_{\lambda \to 0} \lambda f'\left(\frac{P_i^\lambda(Left|\theta, a_{-i})}{\alpha_i^\lambda(Left)}\right) - \lambda f'\left(\frac{P_i^\lambda(Right|\theta, a_{-i})}{\alpha_i^\lambda(Right)}\right).$$

Proof of the claim. Take $\theta \in \Theta$ and $a_{-i} \in A_{-i}$. It follows from Claim 3 that ρ^λ has full support. Thus, in particular, $\rho_{-i}^\lambda(\theta, a_{-i}) > 0$ and $P_i^\lambda(Left|\theta, a_{-i}) \in (0,1)$. In addition, the direct experiment P_i^λ is an optimal solution to

$$\max_{P_i} \sum_{\theta, a} u_i(\theta, a) P_i(a_i|\theta, a_{-i}) \rho_{-i}^\lambda(\theta, a_{-i}) - \lambda I_f(P_i, \mu_{-i}).$$

Therefore, the following first-order condition must hold:

$$\Delta u_i(\theta, a_{-i})\rho^\lambda_{-i}(\theta, a_{-i}) = \lambda \frac{\partial \lambda I_f(P_i^\lambda, \mu_{-i})}{\partial P_i^\lambda(Left|\theta, a_{-i})}.$$

Simple calculus shows that $\partial \lambda I_f(P_i^\lambda, \mu_{-i})/\partial P_i^\lambda(Left|\theta, a_{-i})$ is equal to

$$\rho^\lambda_{-i}(\theta, a_{-i})\left(f'\left(\frac{P_i^\lambda(Left|\theta, a_{-i})}{\alpha_i^\lambda(Left)}\right) - f'\left(\frac{P_i^\lambda(Right|\theta, a_{-i})}{\alpha_i^\lambda(Right)}\right) + o(1/\lambda)\right)$$

where $o(1/\lambda)$ grows much slower than $1/\lambda$. The desired result follows. □

Claim 5. If $\Delta u_i(\theta, a_{-i}) > 0$, then $P_i^\lambda(Left|\theta, a_{-i}) \to 1$.

Proof of the claim. By contradiction, suppose that $P_i^\lambda(Left|\theta, a_{-i}) \to t \in [0, 1)$. If $t \in (0, 1)$, then

$$\lim_{\lambda\to 0} \lambda f'\left(\frac{t}{\alpha_i(Left)}\right) - \lambda f'\left(\frac{1-t}{\alpha_i(Right)}\right) = 0 < \Delta u_i(\theta, a_{-i}),$$

which contradicts Claim 4. If $t = 0$, then

$$\lim_{\lambda\to 0} \lambda f'\left(\frac{t}{\alpha_i(Left)}\right) \leq 0 < \Delta u_i(\theta, a_{-i}),$$

which, again, contradicts Claim 4. □

Claim 6. If $Left$ is risk dominant at θ, then

$$\rho(Left, Left|\theta) \geq \sum_{\theta'} \rho(Left, Left, \theta').$$

Proof of the claim. Define $\alpha_i^\lambda(Left|\theta)$ and $\alpha_i(Left|\theta)$ by

$$\begin{aligned}
\alpha_i^\lambda(Left|\theta) &= \rho^\lambda(Right, Left|\theta) + \rho^\lambda(Left, Left|\theta),\\
\alpha_i(Left|\theta) &= \lim_{\lambda\to 0} \alpha_i^\lambda(Left|\theta) = \rho(Right, Left|\theta) + \rho(Left, Left|\theta).
\end{aligned}$$

In the limit, the players must perfectly coordinate their actions, which implies that

$$\alpha_i(Left|\theta) = \rho(Left, Left|\theta) \quad \text{and} \quad \alpha_i(Left) = \sum_{\theta'} \rho(Left, Left, \theta').$$

Hence, to verify the claim, it is enough to show that

$$\alpha_i(Left|\theta) \geq \alpha_i(Left). \tag{15}$$

If $\theta_{rr} - \theta_{lr} < 0$, then the unique complete-information equilibrium is $(Left, Left)$, which implies that $\alpha_i(Left|\theta) = 1$, which in turn implies that (15) trivially holds.

Assume now that $\theta_{rr} - \theta_{lr} > 0$, so that both $(Left, Left)$ and $(Right, Right)$ are strict complete-information equilibria. By Claims 4 and 5,

$$\theta_{ll} - \theta_{rl} = -\lim_{\lambda \to 0} \lambda f'\left(\frac{P_i^\lambda(Right|\theta, Left)}{\alpha_i^\lambda(Right)}\right),$$
$$\theta_{rr} - \theta_{lr} = -\lim_{\lambda \to 0} \lambda f'\left(\frac{P_i^\lambda(Left|\theta, Right)}{\alpha_i^\lambda(Left)}\right).$$

Since $Left$ is risk dominant at θ, $\theta_{ll} - \theta_{rl} > \theta_{rr} - \theta_{lr}$. Thus, for λ sufficiently small,

$$f'\left(\frac{P_i^\lambda(Right|\theta, Left)}{\alpha_i^\lambda(Right)}\right) \leq f'\left(\frac{P_i^\lambda(Left|\theta, Right)}{\alpha_i^\lambda(Left)}\right).$$

Since f is strictly convex, f' is strictly increasing. It follows that

$$\frac{P_i^\lambda(Left|\theta, Right)}{1 - P_i^\lambda(Left|\theta, Left) + P_i^\lambda(Left|\theta, Right)} \geq \alpha_i^\lambda(Left). \tag{16}$$

Starting from the system of equations

$$\alpha_i^\lambda(Left|\theta) = P_i^\lambda(Left|\theta, Left)\alpha_{-i}^\lambda(Left|\theta) + P_i^\lambda(Left|\theta, Right)(1 - \alpha_{-i}^\lambda(Left|\theta))$$
$$\alpha_{-i}^\lambda(Left|\theta) = P_{-i}^\lambda(Left|\theta, Left)\alpha_i^\lambda(Left|\theta) + P_{-i}^\lambda(Left|\theta, Right)(1 - \alpha_i^\lambda(Left|\theta)),$$

we obtain that $\alpha_i^\lambda(Left|\theta)$ is equal to

$$\frac{P_{-i}^\lambda(Left|\theta, Right)\left(P_i^\lambda(Left|\theta, Left) - P_i^\lambda(Left|\theta, Right)\right) + P_i^\lambda(Left|\theta, Right)}{1 - \left(P_i^\lambda(Left|\theta, Left) - P_i^\lambda(Left|\theta, Right)\right)\left(P_{-i}^\lambda(Left|\theta, Left) - P_{-i}^\lambda(Left|\theta, Right)\right)}.$$

By Claim 5, for λ sufficiently small,

$$P_i^\lambda(Left|\theta, Left) > P_i^\lambda(Left|\theta, Right) \quad \text{and} \quad P_{-i}^\lambda(Left|\theta, Left) > P_{-i}^\lambda(Left|\theta, Right).$$

Thus, we can minorize $\alpha_i^\lambda(Left|\theta)$ and obtain that

$$\alpha_i^\lambda(Left|\theta) \geq \frac{P_i^\lambda(Left|\theta, Right)}{1 - P_i^\lambda(Left|\theta, Left) + P_i^\lambda(Left|\theta, Right)}.$$

It follows from (16) that $\alpha_i^\lambda(Left|\theta) \geq \alpha_i^\lambda(Left)$. □

Claim 7. If $Left$ is less risky at θ than at θ', then

$$\rho(Left, Left|\theta) \geq \rho(Left, Left|\theta').$$

Proof of the claim. If $\theta'_{ll} - \theta'_{rl} < 0$, then trivially

$$\rho(Left, Left|\theta) \geq 0 = \rho(Left, Left|\theta').$$

If $\theta_{rr} - \theta_{lr} < 0$, then trivially

$$\rho(Left, Left|\theta) = 1 \geq \rho(Left, Left|\theta').$$

Now, assume that $\theta_{ll} - \theta_{rl} < 0$. Since $Left$ is less risky at θ than at θ', $\theta'_{ll} - \theta'_{rl} < 0$. Hence, $(Right, Right)$ is the unique complete-information equilibrium both at θ and at θ'. It follows that

$$\rho(Left, Left|\theta) = 0 = \rho(Left, Left|\theta').$$

Next, assume that $\theta'_{rr} - \theta'_{lr} < 0$. Since $Left$ is less risky at θ than at θ', $\theta_{rr} - \theta_{lr} < 0$. Hence, $(Left, Left)$ is the unique complete-information equilibrium both at θ and at θ'. It follows that

$$\rho(Left, Left|\theta) = 1 = \rho(Left, Left|\theta').$$

Finally, assume that both $(Left, Left)$ and $(Right, Right)$ are complete-information equilibria at θ and at θ'. By Claims 4 and 5, we have that for both players i,

$$\theta_{ll} - \theta_{rl} = -\lim_{\lambda \to 0} \lambda f'\left(\frac{P_i^\lambda(Right|\theta, Left)}{\alpha_i^\lambda(Right)}\right),$$
$$\theta_{rr} - \theta_{lr} = -\lim_{\lambda \to 0} \lambda f'\left(\frac{P_i^\lambda(Left|\theta, Right)}{\alpha_i^\lambda(Left)}\right),$$
$$\theta'_{ll} - \theta'_{rl} = -\lim_{\lambda \to 0} \lambda f'\left(\frac{P_i^\lambda(Right|\theta', Left)}{\alpha_i^\lambda(Right)}\right),$$
$$\theta'_{rr} - \theta'_{lr} = -\lim_{\lambda \to 0} \lambda f'\left(\frac{P_i^\lambda(Left|\theta', Right)}{\alpha_i^\lambda(Left)}\right).$$

Since $Left$ is less risky at θ than at θ', we have that for λ sufficiently small,

$$f'\left(\frac{P_i^\lambda(Right|\theta, Left)}{\alpha_i^\lambda(Right)}\right) \leq f'\left(\frac{P_i^\lambda(Right|\theta', Left)}{\alpha_i^\lambda(Right)}\right),$$
$$f'\left(\frac{P_i^\lambda(Left|\theta', Right)}{\alpha_i^\lambda(Left)}\right) \leq f'\left(\frac{P_i^\lambda(Left|\theta, Right)}{\alpha_i^\lambda(Left)}\right).$$

Since f is strictly convex, f' is increasing. Thus,

$$P_i^\lambda(Right|\theta, Left) \leq P_i^\lambda(Right|\theta', Left),$$
$$P_i^\lambda(Left|\theta', Right) \leq P_i^\lambda(Left|\theta, Right).$$

Since this is true for both players i, we conclude that

$$\rho(Left, Left|\theta) \geq \rho(Left, Left|\theta').$$

□

■

References

Alaoui, Larbu and Antonio Penta, "Cost-benefit analysis in reasoning," *Journal of Political Economy*, 2022, *130*, 881–925.

Ali, S. M. and S. D. Silvey, "A general class of coefficients of divergence of one distribution from another," *Journal of the Royal Statistics Society, Series B*, 1966, *28*, 131–142.

Bergemann, Dirk and Juuso Valimaki, "Information acquisition and efficient mechanism design," *Econometrica*, 2002, *70*, 1007–1033.

_ **and Stephen Morris**, "Bayes correlated equilibrium and the comparison of information structures in games," *Theoretical Economics*, 2016, *11*, 487–522.

_ **and** _ , "Information design: A unified perspective," *Journal of Economic Literature*, 2019, *57*, 44–95.

Blackwell, David, "Comparison of experiments," in Jerzy Neyman, ed., *Proceedings of the 2nd Berkeley Symposium on Mathematical Statistics and Probability*, University of California Press 1951, pp. 93–102.

Blume, Lawrence E., "The statistical mechanics of strategic interaction," *Games and Economic Behavior*, 1993, *5*, 387–424.

Caplin, Andrew, Mark Dean, and John Leahy, "Rationally inattentive behavior: Characterizing and generalizing Shannon entropy," *Journal of Political Economy*, 2022, *130*, 1676–1715.

Carlsson, Hans and Eric van Damme, "Global games and equilibrium selection," *Econometrica*, 1993, *61*, 989–1018.

Cover, Thomas M. and Joy A. Thomas, *Elements of Information Theory*, 2nd ed., John Wiley & Sons, 2006.

Csiszar, Imre, "Eine informationstheoretische ungleichung und ihre anwendung auf den beweis der Ergodizitat von markoffschen ketten," *Publ. Math. Inst. Hungar. Acad. Sci.*, 1963, *8*, 85–108.

_ , "Information-type measures of difference of probability distributions and indirect observations," *Stud. Scien. Math. Hung.*, 1967, *2*, 299–318.

Denti, Tommaso, Massimo Marinacci, and Aldo Rustichini, "Experimental cost of information," *American Economic Review*, 2002, *112*, 3106–3123.

__ , __ , **and Luigi Montrucchio**, "A note on rational inattention and rate distortion theory," *Decisions in Economics and Finance*, 2020, *43*, 75–89.

Forges, Francoise, "Five legitimate definitions of correlated equilibrium in games with incomplete information," *Theory and Decision*, 1993, *35*, 277–310.

Harsanyi, John C. and Reinhard Selten, *A General Theory of Equilibrium Selection in Games*, MIT Press, 1988.

Hebert, Benjamin and Jennifer La'O, "Information acquisition, efficiency, and non-fundamental volatility," 2021. Available at https://www.jennifer-la-o.com.

__ **and Michael Woodford**, "Neighborhood-based information costs," *American Economic Review*, 2021, *111*, 3225–55.

Hellwig, Christian and Laura Veldkamp, "Knowing what others know: coordination motives in information acquisition," *Review of Economic Studies*, 2009, *76*, 223–251.

Hoshino, Tetsuya, "Rational inattention in games," 2018. Available at https://sites.google.com/view/tetsuyahoshino.

Kajii, Atsushi and Stephen Morris, "The robustness of equilibria to incomplete information," *Econometrica*, 1997, *65*, 1283–1309.

Liese, Friedrich and Igor Vajda, "On Divergences and Informations in Statistics and Information Theory," *IEEE Transactions on Information Theory*, 2006, *52*, 4394–4412.

Mackowiak, Bartosz, Filip Matejka, and Mirko Wiederholt, "Rational inattention: a review," *Journal of Economic Literature*, Forthcoming.

Matejka, Filip and Alisdair McKay, "Rational inattention to discrete choices: A new foundation for the multinomial logit model," *American Economic Review*, 2015, *105*, 272–298.

Monderer, Dov and Lloyd S. Shapley, "Potential games," *Games and Economic Behavior*, 1996, *14*, 124–143.

Morris, Stephen and Hyun S. Shin, "Rethinking multiple equilibria in macroeconomic modeling," in Ben S. Bernanke and Kenneth Rogoff, eds., *NBER Macroeconomics Annual 2000, Volume 15*, MIT Press, 2001, pp. 139–182.

_ **and** _ , "Global games: theory and applications," in Mathias Dewatripont, Lars P. Hansen, and Stephen Turnovsky, eds., *Advances in Economics and Econometrics (Proceedings of the Eighth World Congress of the Econometric Society)*, Cambridge University Press 2003, pp. 56–114.

_ **and Ming Yang**, "Coordination and continuous stochastic choice," *Review of Economic Studies*, Forthcoming.

Myatt, David P. and Chris Wallace, "Endogenous information acquisition in coordination games," *Review of Economic Studies*, 2012, *79*, 340–374.

Nimark, Kristoffer P. and Savi Sundaresan, "Inattention and belief polarization," *Journal of Economic Theory*, 2019, *180*, 203–228.

Persico, Nicola, "Information acquisition in auctions," *Econometrica*, 2000, *68*, 135–148.

Pomatto, Luciano, Philipp Strack, and Omer Tamuz, "The Cost of information," 2020. Available at arXiv:1812.04211.

Sims, Christopher A., "Implications of rational inattention," *Journal of Monetary Economics*, 2003, *50*, 665–690.

_ , "Rational inattention and monetary economics," in Benjamin M. Friedman and Michael Woodford, eds., *Handbook of Monetary Economics*, Elsevier., 2010, pp. 155–181.

Weinstein, Jonathan and Muhamet Yildiz, "A structure theorem for rationalizability with application to robust predictions of refinements," *Econometrica*, 2007, *75*, 365–400.

Woodford, Michael, "Inattentive valuation and reference-dependent choice," 2012. Available at https://blogs.cuit.columbia.edu/mw2230/.

Yang, Ming, "Coordination with flexible information acquisition," *Journal of Economic Theory*, 2015, *158*, 721–738.