Rational Inattention in Games

Tetsuya Hoshino*

November 10, 2018

Job Market Paper

Check tetsuyahoshino.com for the latest version

Abstract

I study an incomplete-information game with information acquisition. As in the theory of rational inattention, players flexibly choose what information to acquire, but such information is costly. If information costs were zero, players would learn the state, thereby reducing the game to one of complete information. The complete-information game, typically with multiple equilibria, can thus be approximated by a nearby incomplete-information game with small information costs. In this paper, I model these costs using a general nonparametric form, which includes Shannon entropy costs as a special case. The main result is that *any* strict Nash equilibrium of the complete-information game is arbitrarily close to the *unique* equilibrium of a nearby incomplete-information game with *some* small information costs. This result has the implication, for instance, that in a 2×2 coordination game, even the non-risk-dominant equilibrium is robust in this sense.

Keywords: rational inattention, information acquisition, equilibrium selection, Bayes correlated equilibrium, a structure theorem.

JEL classification: C72, D82, D83.

---

*I am grateful to Vijay Krishna, Nageeb Ali, and Kalyan Chatterjee for their continuous support and guidance. I also thank Yu Awaya, Henrique de Oliveira, Tommaso Denti, Edward Green, Keisuke Hirano, Arthur Robson, Ariel Rubinstein, Eran Shmaya, Colin Stewart, Jonathan Weinstein, and Yuichi Yamamoto as well as participants in the Pennsylvania Economic Theory Conference and the Midwest Economic Theory Conference for their comments. All remaining errors are my own.
Penn State University · tzh144@psu.edu · tetsuyahoshino.com

1 Introduction

Economic agents may have incentives to acquire information to make better decisions. What information to acquire—or pay attention to—is itself a choice. If information is costly, economic agents have to balance the benefits and costs of acquiring information. The theory of rational inattention is not only an important paradigm for bounded rationality resulting from limited cognitive ability (Simon, 1955; Sims, 2003) but also provides a useful framework to model information acquisition. In this framework, economic agents choose what information to acquire, but because information is costly—economically or psychologically—they may rationally choose not to be perfectly informed.

In this paper, I study a strategic situation in which multiple players interact, or an incomplete-information game with costly information acquisition, and examine how information acquisition affects equilibrium behavior. I refer to this game as a **rational inattention game**—an **RI game** for short.[1] Players' payoffs in an RI game consist of their payoffs in the incomplete-information game minus the information costs. In a strategic situation, players may want to reduce both kinds of uncertainty: Structural uncertainty about the underlying fundamentals and strategic uncertainty about their opponents' behavior. Players may want to reduce both kinds of uncertainty but, as argued above, they may rationally choose not to eliminate it completely. Information costs depend on what information they acquire, as will be specified in more detail below. The solution concept employed is that of Bayes correlated equilibrium (Bergemann and Morris, 2016), amended to account for information costs.

How should information costs be modeled? In the literature, information costs have often been modeled using Shannon entropy, as suggested by Sims (2003, 2010), and the Shannon information costs have been widely used due to their tractability. However, there are reasons to depart from focusing on such costs. First, information costs may correspond to the costs of hiring an expert, hiring a consulting firm, or performing market research, and a priori, there is no reason for the market price for information or the effort cost to take the entropy form (e.g., Persico, 2000). Second, even when it comes to the issue of information processing, researchers have observed that the Shannon information costs lead to potentially unrealistic conclusions. For example, they imply an "independence-of-irrelevant-alternatives (IIA)" property, which states that for a given state of the world, the ratio of an agent's choice probabilities of two alternatives

[1] I will use the term "rational attention" to describe this entire class of models, while recognizing that the literature often uses this term to refer to the specific case of the Shannon information costs.

does not depend on the payoff from a third (irrelevant) alternative (Fosgerau et al., 2018).[2] Hence, the Shannon information costs are independent of the labelling of alternatives and cannot capture similarities between alternatives. That is, they are context independent (Hobson, 1969). Moreover, experimental evidence suggests that observed behavior is not consistent with the Shannon information costs (Dean and Neligh, 2017).

In this paper, I allow information costs to take a general nonparametric form, which includes the Shannon information costs as a special case. I require information costs to satisfy some fundamental properties which rule out ad hoc settings and provide normative support for using the general information costs. For example, I require that the cost of better information be higher and that the cost of no information be zero. In Section 3.3, I will elaborate on the modeling and interpretations of information acquisition as well as information costs in greater detail.

The Main Result. If information costs were zero, players would acquire full information, so that an RI game would effectively be a complete-information game with a true state. The complete-information game can thus be approximated by a nearby RI game with small information costs.

This paper addresses the robustness of equilibria to the specification of information costs or the endogenous information structures thereof. That is, the question is: How do the equilibria of the nearby RI game vary as the information costs change slightly? The main result—a **structure theorem**—is:

> **Theorem.** *For any incomplete-information game,* any *strict Nash equilibrium of the complete-information game with a true state is arbitrarily close to the* unique *equilibrium of a nearby rational inattention game with* some *small information costs.*

Why is any strict Nash equilibrium of the underlying game uniquely selected by some small information costs? Each player can choose an optimal action with high probability when information costs are small. In a strategic setting, a player care about a small risk that the opponent may make a mistake. This small risk affects the player's

[2] The IIA property discusses how choices vary across different states of the world. This is different from the exercise in Matějka and McKay (2015), who consider how choices vary as priors or choice sets change. To examine the robustness of equilibria, the change in behavior across the states of the world is more relevant; a researcher may need to uncover how sensitive players' behavior is to the change in the state of the world.

behavior, which, in turn, changes the opponents' behavior. Again, it drives the original player to revise his behavior, and so on. This fear of these small risks of mistakes is amplified through strategic interaction, which results in a large difference in an equilibrium play.

A (partial) converse of this result is also true: Any non-Nash equilibrium action profile could be played with probability close to zero in any equilibrium of a nearby RI game with any small information costs. If the complete-information games do not have weak (pure-strategy) Nash equilibria, this converse is the full converse.[3]

The main result holds in *any* finite game with the **richness assumption**, which I will now discuss, as well as a regularity assumption which is satisfied generically. The richness assumption requires that every action of every player be strictly dominant at some state. The same assumption was also made by Weinstein and Yildiz (2007), but interestingly, it played different roles in these two studies. While they used it to exploit an infection argument on players' belief hierarchies, I use it to make players' behavior sensitive to perturbations of information costs.

An implication of this result is that the nature of players' information acquisition technology must be carefully analyzed to make economically justifiable predictions, particularly to make sharper predictions than strict Nash equilibria. If one interprets information costs as a model of players' limited cognitive ability or bounded rationality, the nature of their bounded rationality must be analyzed to make predictions.

Related Literature. This study is closely related to the literature of global games because both examine the robustness of equilibria to information perturbation. The electronic mail game of Rubinstein (1989) illustrates that the prediction of the payoff-dominant equilibrium may be highly sensitive to the specification of information structures. The global game of Carlsson and van Damme (1993) shows that in a two-player-two-action coordination game with "small" incomplete information, the risk-dominant equilibrium is uniquely selected. Weinstein and Yildiz (2007) argued that these particular selections of equilibria are due to particular selections of information structure—in other words, for any rationalizable action profile, there is some nearby type at which the action profile is uniquely selected.

These studies considered the exogenous perturbation of information structures, but

[3] Any action profile, even non-rationalizable one, is played in an equilibrium if information costs can be taken arbitrarily large. It suffices to calibrate marginal benefits and costs so that players choose such information structures that they take the actions (Denti, 2018). However, these information structures will be very different from that of the complete-information game.

recent studies, like mine, make information structures endogenous. In this introduction, I elaborate on studies taking the rational-inattention approach.[4] The existing studies mainly focused on coordination games. Yang (2015) and Denti (2018) allowed players to choose what information to acquire. Both of them employed the Shannon information costs, but Yang (2015) required that players' information be independent across them, while Denti (2018) did not. As the information becomes cheap, multiple equilibria arise in the former, while the risk-dominant equilibrium is uniquely selected in the latter. Morris and Yang (2016), like Yang (2015), required that players' information be independent but did not use the Shannon information costs. They show that a particular equilibrium is uniquely played as the information became cheap. They interpret this result as a warning about the use of the Shannon information costs in strategic settings. I generalize the result of equilibrium sensitivity to the specification of information acquisition in this direction. This result draws attention to the specification of information acquisition or endogenous information structures thereof in strategic settings.

Layout. The rest of the paper is organized as follows. Section 2 illustrates the logic of the main result, using an example. Section 3 sets up the model, and Section 4 presents the main result. Section 5 concludes. Appendix A contains proofs from Section 4.

2 An Example

In this section, I illustrate the logic of the main result, using an example.

Consider a two-player investment game, in which player $i = 1, 2$ chooses an action a_i from the set $A_i = \{1, 0\}$, where action 1 is "Invest" and action 0 is "Not." As usual, let $A = A_1 \times A_2$. There are three states $\Theta = \{-\frac{1}{3}, \frac{2}{3}, \frac{5}{3}\}$ with a full-support common prior $\psi \in \Delta(\Theta)$. Payoff functions u_i are defined by Table 1.

θ	1	0
1	$1-\theta, 1-\theta$	$-\theta, 0$
0	$0, -\theta$	$0, 0$

Table 1: Payoff Structure of the Investment Game

[4] Another strand of the literature makes information acquisition endogenous but rigid, where players receive Gaussian information (Hellwig and Veldkamp, 2009; Myatt and Wallace, 2011). Information structures are endogenous also because of signalling (Angeletos et al., 2006), learning (Angeletos et al., 2007; Basu et al., 2018), and communication (Hoshino, 2018).

Action 1 is strictly dominant at state $\theta = -\frac{1}{3}$, and action 0 is strictly dominant at state $\theta = \frac{5}{3}$. There are two pure-strategy Nash equilibria $(1, 1)$ and $(0, 0)$ at state $\theta = \frac{2}{3}$; the former is payoff dominant and the latter is risk dominant.

To choose an action a_i, player i may acquire information about the state θ and the action a_j that player $j \neq i$ is going to take. Note that this game is itself a static game. By acquiring information about the action a_j that player j is going to take, I mean that player i's behavior is correlated to player j's behavior. See Section 3.3 for more discussion on the modeling and interpretation.

Rational inattention provides a useful framework to model this situation. It models player i as if he were choosing his conditional choice probabilities $p_i(\cdot \mid a_j, \theta)$ for each (a_j, θ). This choice reflects his choice of information and results in the form of mistakes. For example, if he acquires no information, he chooses a fixed action a_i for any pair (a_j, θ), whereas if he acquires precise information, he makes more effort, so that a better conditional $p_i(\cdot \mid a_j, \theta)$, which brings more expected payoff, should be most costly.

These information costs are modeled as a function $c : \Delta(A \times \Theta) \to \mathbb{R}_+$ that specifies the cost $c(p) \geq 0$ of a joint p, and the joint is defined by player i's conditional choice probabilities and a given marginal, denoted $p(a_j, \theta)$. Player i maximizes the expected payoff minus the costs:

$$\sum_{a_j, \theta} p(a_j, \theta) \sum_{a_i} p_i(a_i \mid a_j, \theta) u_i(a_i, a_j, \theta) - \lambda c(p), \tag{1}$$

where $\lambda > 0$ denotes a scale factor.

I describe an appropriate solution concept. If a state θ were common knowledge, players' choice of actions would be correlated; thus, the solution concept would be correlated equilibrium. Now there is incomplete information about a state θ. Then, a solution concept is, roughly speaking, an incomplete-information version of correlated equilibrium. In particular, I use Bayes correlated equilibrium (Bergemann and Morris, 2016), amended to account for information costs. An equilibrium is defined as a joint $p^* \in \Delta(A \times \Theta)$ such that its marginal $p^*(\theta)$ is equal to the prior and that the conditional choice probabilities $p_i^*(\cdot \mid a_j, \theta)$ are optimal for player i given the marginal $p^*(a_j, \theta)$.

Now that I have described the primitives, I illustrate the main result by showing that the equilibrium is highly sensitive to the specification of information costs. To begin with, I translate an equilibrium p^* into the language of Markov chain. I emphasize that *this game is a static game but interpreted as if it were a dynamic game*. In Figure 1, node a_1a_2 is action profile (a_1, a_2), and arrow $a_1a_2 \to a_1'a_2$ is a transition from node

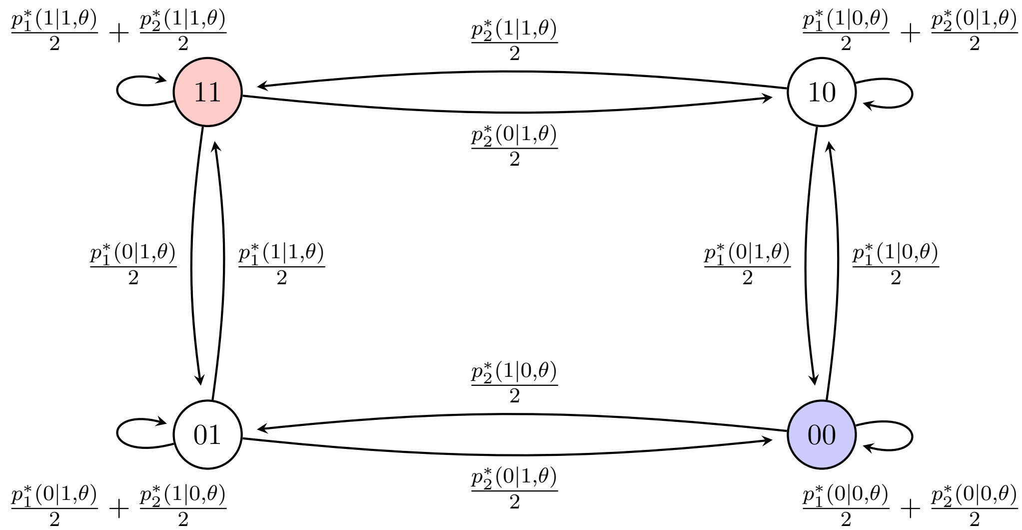

Figure 1: Markov Chain at State θ

$a_1 a_2$ to $a_1' a_2$. Define a transition rule as follows: Conditional on a state θ, at node $a_i a_j$, one player is chosen with equal probability and then the chosen player, say i, chooses an action according to the conditional $p^*(\cdot \mid a_j, \theta)$. To be explicit about the chosen player, I will put subscript i and write $p_i^*(\cdot \mid a_j, \theta)$. For example, arrow $11 \to 01$ has probability $\frac{p_1^*(0|1,\theta)}{2}$ since player 1 chooses action 0 given player 2's action 1. This transition rule defines a Markov chain; by construction, the equilibrium distribution $p^*(\cdot \mid \theta) \in \Delta(A)$ is a stationary distribution of the Markov chain. This stationary distribution is unique (under a mild condition), and this result will be used to show the equilibrium uniqueness of the game.

The Shannon Information Cost. The most frequently used information cost function in the literature is based on Shannon entropy:[5]

$$c_H(p) := \sum_{a_j,\theta} p(a_j,\theta) \underbrace{\sum_{a_i} p_i(a_i \mid a_j,\theta)\Big(\log p_i(a_i \mid a_j,\theta) - \log p_i(a_i)\Big)}_{=:\, H},$$

where $p_i(a_i)$ is player i's *unconditional* choice probability of action a_i. Loosely speaking, H measures how far the conditional $p_i(\cdot \mid a_j, \theta)$ is from the unconditional $p_i(\cdot)$, using the log difference. It depends on the pair (a_j, θ), but I abuse notation by not writing the dependence explicitly.

[5]See Cover and Thomas (2006) for the information-theoretic foundation of Shannon entropy.

Given the information costs $c = c_H$, player i's optimal conditional choice probabilities take the logit-like form (Matějka and McKay, 2015): For the reciprocal $\beta = 1/\lambda$,

$$p_i(a_i \mid a_j, \theta) = \frac{p_i(a_i)e^{\beta u_i(a_i,a_j,\theta)}}{p_i(1)e^{\beta u_i(1,a_j,\theta)} + p_i(0)e^{\beta u_i(0,a_j,\theta)}}.$$

These conditional choice probabilities depend on the unconditional ones, but they are endogenously determined at an equilibrium. It is difficult to identify the unconditional ones, but their specific values are, in fact, not important in this example. Here, I assume the unconditional: $p_i(0) = \frac{1}{3}$ and $p_i(1) = \frac{2}{3}$. Substituting them, I obtain the following:

$$p_i\big(0 \mid 1, \tfrac{2}{3}\big) = \frac{1}{2e^{\frac{1}{3}\beta} + 1}, \qquad p_i\big(1 \mid 0, \tfrac{2}{3}\big) = \frac{2e^{-\frac{2}{3}\beta}}{2e^{-\frac{2}{3}\beta} + 1}. \tag{2}$$

These are probabilities of miscoordination, and note that $p_i(0 \mid 1, \frac{2}{3}) \to 0$ and $p_i(1 \mid 0, \frac{2}{3}) \to 0$ as $\lambda \to 0$ (i.e, $\beta \to \infty$). That is, the probability that a player makes a mistake vanishes as the information costs become cheap.

Consider the ratio of the error probabilities: $R := \frac{p_i(1|0,\frac{2}{3})}{p_i(0|1,\frac{2}{3})}$. Due to the tractable formula (2), it is straightforward to see that $R \to 0$ as $\lambda \to 0$. It implies that the Markov chain stays at node 00 much longer than at node 11 (once it reaches the respective nodes) for a small λ, and the (unique) stationary distribution assigns high probability to node 00; at a unique equilibrium, players take action profile $(0, 0)$ at state $\theta = \frac{2}{3}$.

Another Information Cost. What if R is large for a small λ? Then, the Markov chain stays at node 11 much longer than at node 00, and the (unique) stationary distribution assigns high probability to node 11; players take action profile $(1, 1)$ with high probability. Then, the question is: Does there exist such an information cost function? If so, is it a realistic one?

Consider the following information cost function:

$$c_G(p) := \sum_{a_j,\theta} p(a_j, \theta) \underbrace{\Big(-\log\Big(1 - \frac{1}{2}\sum_{a_i}\big|\sqrt{p_i(a_i \mid a_j, \theta)} - \sqrt{p_i(a_i)}\big|^{\kappa}\Big)\Big)}_{=:\,G},$$

where $\kappa \geq 2$ is a parameter of curvature. Here, G measures how far the conditional $p_i(\cdot \mid a_j, \theta)$ is from the unconditional $p_i(\cdot)$, using the square root difference.

For the sake of comparison, I assume the same unconditional: $p_i(0) = \frac{1}{3}$ and $p_i(1) = \frac{2}{3}$. These $\lambda H, \lambda G$ can be regarded as functions of the conditional $p_i(\cdot \mid a_j, \theta)$. Moreover,

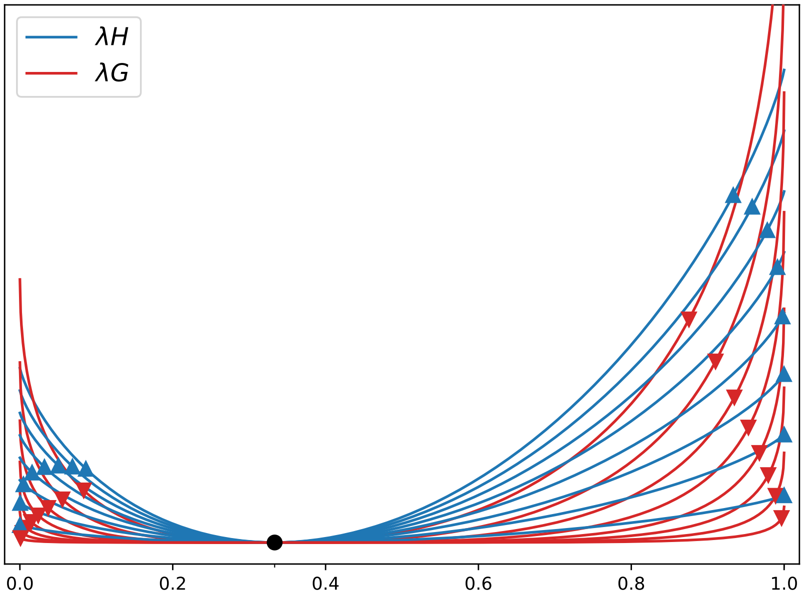

Figure 2: the functions $\lambda H, \lambda G$ as the scale λ vanishes

since actions are binary, they are regarded as functions of probability $p_i(0 \mid a_j, \theta)$. Their graphs are plotted in Figure 2, where the lower graphs are associated with smaller λ's.

Since both the expected payoff and the information cost $c = c_G$ are linear in marginal probabilities $p(a_j, \theta)$, to maximize the objective (1), it suffices to maximize the sub-objective $\sum_{a_i} p_i(a_i \mid a_j, \theta) u_i(a_i, a_j, \theta) - \lambda G(p_i(\cdot \mid a_j, \theta))$ for each pair (a_j, θ). Given the pair $(a_j, \theta) = (0, \frac{2}{3})$, for example, his subproblem is:

$$\max_{p_i(\cdot|0,\frac{2}{3})} \quad \underbrace{\left(u_i(0,0,\tfrac{2}{3}) - u_i(1,0,\tfrac{2}{3})\right)}_{=\frac{2}{3}} p_i(0 \mid 0, \tfrac{2}{3}) - \lambda G(p_i(\cdot \mid 0, \tfrac{2}{3})).$$

The optimal $p_i(0 \mid 0, \frac{2}{3})$ balances the marginal benefit $\frac{2}{3}$ with the marginal cost. In Figure 2, the marginal cost is equal to the slope of the function λG. Similarly, given the pair $(a_j, \theta) = (1, \frac{2}{3})$, the optimal $p_i(0 \mid 1, \frac{2}{3})$ balances the marginal benefit $u_i(0, 1, \frac{2}{3}) - u_i(1, 1, \frac{2}{3}) = -\frac{1}{3}$ with the marginal cost.

In Figure 2, both $\lambda H, \lambda G$ have slope $\frac{2}{3}$ at the right markers and slope $-\frac{1}{3}$ at the left. These markers indicate the optimal $p_i(0 \mid 0, \frac{2}{3})$ and $p_i(0 \mid 1, \frac{2}{3})$. The distances from the right markers to 1 are equal to probabilities $p_i(1 \mid 0, \frac{2}{3}) = 1 - p_i(0 \mid 0, \frac{2}{3})$, while those from the left markers to 0 are equal to probabilities $p_i(0 \mid 0, \frac{2}{3})$.

To see how the ratio R varies, I compare the convergence rate of the right markers with that of the left. As seen above, R is small under the cost λc_H for a small λ.

Indeed, the right ▲'s converge faster than the left ▲'s. In contrast, the left ▼'s converge faster than the right ▼'s, which implies that R is large under the cost λc_G for a small λ. Thus, the Markov chain of Figure 1 stays at node 11 much longer than at node 00, and its (unique) stationary distribution assigns high probability to node 11. Hence, the equilibrium action distribution at state $\theta = \frac{2}{3}$ assigns high probability to action profile $(1, 1)$ for a small λ.

> **Result.** *At state $\theta = \frac{2}{3}$, equilibrium $(0, 0)$ is uniquely selected by the information cost λc_H in the limit $\lambda \to 0$, whereas equilibrium $(1, 1)$ is uniquely selected by the information cost λc_G in the limit $\lambda \to 0$.*

Moreover, the converse is also true: At state $\theta = \frac{2}{3}$, neither action profile $(1, 0)$ nor $(0, 1)$ will be played with small probability for any small information costs. As illustrated in Figure 2, player i's conditional choice probabilities are such that $p_i(0 \mid 1, \frac{2}{3}) \to 0$ and $p_i(1 \mid 0, \frac{2}{3}) \to 0$ as $\lambda \to 0$. Hence, when player j takes action a_j player i will match his action $a_i = a_j$ with high probability, so that neither action profile $(1, 0)$ nor $(0, 1)$ will be played with small probability.

Properties of the Information Costs. Next, I review the properties of the information costs $\lambda c_H, \lambda c_G$. First, better information, which brings more expected payoff, is more costly in both cases. To see this, observe that both H, G are increasing in the distance from the unconditional $p_i(0) = \frac{1}{3}$. These information costs thus capture the trade-off of information acquisition. Second, the cost of acquiring no information is zero in both cases. Observe that the minimum cost is attained at the unconditional $p_i(0) = \frac{1}{3}$, and it is equal to zero. Third, both information costs approximate the complete-information game as the scale λ vanishes. Note that $\sup |H(\cdot)| < \infty$ and $\sup |G(\cdot)| < \infty$ for any parameter κ and that their sup norms are bounded: $\|\lambda c_H\| := \lambda \sup |c_H(p)| \leq \lambda \sup |H(\cdot)|$ and $\|\lambda c_G\| := \lambda \sup |c_G(p)| \leq \lambda \sup |G(\cdot)|$. As a result, $\|\lambda c_H\| \to 0$ and $\|\lambda c_G\| \to 0$ as $\lambda \to 0$. That is, a player will be able to acquire any information at cheap costs. These properties provide normative support for using the information costs λc_G. Hence, the result is not an artifact of ad hoc choice of information costs.

Difficulties Missed in the Example. This example misses several difficulties that will arise in the general model. In this example, the unconditional choice probabilities are assumed to be $p_i(0) = \frac{1}{3}$ and $p_i(1) = \frac{2}{3}$, but they must be endogenously determined in an equilibrium. This endogeneity complicates the analysis, creating another difficulty.

The values of the optimal conditional choice probabilities, marked in Figure 2, depend on those of the unconditional ones. If, for example, the unconditional $p_i(0)$ were close, or even equal, to zero then a player would not acquire any information, irrespective of the functional form of information costs. Then, the logic described above would no longer work. However, such a situation never emerges under the richness assumption, which requires that every action be strictly dominant for some state. It is satisfied in this example: Action 1 is strictly dominant at state $\theta = -\frac{1}{3}$ and action 0 is strictly dominant at state $\theta = \frac{5}{3}$.

3 The Model

3.1 The Basic Game

Consider a two-player game.[6] Let $I = \{1, 2\}$ be the set of players, where I write i for a generic player and $j \neq i$ for the opponent. Player i chooses an action a_i from a finite set A_i. As usual, let $A = A_1 \times A_2$. There is a finite set of payoff states Θ, where I write θ for a generic element. There is incomplete information about the state, and there is a full-support common prior $\psi \in \Delta(\Theta)$. Player i's payoff is given by a function $u_i : A \times \Theta \to \mathbb{R}$. This incomplete-information game $G := \langle I, (A_i)_i, (\Theta, \psi), (u_i)_i \rangle$ is called a **basic game**.

Let $G_\theta := \langle I, (A_i)_i, (u_i(\cdot, \theta))_i \rangle$ be the complete-information game with a state θ, and player i's best response to player j's action a_j is denoted by

$$\mathrm{BR}_i(a_j, \theta) := \operatorname*{argmax}_{a_i} u_i(a_i, a_j, \theta).$$

Action profile $a^* = (a_1^*, a_2^*)$ is a Nash equilibrium of the game G_θ if it is the case that $a_i^* \in \mathrm{BR}_i(a_j^*, \theta)$ for each i, and it is a strict Nash equilibrium if the best response $\mathrm{BR}_i(a_j^*, \theta)$ is a singleton for each i.

3.2 The Rational Inattention Game

Now I allow players to acquire—or pay attention to—information before choosing actions. Specifically, to make a better decision player i has an incentive to acquire information about the pair (a_j, θ). If the pair (a_j, θ) were known, he would choose

[6] I focus on a two-player game to ease exposition. All results extend to the general finite-player game.

the best response $\mathrm{BR}_i(a_j, \theta)$. If instead only noisy information is available, he has to choose an action a_i on the basis of that information. The form of noisy information exogenously determines his choice, or the action distribution $p_i(\cdot \mid a_j, \theta) \in \Delta(A_i)$.

Strategies. Rational inattention allows for a flexible approach: Player i acts as if he were choosing his conditional choice probabilities $p_i(\cdot \mid a_j, \theta)$ for each (a_j, θ). Refer to them as conditionals for short, and let $\mathcal{P}_i := (p_i(\cdot \mid a_j, \theta))_{a_j, \theta}$ be the conditionals. I will often refer to it as his strategy. His strategy reflects his choice of what information to acquire, and results in the form of mistakes. See Section 3.3 for more discussion on the modeling.

Player i cannot control the opponent's action nor a state when choosing his strategy. Hence, a marginal $p_{j\theta} \in \Delta(A_j \times \Theta)$ is exogenously given to him. Note that it must be consistent—i.e., its marginal on the state space Θ must coincide—with the prior ψ: $\mathrm{marg}_\Theta\, p_{j\theta} = \psi$.

Player i's strategy $\mathcal{P}_i$ and the marginal $p_{j\theta}$ define a joint $p \in \Delta(A \times \Theta)$, and yield the following expected payoff:

$$\mathbb{E}_p[u_i(a_i, a_j, \theta)] = \sum_{a_j, \theta} p_{j\theta}(a_j, \theta) \sum_{a_i} p_i(a_i \mid a_j, \theta) u_i(a_i, a_j, \theta).$$

Information Costs. Information is costly. For example, a player has to make an effort to acquire precise information, so that a concentrated conditional $p_i(\cdot \mid a_j, \theta)$ must be associated with high cost. Information costs for player i are modeled by a function $c_i : \Delta(A \times \Theta) \to \mathbb{R}_+$ that specifies the cost $c_i(p) \geq 0$ of a joint p. Given a marginal $p_{j\theta} \in \Delta(A_j \times \Theta)$, player i incurs the cost $c_i(p_i)$ to choose a strategy $\mathcal{P}_i$ that induces the joint p_i. Information costs can thus be interpreted as a function of a strategy $\mathcal{P}_i$ given the marginal $p_{j\theta}$, and I will often write $c_i(\mathcal{P}_i, p_{j\theta})$ instead.

I model information costs using a general nonparametric form, but I require that they satisfy some fundamental properties which rule out ad hoc settings and provide normative support for using the general form. The first requirement is that better information, which brings more expected payoff, be more costly. This is formalized as follows:

Definition 1. Given a marginal $p_{j\theta}$, player i's information cost c_i is **monotone** if for any strategies $\mathcal{P}_i$, $\mathcal{Q}_i$ such that the former is more profitable than the latter:

$$\sum_{a_j,\theta} p_{j\theta}(a_j,\theta) \sum_{a_i} p_i(a_i \mid a_j,\theta) u_i(a_i,a_j,\theta) > \sum_{a_j,\theta} p_{j\theta}(a_j,\theta) \sum_{a_i} q_i(a_i \mid a_j,\theta) u_i(a_i,a_j,\theta),$$

the cost of the former is (weakly) greater than that of the latter: $c_i(\mathcal{P}_i, p_{j\theta}) \geq c_i(\mathcal{Q}_i, p_{j\theta})$.

The second requirement is that the cost of acquiring no information—i.e., paying null attention—be zero. Note that a player who acquires no information chooses a fixed action a_i for any pair (a_j, θ). This is formalized as follows:

Definition 2. Player i's strategy $\mathcal{P}_i$ is a **null-information strategy** if for each $a_i \in A_i$, each $a_j \in A_j$, and each $\theta \in \Theta$,

$$p_i(a_i) = p_i(a_i \mid a_j, \theta),$$

where $p_i(a_i)$ is his unconditional choice probability of action a_i. Given any marginal $p_{j\theta}$, the information cost for the null-information strategy $\mathcal{P}_i$ is zero: $c_i(\mathcal{P}_i, p_{j\theta}) = 0$.

The Rational Inattention Game. In this paper, the pair $\langle G, c\rangle$ is referred to as a **rational inattention game**—an **RI game** for short—that consists of the basic game G and the information costs $c = (c_1, c_2)$.[7]

In the RI game $\langle G, c\rangle$, player i's objective is the expected payoff minus the information cost. Formally, his problem is: Given a marginal $p_{j\theta}$,

$$\max_{\mathcal{P}_i} \quad \sum_{a_j,\theta} p_{j\theta}(a_j,\theta) \sum_{a_i} p_i(a_i \mid a_j,\theta) u_i(a_i,a_j,\theta) - c_i(\mathcal{P}_i, p_{j\theta}). \tag{3}$$

An appropriate solution concept for the RI game $\langle G, c\rangle$ requires that each player's strategy be optimal given the marginals over the opponents' behavior and the state. Then, the relevant solution concept is the same as the Bayes correlated equilibrium of Bergemann and Morris (2016), but amended to account for information costs.

Definition 3. An **equilibrium** of an RI game $\langle G, c\rangle$ is a joint $p^* \in \Delta(A \times \Theta)$ such that:

(i) It is consistent with the prior ψ: $\mathrm{marg}_{\Theta}\, p^* = \psi$.

[7] I will use the term "rational attention" to describe this entire model class, while recognizing that the literature often uses this term to refers to the specific case of the Shannon information costs.

(ii) For each $i \in I$, player i's strategy $\mathcal{P}_i^* = (p^*(\cdot \mid a_j, \theta))_{a_j,\theta}$ is optimal to the marginal $p_{j\theta}^* = \mathrm{marg}_{A_j \times \Theta}\, p^*$: For any strategy $\mathcal{P}_i = (p_i(\cdot \mid a_j, \theta))_{a_j,\theta}$,

$$\sum_{a_j,\theta} p_{j\theta}^*(a_j, \theta) \sum_{a_i} p^*(a_i \mid a_j, \theta) u_i(a_i, a_j, \theta) - c_i(\mathcal{P}_i^*, p_{j\theta}^*)$$
$$\geq \sum_{a_j,\theta} p_{j\theta}^*(a_j, \theta) \sum_{a_i} p_i(a_i \mid a_j, \theta) u_i(a_i, a_j, \theta) - c_i(\mathcal{P}_i, p_{j\theta}^*).$$

3.3 Discussion of the Rational Inattention Game

In this subsection, I discuss interpretations of an RI game. My presentation of the RI game follows Sims (2003, 2010), in which players are modeled to choose conditional choice probabilities directly. In the RI game, a player's choice of action can be correlated with the opponent's, conditional on a state. In what follows, I discuss the interpretations of the correlation as well as information costs.

Robustness Interpretation. Consider two firms that may acquire information from an expert, a consulting firm, market research, and so on before choosing their prices. Consider a researcher who plans to model this situation. She knows that the firms may acquire information but does not know what information they can access. How should she model such a situation, especially if she wishes to make a prediction robust to the specification of endogenous information structures? Unless she knows that the firms' information is independent, she may not want to assume the independence; thus, she may need to take into consideration the possibility that the firms acquire more general information. Information costs may correspond to the costs of hiring an expert, hiring a consulting firm, or performing market research, and a priori, there is no reason for the market price for information or the effort cost to take the entropy form. Even when it comes to the issue of information processing, recent studies have observed that the Shannon information costs lead to potentially unrealistic conclusions (e.g., Fosgerau et al., 2018; Dean and Neligh, 2017). From the perspective of robustness, therefore, it is worth allowing information costs to take a general nonparametric form.

In conclusion, to examine the robustness to the specification of information acquisition, or endogenous information structures thereof, I allow players to acquire any information and their information costs to take a general nonparametric form.[8] How-

[8]This motivation is reminiscent of that of Weinstein and Yildiz (2007). While Carlsson and van Damme (1993) consider a particular way to perturb information in global games, Weinstein and Yildiz (2007) allow for any kind of perturbation.

ever, these costs must satisfy fundamental properties (Definitions 1 and 2) to rule out ad hoc settings and provide normative support for using the general information costs.

The same argument holds if players' information acquisition is interpreted as the bounded rationality resulting from their limited cognitive ability. The researcher knows that they are boundedly rational but not how boundedly rational they are. I allow for various kinds of bounded rationality. For instance, I do not require that players' mistakes be independent nor correlated in a particular way.

Nash-Equilibrium Interpretation. This interpretation originates in the self-enforcing requirement of Nash equilibria of complete-information games. In a Nash equilibrium, players have common knowledge about the actions that they are going to take. They have somehow reached a self-enforcing agreement on the Nash equilibrium play. In an RI game, players reach an agreement by acquiring information about each other's behavior as well as about a payoff state. Even according to this interpretation, a researcher should not restrict the way the players reach an agreement for robust predictions.

Mediator Interpretation. Lastly, I mention a mediator interpretation. Suppose that there is a mediator who sends information, such as action recommendations, to players and that players have to pay costs to receive the information. In this interpretation, the correlation of their choice of actions reduces to that of the information provided to them. The mutual optimality of Definition 3 also corresponds to incentive compatibility.

4 The Main Result

In this section, I present the main result. Roughly speaking, it states that: Any strict Nash equilibrium of the complete-information game G_θ with a true state θ is arbitrarily close to the unique equilibrium of a nearby RI game with some small information costs.

4.1 The Statement

I start with the notion of a nearby RI game with small information costs. If players' information costs were zero, they would acquire full information, so that the RI game would reduce to the underlying complete-information game G_θ with a true state θ. This leads to an idea to approximate the game G_θ, using an RI game $\langle G, c\rangle$ such that all players' information costs are "small" for any information. To formalize this idea, I use the sup norm:

$$\|c_i\| := \sup |c_i(\mathcal{P}_i, p_{j\theta})|,$$

where the supremum is taken over all pairs $(\mathcal{P}_i, p_{j\theta})$ such that the unconditional choice probability $p_i(a_i)$ of each action a_i is bounded away from zero, where the unconditional $p_i(\cdot)$ is defined by the pair $(\mathcal{P}_i, p_{j\theta})$.[9] Since no information can induce a player to take an action that is assigned the unconditional choice probability 0, the cost of such information, if any, should be infinite; therefore, the supremum is infinite. This requirement is to avoid these trivial exceptions, but it can be ignored since it will be automatically satisfied at any equilibrium (Lemma 3).

A **nearby RI game** refers to an RI game with all players' information costs being close to zero in the norm. The game G_θ is approximated by a sequence of nearby RI games $(\langle G, c^n \rangle)_n$ such that for each i, the sequence $(c_i^n)_n$ vanishes in the norm: $\|c_i^n\| \to 0$ as $n \to \infty$.

Two assumptions are made for the main result. First, I impose an assumption on the underlying complete-information games G_θ.

Assumption 1 (Strictness)**.** For each $i \in I$ and each $(a_j, \theta) \in A_j \times \Theta$, player i's best response $\mathrm{BR}_i(a_j, \theta)$ is a singleton.

This assumption is satisfied generically. When a player is indifferent between multiple actions, the tie can be broken by an arbitrarily small payoff perturbation. Similarly, any pure-strategy Nash equilibrium of a complete-information game G_θ can be made strict by an arbitrarily small payoff perturbation.

Second, I assume that the state space Θ is rich enough so that every action of every player is strictly dominant at some state.

Assumption 2 (Richness)**.** For each $i \in I$ and each $a_i \in A_i$, there exists $\theta^{a_i} \in \Theta$ such that action a_i is strictly dominant for player i:

$$u_i\Big(a_i, a_j, \theta^{a_i}\Big) > u_i\Big(a_i', a_j, \theta^{a_i}\Big) \qquad \forall\, a_i' \in A_i \setminus \{a_i\}, \quad \forall\, a_j \in A_j.$$

Now I state the main result—a **structure theorem**:

[9] In Section 2, if H, G had the unconditional $p_i(0) = 0$ and $p_i(1) = 1$, they would take infinite values at any conditional $p_i(0 \mid a_j, \theta) > 0$. The unconditional choice probabilities $p_i(0), p_i(1)$ are required to be bounded away from zero.

Theorem 1. *For each $\theta \in \Theta$, let $a_\theta \in A$ be any strict Nash equilibrium of the complete-information game G_θ. There exists a sequence of information costs $(c^n)_n$ such that:*

(i) For each $i \in I$, the information cost c_i^n satisfies Definitions 1 and 2 for each $n \in \mathbb{N}$, and $\|c_i^n\| \to 0$ as $n \to \infty$.

(ii) For each $n \in \mathbb{N}$, there exists a unique equilibrium p^n in the RI game $\langle G, c^n \rangle$.

(iii) For each $\theta \in \Theta$, $p^n(a_\theta \mid \theta) \to 1$ as $n \to \infty$.

Remark 1. Assumption 2 is a tight condition for Theorem 1. In Section 4.4, I show by example that in an RI game *without* the assumption, there can exist a null-information equilibrium and hence that there can be multiple equilibria.

4.2 The Proof

Whereas the literature has mainly studied coordination games and used their structures to characterize equilibria, I study general games, so that I take a new approach. First, I interpret an equilibrium of an RI game as a stationary distribution of some Markov chain and provide a condition for the uniqueness of the stationary distribution (Lemma 1), but the condition is imposed on conditional choice probabilities, which players choose endogenously. Next, I construct information costs so that players are willing to choose such conditional choice probabilities. As discussed at the end of Section 2, if some *unconditional* choice probability took an extreme value then any information costs (satisfying Definitions 1 and 2) could not induce the player to choose the desired conditional choice probabilities, since he would not acquire any information. As a preliminary step toward the construction of information costs, I give a partial characterization of equilibrium unconditional choice probabilities (Lemmas 2 and 3). With the help of this result, I construct the desired information costs (Lemma 4). Finally, I complete the proof by combining these lemmas.

The Markov Chain Interpretation of Equilibria of an RI Game. Let p^* be an equilibrium of an RI game $\langle G, c \rangle$. For now, I assume that given the marginal $p^*_{j\theta} = \operatorname{marg}_{A_j \times \Theta} p^*$, player i's equilibrium strategy $\mathcal{P}_i^* = (p^*(\cdot \mid a_j, \theta))_{a_j,\theta}$ is a *unique* solution to Problem (3). His information costs will be constructed so that this assumption is satisfied.

Player i's conditional $p^*(\cdot \mid a_j, \theta)$ can be interpreted as his noisy best responses to player j's action a_j at a state θ. His information cost c_i prevents him from taking the best response $\mathrm{BR}_i(a_j, \theta)$, causing mistakes. Along this interpretation, it is natural to examine whether the iteration of the noisy best responses leads to a suitable outcome. The iteration can be described as a Markov chain, and an equilibrium action distribution $p^*(\cdot \mid \theta) \in \Delta(A)$ can be sampled via a Markov chain Monte Carlo algorithm—specifically, the Gibbs Sampler.

For each $a \in A$, say that action profile $a' \in A$ is its neighbor if these two a, a' differ in at most one element, and let $\mathcal{N}(a)$ be the set of all neighbors of action profile a. By definition, each a is a neighbor of its own. For each θ, the algorithm defines a matrix $\mathbf{P}_\theta \in \mathbb{R}^{|A| \times |A|}$, where I write $\mathbf{P}_\theta(a, a')$ for its (a, a')-element, by

$$\mathbf{P}_\theta(a, a') := \begin{cases} \dfrac{1}{2} \displaystyle\sum_i p^*(a'_i \mid a_j, \theta) & \text{if } a' \in \mathcal{N}(a) \\ 0 & \text{if } a' \notin \mathcal{N}(a). \end{cases}$$

It is straightforward to show that this matrix $\mathbf{P}_\theta$ is a transition matrix, with each row summing to 1 (Lemma 1). Hence, it defines a (homogeneous) Markov chain.

Let $\mathbf{P}_\theta$ denote the Markov chain. Given any initial distribution $p^{(0)}(\cdot \mid \theta) \in \Delta(A)$, identified as a row vector of length $|A|$, the evolution of distributions is described by $p^{(t+1)}(\cdot \mid \theta) = p^{(t)}(\cdot \mid \theta)\mathbf{P}_\theta$ for each $t = 0, 1, 2, \ldots$. A distribution $p^*(\cdot \mid \theta)$ is said to be stationary if equation $p^*(\cdot \mid \theta) = p^*(\cdot \mid \theta)\mathbf{P}_\theta$ holds. Note that a stationary distribution is not necessarily unique in general. By construction, the equilibrium action distribution $p^*(\cdot \mid \theta)$ is stationary with respect to the Markov chain $\mathbf{P}_\theta$.[10] In the lemma below, a sufficient condition is provided for the Markov chain $\mathbf{P}_\theta$ to be globally stable—i.e., a stationary distribution is unique and independent of the choice of $p^{(0)}(\cdot \mid \theta)$. This result will be used to establish the equilibrium uniqueness.

Lemma 1. *For an equilibrium p^* of an RI game $\langle G, c \rangle$, the matrix $\mathbf{P}_\theta$ is a transition matrix for each $\theta \in \Theta$. For each $\theta \in \Theta$, if there exists $t \in \mathbb{N}$ such that given any $a, a' \in A$, there exists $b \in A$ such that $\mathbf{P}_\theta^t(a, b) > 0$ and $\mathbf{P}_\theta^t(a', b) > 0$, then the distribution $p^*(\cdot \mid \theta)$ is a unique stationary distribution of the Markov chain $\mathbf{P}_\theta$.*

The key to this lemma is that players may choose suboptimal actions with (possibly small but) non-zero probabilities—i.e., make errors. For example, in additive random utility models, this is the case whenever random utility shocks have wide support, as in the logit and probit cases (McFadden, 1984; Train, 2009).

[10] The converse is not true in general.

Remark 2. The Markov chain formulation of the equilibrium action distribution of the RI game is apparently dynamic, but it is just a mathematical interpretation of the equilibrium action distributions. The RI game itself is a static game, but interpreted as if it were a dynamic game.

Partial Characterization of Equilibrium Unconditional Choice Probabilities. Lemma 1 requires that players make mistakes with small but non-zero probabilities. Hence, I will now construct information costs so that they are willing to accept small but non-zero probabilities of mistakes. For this construction, as explained above, one needs to bound equilibrium *unconditional* choice probabilities. In what follows, I will discuss the bounds. A difficulty in this step is that player i's behavior depends not only on his information costs but also on player j's behavior—i.e., their strategic behavior is determined interdependently. This interdependence creates another difficulty. It may give rise to the possibility that all players choose null-information strategies in an equilibrium. Such an equilibrium is referred to as a **null-information equilibrium**. The existence of null-information equilibria can cause equilibrium multiplicity, which violates Theorem 1. At the end of this step, however, the non-existence of null-information equilibria is shown as a byproduct. That is, players will acquire some information in any equilibrium.

The key is Assumption 2. It is useful to overview why it drives players to acquire information, i.e., pay attention to each other's behavior as well to about a state. The logic is illustrated as follows. Suppose that player i (he) pays attention to a state. Then, player j (she) believes that his behavior depends on the state, which motivates her to pay attention to his behavior as well as to the state. This, in turn, implies that since he is aware of her attention to his attention (to the state), he is now motivated to pay attention to the attention of hers. This repeats ad infinitum. The infinite regress of paying attention may be reminiscent of the electronic mail game of Rubinstein (1989), but the attention allocations are endogenously chosen by the players.

Under Assumption 2, there exists a state θ^{a_i} such that action a_i is strictly dominant for player i. It is intuitive that conditional on the state θ^{a_i}, player i will assign high probability to action a_i when information costs are cheap.

Lemma 2. *For each $i \in I$, let $(c_i^n)_n$ be a sequence of player i's information costs such that $\|c_i^n\| \to 0$ as $n \to \infty$. For any marginal $p_{j\theta}^n$, player i's optimal strategy $\mathcal{P}_i^n$ under the* *information cost c_i^n is such that: For each $a_i \in A_i$,*

$$\lim_{n\to\infty} p_i^n\left(a_i \mid \theta^{a_i}\right) = 1,$$

where $p_i^n(a_i \mid \theta^{a_i})$ is the choice probability of action a_i conditional on the state θ^{a_i}.

Even though these conditional choice probabilities converge, uncondiditiontional ones do not necessarily. Conditional on the state θ^{a_i}, player i prefers to choose (strictly dominant) action a_i irrespective of player j's behavior, but his *unconditional* choice of action may depend on her behavior. Hence, his unconditional choice probabilities are determined interdependently. This interdependence complicates the analysis, making it difficult to test whether the unconditional ones converge as information costs vanish. For this reason, I use the upper and lower limit. Despite this indeterminacy, all unconditional choice probabilities must be bounded away from 0 and 1.

Lemma 3. *For each $i \in I$, let $(c_i^n)_n$ be a sequence of player i's information costs such that $\|c_i^n\| \to 0$ as $n \to \infty$. Then, given any marginal $p_{j\theta}^n$, player i's optimal strategy $\mathcal{P}_i^n$ under the information cost c_i^n is such that: For each $a_i \in A_i$,*

$$\psi\left(\theta^{a_i}\right) \leq \liminf_{n\to\infty} p_i^n(a_i) \leq \limsup_{n\to\infty} p_i^n(a_i) \leq 1 - \sum_{a_i' \neq a_i} \psi\left(\theta^{a_i'}\right),$$

where $p_i^n(a_i)$ is the unconditional choice probability of action a_i.

By Lemma 3, a player chooses a moderate—not close to 0 nor 1—unconditional choice probabilities and then pays attention to information as information costs become cheap. Such unconditional choice probabilities allow the player to flexibly adjust his choice probabilities from the unconditional ones to conditional ones by acquiring information. It suggests that their conditional ones are, to some extent, sensitive to perturbations of information costs.

As a corollary of Lemmas 2 and 3, I establish the non-existence of null-information equilibria when information costs are small enough.

Corollary 1. *For some $i \in I$ and for any small enough $\epsilon > 0$, let c_i be player i's information cost such that $\|c_i\| < \epsilon$. Then, there does not exist a null-information equilibrium.*

The Construction of Information Costs. Now that the preliminary result on the unconditional choice probabilities have been obtained, I will now construct information costs so that they induce players to choose particular conditional choice probabilities. I start with defining these conditional ones, named target strategies.

If player i's information costs were zero—i.e., he could acquire any information at the cost of zero—he would choose the best response $\mathrm{BR}_i(a_j, \theta)$ with probability 1 if player j chooses action a_j at a state θ. It is then intuitive to expect that when information costs are small, his optimal conditional choice probability should be close to the best response. This is the idea of target strategies.

To formalize it, I introduce the following notation: For each i and each (a_j, θ),

$$\mathbf{1}_{\mathrm{BR}_i(a_j,\theta)}(a_i) := \begin{cases} 1 & \text{if } a_i = \mathrm{BR}_i(a_j, \theta) \\ 0 & \text{if } a_i \neq \mathrm{BR}_i(a_j, \theta), \end{cases}$$

where the best response is unique under Assumption 1. Target strategies are defined as conditional choice probabilities that are close to this strategy $\mathbf{1}_{\mathrm{BR}_i(\cdot,\cdot)}(\cdot)$.

Definition 4. A sequence of **target joints** is a sequence of joints $(\tau^n)_n$ such that:

(i) For each $n \in \mathbb{N}$, the joint τ^n is consistent with the prior ψ: $\mathrm{marg}_\Theta\, \tau^n = \psi$.

(ii) For each $i \in I$ and each $(a_j, \theta) \in A_j \times \Theta$, $\tau^n(\cdot \mid a_j, \theta) \to \mathbf{1}_{\mathrm{BR}_i(a_j,\theta)}(\cdot)$ as $n \to \infty$, whenever the conditional choice probability is defined.

The conditionals $\mathcal{T}_i^n := (\tau^n(\cdot \mid a_j, \theta))_{a_j,\theta}$ is referred to as the **target strategy** for player i.

I show that a player is willing to choose his target strategy as his (unique) optimal strategy under some information cost.

Lemma 4. *Let $(\tau^n)_n$ be a sequence of target joints. For each $i \in I$, there exists a sequence of information costs $(c_i^n)_n$ such that $\|c_i^n\| \to 0$ as $n \to \infty$ and that player i's unique optimal strategy is the target strategy $\mathcal{T}_i^n$ given any marginal $p_{j\theta}^n$.*

The key idea to this lemma is simple, as I will now illustrate. Fix player j's action a_j and a state θ. Player i's relative loss from choosing a suboptimal action $a_i \neq \mathrm{BR}_i(a_j, \theta)$, compared to the optimal action $\mathrm{BR}_i(a_j, \theta)$, is at most $\Delta u_i(a_j, \theta) := \max_{a_i} u_i(a_i, a_j, \theta) - \min_{a_i'} u_i(a_i', a_j, \theta)$. Suppose that player i deviates by decreasing the probability to choose the suboptimal action a_i from the target $\tau(a_i \mid a_j, \theta)$ to zero. Then, the (additional) benefit from this deviation is at most $\tau(a_i \mid a_j, \theta)\Delta u_i(a_j, \theta)$. To prevent this deviation, it suffices to design the information costs for the deviating strategy greater than the benefit. Given such information costs, player i will accept small error probabilities (to choose suboptimal actions) implied by the target strategies.

However, this intuition misses a difficulty. Suppose that his *unconditional* choice probability of an suboptimal action a_i is closer, or even equal, to zero than the target level $\tau(a_i \mid a_j, \theta)$. Then, any information costs (satisfying Definitions 1 and 2) cannot induce him to increase the conditional choice probability of action a_i (to the target level). This is because he has to incur information costs to choose a less profitable strategy. However, it will be shown that such a situation never happens, using Lemmas 2 and 3.

The Proof of Theorem 1. Now that all lemmas needed have been obtained, Theorem 1 is proven. I will now summarize the idea, relegating the detail to Appendix A. For any strict Nash equilibrium a_θ of the complete-information game G_θ with a state θ, take a target joint τ assigning high probability to that action profile a_θ. By Lemma 4, there exists some information cost c_i such that the target strategy $\mathcal{T}_i$ is optimal to player i. As shown in Lemma 4, it is optimal to any marginal $p_{j\theta}$, not only to the marginal $\operatorname{marg}_{A_j \times \Theta} \tau$, under Assumption 2. Hence, an opponent's strategy affects a player's unconditional choice probabilities but not his conditional ones. Lemma 1 shows that the Markov chain associated to the target strategies has a unique stationary distribution $\tau(\cdot \mid \theta)$ for each θ, which leads to the target joint τ being the unique equilibrium, which completes the proof of Theorem 1.

4.3 The Converse

It is established that any strict Nash equilibrium of an underlying game can be uniquely selected by some information costs that vanish. Now I show its partial converse: Any non-Nash equilibrium action profile will not played in any equilibrium of a nearby RI game with any small information costs.

Theorem 2. *For each $\theta \in \Theta$, let $\tilde{a}_\theta \in A$ be any action profile that is* not *a Nash equilibrium of the complete-information game G_θ. Let $(c^n)_n$ be any sequence of information costs, and suppose that:*

(i) For each $i \in I$, $\|c_i^n\| \to 0$ as $n \to \infty$.

(ii) For each $n \in \mathbb{N}$, there exists an equilibrium p^n in the RI game $\langle G, c^n \rangle$.

Then, $p^n(\tilde{a}_\theta \mid \theta) \to 0$ as $n \to \infty$.

This result is intuitive. In the investment game of Section 2, action profiles $(1, 0)$ and $(0, 1)$ are rationalizable but not Nash equilibria at state $\theta = \frac{2}{3}$. Neither of them is selected by any sequence of information costs that vanish, because when player j takes action a_j, player i will match his action $a_i = a_j$ with high probability if his information costs are small. Theorem 2 is obtained by generalizing this observation.

4.4 Discussion

In this subsection, I discuss the implications of Theorems 1 and 2 as well as the related literature. I will also discuss the role of Assumption 2.

4.4.1 The Related Literature

In this paper, I refer to the main result as a **structure theorem**, because it is related to the structure theorem of Weinstein and Yildiz (2007). In what follows, I compare mine with theirs and then discuss other related studies.

The Structure Theorem with Endogenous Information Structure. Weinstein and Yildiz (2007) showed that any *rationalizable action profile* of an underlying game is uniquely selected by some *exogenous* information structure. Players' beliefs are specified exogenously, in a way that the tail beliefs assign positive probability of an opponent having first-order beliefs that imply dominant actions. This specification prevents players from acquiring approximate common knowledge of a payoff state; it enables any rationalizable action profile to be selected.

I show that any *strict Nash equilibrium* of an underlying game is uniquely selected by some *endogenous* information structure. Players' beliefs are determined at an equilibrium. I specify information cost functions. Irrespective of the functional forms, as the costs vanish, the players will be able to acquire more precise information, which drives the converse result: Any non-Nash equilibrium action profile is not played in any equilibrium of an RI game as information costs vanish.

There is another difference. For any rationalizable action profile a of an underlying game, Weinstein and Yildiz (2007) constructed a state space and a prior under which a type profile t can be realized such that players' interim beliefs are close to the common-knowledge beliefs and that action profile a is uniquely rationalizable. Note that the probability of the type profile t being realized may be small and that their structure theorem is silent about a play when another type profile $t' \neq t$ is realized. Hence, the probability that action profile a is uniquely rationalizable may be small. However, I show that any strict Nash equilibrium is played with probability close to 1 as a unique equilibrium of some nearby RI game.

Robustness to Incomplete Information. Kajii and Morris (1997) introduced a notion of robustness of a given Nash equilibrium of a complete-information game to incomplete information. They defined the Nash equilibrium to be robust if, in any incomplete-information game with a common prior assigning high probability to the event that the payoffs are as described in the complete-information game, the incomplete-information game will have a Bayesian Nash equilibrium that generates behavior close to the Nash equilibrium. Robust equilibria do not necessarily exist in general, but in a two-player-two-action coordination game, the risk-dominant equilibrium is uniquely robust.

In the investment game of Section 2, equilibrium $(0,0)$ is uniquely robust at state $\theta = \frac{2}{3}$. This does not contradict my result that action profile $(1,1)$ can be played under the information costs λc_G with probability close to 1 for a small λ, because action profile $(0,0)$ is still played with non-zero but small probability. Kajii and Morris (1997) argued that any incomplete-information game with the common prior $\psi(\theta = \frac{2}{3})$ close to 1 has a Bayesian Nash equilibrium at which action profile $(0,0)$ is played with high probability. It may have other equilibria, at which action profile $(1,1)$ can be played.

4.4.2 The Role of the Richness Assumption

I illustrate that Assumption 2 is a tight condition for Theorem 1, using simple examples.

Example 1. Consider the investment game of Section 2, but now let the prior be such that $\psi(\theta = \frac{2}{3}), \psi(\theta = \frac{5}{3}) > 0$, but $\psi(\theta = -\frac{1}{3}) = 0$. The payoffs are summarized in Table 2. Since action 1 is not strictly dominant in either state, Assumption 2 does not hold.

$\theta = \frac{2}{3}$	1	0
1	$\frac{1}{3}, \frac{1}{3}$	$-\frac{2}{3}, 0$
0	$0, -\frac{2}{3}$	$0, 0$

$\theta = \frac{5}{3}$	1	0
1	$-\frac{2}{3}, -\frac{2}{3}$	$-\frac{5}{3}, 0$
0	$0, -\frac{5}{3}$	$0, 0$

Table 2: Payoff Structure without Assumption 2

In this example, *there exists a null-information equilibrium such that action profile* $(0,0)$ *is played at each state* $\theta = \frac{2}{3}, \frac{5}{3}$. Thus, action profile $(1,1)$ is not a unique equilibrium outcome at state $\theta = \frac{2}{3}$. To see this, suppose that a marginal $p^*_{j\theta}$ is: $p^*_{j\theta}(0,\theta) = \psi(\theta)$ for each θ (i.e., $\text{marg}_{A_j} p^*_{j\theta}(0) = 1$). Then, player i will choose to pay null attention to a state because, since action profile $(0,0)$ is an equilibrium at each state, player i should choose action 0 irrespective of a state. Thus, his best response is a null-information strategy $p_i(0 \mid a_j, \theta) = p_i(0)$ for each (a_j, θ). The same logic is applied to the opponent. ■

Example 2. Consider the investment game of Section 2, but now let the prior be such that $\psi(\theta = -\frac{1}{3}), \psi(\theta = \frac{5}{3}) > 0$, but $\psi(\theta = \frac{2}{3}) = 0$. The payoffs are summarized in Table 3. Since actions $1, 0$ are strictly dominant at states $\theta = -\frac{1}{3}, \frac{5}{3}$ respectively, Assumption 2 holds.

$\theta = -\frac{1}{3}$	1	0
1	$\frac{4}{3}, \frac{4}{3}$	$\frac{1}{3}, 0$
0	$0, \frac{1}{3}$	$0, 0$

$\theta = \frac{5}{3}$	1	0
1	$-\frac{2}{3}, -\frac{2}{3}$	$-\frac{5}{3}, 0$
0	$0, -\frac{5}{3}$	$0, 0$

Table 3: Payoff Structure with Assumption 2

In this example, *there does not exist a null-information equilibrium if a player's information costs are small enough.* Since a player wants to match his action with the state irrespective of the opponent's behavior, an additional payoff that he gains from increasing the probability of correct matching is strictly positive and dependent only on the prior over the states. Thus, it exceeds information costs that are small enough (i.e., $\|c_i\| < \epsilon$ for a small enough $\epsilon > 0$), so that he will pay some attention to the state. ■

The lack of Assumption 2 may give rise to a null-information equilibrium by depriving players of incentives to acquire information about a state. Players care about the state but not the opponents' behavior in Example 2, but in the general model they care about both. In the general model, as discussed in Section 4.2, Assumption 2 triggers the infinite regress of players' paying attention to the opponents' attention. This makes their incentive of acquiring information sensitive to perturbations of information costs, thereby creating a room to influence their behavior.[11]

5 Conclusion

I studied an incomplete-information game with information acquisition. As in the theory of rational inattention, players are allowed to flexibly choose information to

[11] Assumption 2 is the key to the equilibrium uniqueness in the literature of equilibrium robustness to the exogenous specification of beliefs (e.g., Carlsson and van Damme, 1993; Morris and Shin, 2003; Weinstein and Yildiz, 2007), but its use in this literature is different from that in this paper. The literature has used this assumption for an infection argument.

acquire—or pay attention to—but such information is costly. These costs have often been modeled using Shannon entropy in the literature, but there are several reasons to depart from focusing on such costs. Information acquisition may correspond to the costs of hiring an expert or performing market research, and there is no reason for the price for information to take the entropy form; moreover, even when it comes to the issue of information processing, recent studies have observed that Shannon information costs lead to potentially unrealistic implications.

In this paper, I examined the robustness of equilibria to information acquisition: How sensitive are the equilibria to the specification of information costs or endogenous information structures thereof? I modeled information costs using a general nonparametric form, including the Shannon information costs as a special case. The main result is: *Any* strict Nash equilibrium of the complete-information game with a true state is arbitrarily close to the *unique* equilibrium of a nearby incomplete-information game with *some* small information costs. Its partial converse is also true: *Any* non-Nash equilibrium action profile will not be played in *any* equilibrium of a nearby incomplete-information game with *any* small information costs. Note that one can interpret information costs as a model of players' limited cognitive ability or bounded rationality. An implication of these results is, therefore, that the nature of players' information acquisition technology or bounded rationality must be carefully analyzed to make economically justifiable predictions.

Numerous questions remain open for future work. For example, one may wish to investigate how assumptions on costly information acquisition translate into equilibrium predictions. In particular, regarding the issue of information processing, there is growing literature on properties that may be required of attention cost functions (e.g., Caplin et al., 2017; Dean and Neligh, 2017). While these studies focus on single-agent decision-making, it is of interest to study how these properties may restrict equilibrium predictions.

One may also wish to find some unifying approach that informs us as to why risk-dominant equilibria are selected, for example, in global games (e.g., Carlsson and van Damme, 1993; Morris and Shin, 2003), evolutionary games (e.g., Kandori et al., 1993; Young, 1993), and quantal response equilibria with a particular property (e.g., McKelvey and Palfrey, 1995; Goeree et al., 2016). Appendix B offers a step toward this direction.

A Appendix: Omitted Proofs

Proof of Lemma 1. It is obvious that all elements of the matrix $\mathbf{P}_\theta$ are nonnegative. To see that each of its row sums to 1, notice that for each $a \in A$,

$$\sum_{a' \in A} \mathbf{P}_\theta(a, a') = \sum_{a' \in \mathcal{N}(a)} \mathbf{P}_\theta(a, a') + \sum_{a' \notin \mathcal{N}(a)} \underbrace{\mathbf{P}_\theta(a, a')}_{=0}$$
$$= \frac{1}{2} \sum_i \sum_{a'_i} p^*(a'_i \mid a_j, \theta) = 1.$$

Under the assumption for this lemma, the stability result of the Dobrushin coefficient can be applied, and thus the Markov chain $\mathbf{P}_\theta$ is globally stable for each θ.[12] Thus, it suffices to show that the equilibrium distribution $p^*(\cdot \mid \theta)$ is stationary to the Markov chain $\mathbf{P}_\theta$ for each θ. To see this, I prove the detailed balance: $p^*(a \mid \theta)\mathbf{P}_\theta(a, a') = p^*(a' \mid \theta)\mathbf{P}_\theta(a', a)$ for any $a, a' \in A$. This equation is obviously true in either case that $a = a'$ or that $a \notin \mathcal{N}(a')$, or equivalently $a' \notin \mathcal{N}(a)$. For any $a, a' \in A$ such that $a_i \neq a'_i$ and $a_j = a'_j$, it follows, from the Bayes' rule, that

$$p^*(a \mid \theta)\mathbf{P}_\theta(a, a') = \frac{1}{2} p^*(a_i, a_j \mid \theta) p^*(a'_i \mid a_j, \theta)$$
$$= \frac{1}{2} p^*(a_i \mid a_j, \theta) p^*(a'_i, a_j \mid \theta) = p^*(a' \mid \theta)\mathbf{P}_\theta(a', a).$$

From the detailed balance, it follows that for each $a \in A$,

$$\sum_{a' \in A} p^*(a' \mid \theta)\mathbf{P}_\theta(a', a) = \sum_{a' \in A} p^*(a \mid \theta)\mathbf{P}_\theta(a, a') = p^*(a \mid \theta),$$

which implies the stationarity. ■

Proof of Lemma 2. Suppose, for a contradiction, that there exist player i, action a_i^*, and state $\theta^{a_i^*}$ such that:

(i) Action a_i^* is strictly dominant for player i at the state $\theta^{a_i^*}$:

$$u_i\left(a_i^*, a_j, \theta^{a_i^*}\right) > u_i\left(a_i, a_j, \theta^{a_i^*}\right) \qquad \forall\, a_i \in A_i \setminus \{a_i^*\} \quad \forall\, a_j \in A_j.$$

[12] Indeed, the assumption for this lemma implies that the Dobrushin coefficient of the Markov kernel to the power of some t is strictly positive.

(ii) Given a marginal $p^n_{j\theta}$, player i's optimal strategy $\mathcal{P}^n_i := (p^n_i(\cdot \mid a_j, \theta))_{a_j,\theta}$ is:

$$\liminf_{n\to\infty} p^n_i\left(a^*_i \mid \theta^{a^*_i}\right) < 1, \tag{4}$$

where $p^n_i(a^*_i \mid \theta^{a^*_i})$ is the choice probability of action a^*_i conditional on the state $\theta^{a^*_i}$.

Inequality (4) is assumed for a contradiction. I am using the lower limit because the convergence of the sequence $(p^n_i(a^*_i \mid \theta^{a^*_i}))_n$ has not yet been established at this stage.

In this proof, I fix the marginal $p^n_{j\theta}$, so that I write $c_i(\mathcal{P}^n_i)$ for the cost $c_i(\mathcal{P}^n_i, p^n_{j\theta})$. Player i's objective at the strategy $\mathcal{P}^n_i$ is:

$$U_i(\mathcal{P}^n_i) := \sum_{a_j,\theta} p^n_{j\theta}(a_j,\theta) \sum_{a_i} p^n_i(a_i \mid a_j,\theta) u_i(a_i,a_j,\theta) - c^n_i(\mathcal{P}^n_i).$$

Now I define another strategy $\mathcal{Q}^n_i := (q^n_i(\cdot \mid a_j,\theta))_{a_j,\theta}$. Fix a small $\gamma > 0$, and consider the strategy $\mathcal{Q}^n_i$ such that:

(i) For each $n \in \mathbb{N}$ and each $(a_j,\theta) \in A_j \times (\Theta \setminus \{\theta^{a^*_i}\})$, $q^n_i(\cdot \mid a_j,\theta) = p^n_i(\cdot \mid a_j,\theta)$.

(ii) There exists $N_\gamma \in \mathbb{N}$ such that:

  (a) For each $n \geq N_\gamma$, $q^n_i(a_i) \geq \gamma$ for each $a_i \in A_i$, where the unconditional $q^n_i(\cdot)$ is defined by the strategy $\mathcal{Q}^n_i$ and the marginal $p^n_{j\theta}$.

  (b) For each $n \geq N_\gamma$, $q^n_i(a^*_i \mid \theta^{a^*_i}) = 1$.

Note that $q^n_i \neq p^n_i$ for each $n \geq N_\gamma$ and that $c^n_i(\mathcal{Q}^n_i) \to 0$ as $n \to \infty$.

Then, I will show that there exists some $n \in \mathbb{N}$ such that:

$$U_i(\mathcal{Q}^n_i) - U_i(\mathcal{P}^n_i) > 0, \tag{5}$$

which contradicts the optimality of the strategy $\mathcal{P}^n_i$.

To ease exposition, let, for each $n \in \mathbb{N}$ and each $(a_j,\theta) \in A_j \times \Theta$,

$$V^n_i(a_j,\theta) := \sum_{a_i} \left(q^n_i(a_i \mid a_j,\theta) - p^n_i(a_i \mid a_j,\theta)\right) u_i(a_i,a_j,\theta).$$

From Condition (i) of the strategy $\mathcal{Q}^n_i$, it follows that

$$\sum_{a_j,\theta} p^n_{j\theta}(a_j,\theta) V^n_i(a_j,\theta) = \sum_{a_j} p^n_{j\theta}\left(a_j,\theta^{a^*_i}\right) V^n_i\left(a_j,\theta^{a^*_i}\right)$$

$$= \sum_{a_j} \underbrace{\operatorname{marg}_\Theta p^n_{j\theta}\left(\theta^{a_i^*}\right)}_{=\psi(\theta^{a_i^*}) \quad \forall n} p^n_{j\theta}\left(a_j \mid \theta^{a_i^*}\right) V_i^n\left(a_j, \theta^{a_i^*}\right)$$

$$= \psi\left(\theta^{a_i^*}\right) \sum_{a_j} p^n_{j\theta}\left(a_j \mid \theta^{a_i^*}\right) V_i^n\left(a_j, \theta^{a_i^*}\right).$$

Then,

$$\begin{aligned} U_i(\mathcal{Q}_i^n) - U_i(\mathcal{P}_i^n) &= \sum_{a_j,\theta} p^n_{j\theta}(a_j, \theta) V_i^n(a_j, \theta) - \underbrace{\left(c_i^n(\mathcal{Q}_i^n) - c_i^n(\mathcal{P}_i^n)\right)}_{\leq c_i^n(\mathcal{Q}_i^n) \leq \|c_i^n\|} \\ &\geq \psi\left(\theta^{a_i^*}\right) \sum_{a_j} p^n_{j\theta}\left(a_j \mid \theta^{a_i^*}\right) V_i^n\left(a_j, \theta^{a_i^*}\right) - \|c_i^n\|. \end{aligned} \tag{6}$$

From Condition (ii)b, it follows that for each $n \geq N_\gamma$,

$$\begin{aligned} V_i^n\left(a_j, \theta^{a_i^*}\right) = &\left(1 - p_i^n\left(a_i^* \mid a_j, \theta^{a_i^*}\right)\right) u_i\left(a_i^*, a_j, \theta^{a_i^*}\right) \\ &- \sum_{a_i \neq a_i^*} p_i^n\left(a_i \mid a_j, \theta^{a_i^*}\right) u_i\left(a_i, a_j, \theta^{a_i^*}\right). \end{aligned} \tag{7}$$

Conditional on the state $\theta^{a_i^*}$, player i's second best response to player j's action a_j is denoted by $a_i^{**}(a_j) \in \operatorname{argmax}_{a_i \neq a_i^*} u_i(a_i, a_j, \theta^{a_i^*})$. Then,

$$\begin{aligned} \sum_{a_i \neq a_i^*} p_i^n\left(a_i \mid a_j, \theta^{a_i^*}\right) u_i\left(a_i, a_j, \theta^{a_i^*}\right) &\leq \sum_{a_i \neq a_i^*} p_i^n\left(a_i \mid a_j, \theta^{a_i^*}\right) u_i\left(a_i^{**}(a_j), a_j, \theta^{a_i^*}\right) \\ &= \left(1 - p_i^n\left(a_i^* \mid a_j, \theta^{a_i^*}\right)\right) u_i\left(a_i^{**}(a_j), a_j, \theta^{a_i^*}\right). \end{aligned}$$

Hence,

$$(7) \geq \left(1 - p_i^n\left(a_i^* \mid a_j, \theta^{a_i^*}\right)\right)\left(u_i\left(a_i^*, a_j, \theta^{a_i^*}\right) - u_i\left(a_i^{**}(a_j), a_j, \theta^{a_i^*}\right)\right). \tag{8}$$

Since action a_i^* is strictly dominant at the state $\theta^{a_i^*}$ and since $a_i^* \neq a_i^{**}(a_j')$ for each $a_j' \in A_j$, it follows that

$$\underline{u} := \min_{a_j'} \left\{ u_i\left(a_i^*, a_j', \theta^{a_i^*}\right) - u_i\left(a_i^{**}(a_j'), a_j', \theta^{a_i^*}\right) \right\} > 0.$$

From inequality (8), it follows that for each $a_j \in A_j$,

$$V_i^n\left(a_j, \theta^{a_i^*}\right) \geq \left(1 - p_i^n\left(a_i^* \mid a_j, \theta^{a_i^*}\right)\right) \underline{u}.$$

Hence,

$$
\begin{aligned}
(6) &> \psi\left(\theta^{a_i^*}\right) \sum_{a_j} p_{j\theta}^n\left(a_j \mid \theta^{a_i^*}\right)\left(1 - p_i^n\left(a_i^* \mid a_j, \theta^{a_i^*}\right)\right)\underline{u} - \|c_i^n\| \\
&= \psi\left(\theta^{a_i^*}\right)\left(1 - p_i^n\left(a_i^* \mid \theta^{a_i^*}\right)\right)\underline{u} - \|c_i^n\|. \qquad (9)
\end{aligned}
$$

Since $\lim_n \|c_i^n\| = 0$ by assumption and since the unconditional q_i^n is full support by Condition (ii)a of the strategy $\mathcal{Q}_i^n$, it follows that for any $\epsilon > 0$, there exists $N_\epsilon \in \mathbb{N}$ such that for each $n \geq N_\epsilon$,

$$
\|c_i^n\| < \epsilon.
$$

From inequality (4), it follows that for any $\delta > 0$, there exists $N_\delta \in \mathbb{N}$ such that for each $n \geq N_\delta$,

$$
p_i^n\left(a_i^* \mid \theta^{a_i^*}\right) < \liminf_{n\to\infty} p_i^n\left(a_i^* \mid \theta^{a_i^*}\right) + \delta.
$$

Let $\bar{N} := \max\{N_\epsilon, N_\delta, N_\gamma\}$, and then for each $n \geq \bar{N}$, it follows that

$$
\begin{aligned}
(9) &= \psi\left(\theta^{a_i^*}\right)\left(1 - p_i^n\left(a_i^* \mid \theta^{a_i^*}\right)\right)\underline{u} - \|c_i^n\| \\
&> \psi\left(\theta^{a_i^*}\right)\left(1 - \liminf_{n\to\infty} p_i^n\left(a_i^* \mid \theta^{a_i^*}\right) - \delta\right)\underline{u} - \epsilon \\
&= \psi\left(\theta^{a_i^*}\right)\left(1 - \liminf_{n\to\infty} p_i^n\left(a_i^* \mid \theta^{a_i^*}\right)\right)\underline{u} - \left(\psi\left(\theta^{a_i^*}\right)\underline{u}\delta + \epsilon\right).
\end{aligned}
$$

Note that the prior is $\psi(\theta^{a_i^*}) > 0$, that $1 - \liminf_n p_i^n(a_i^* \mid \theta^{a_i^*}) > 0$ by inequality (4), that $\underline{u} > 0$, and that both $\epsilon, \delta > 0$ are taken arbitrarily small. Hence, the last line must be strictly positive, which implies inequality (5), but it contradicts the optimality of the strategy $\mathcal{P}_i^n$. Hence, $\liminf_n p_i^n(a_i \mid \theta^{a_i}) = 1$ for each $a_i \in A_i$. It then follows that the limit exists and equals to 1. That is, $\lim_n p_i^n(a_i \mid \theta^{a_i}) = 1$. ■

Proof of Lemma 3. For each $a_i \in A_i$ and each $n \in \mathbb{N}$, note that

$$
\begin{aligned}
p_i^n(a_i) &= \sum_{a_j,\theta} p_{j\theta}^n(a_j, \theta) p_i^n(a_i \mid a_j, \theta) \\
&= \sum_\theta \underbrace{\operatorname{marg}_\Theta p_{j\theta}^n(a_j, \theta)}_{=\psi(\theta) \quad \forall n} \underbrace{\sum_{a_j} p_{j\theta}^n(a_j \mid \theta) p_i^n(a_i \mid a_j, \theta)}_{= p_i^n(a_i \mid \theta)}
\end{aligned}
$$

$$
\begin{aligned}
&= \sum_{\theta} \psi(\theta) p_i^n(a_i \mid \theta) \\
&\geq \psi\Big(\theta^{a_i}\Big) p_i^n\Big(a_i \mid \theta^{a_i}\Big),
\end{aligned}
$$

where the state θ^{a_i} is a state at which action a_i is strictly dominant for player i. Hence,

$$
\begin{aligned}
\liminf_{n\to\infty} p_i^n(a_i) &\geq \liminf_{n\to\infty} \psi\Big(\theta^{a_i}\Big) p_i^n\Big(a_i \mid \theta^{a_i}\Big) \\
&= \psi\Big(\theta^{a_i}\Big) \liminf_{n\to\infty} p_i^n\Big(a_i \mid \theta^{a_i}\Big) \\
&= \psi\Big(\theta^{a_i}\Big),
\end{aligned}
$$

where $\liminf_n p_i^n(a_i \mid \theta^{a_i}) = \lim_n p_i^n(a_i \mid \theta^{a_i}) = 1$ by Lemma 2.

For each $a_i \in A_i$ and each $n \in \mathbb{N}$,

$$
p_i^n(a_i) = 1 - \sum_{a_i' \neq a_i} p_i^n\Big(a_i'\Big).
$$

Hence,

$$
\begin{aligned}
\limsup_{n\to\infty} p_i^n(a_i) &= 1 - \liminf_{n\to\infty} \sum_{a_i' \neq a_i} p_i^n\Big(a_i'\Big) \\
&\leq 1 - \sum_{a_i' \neq a_i} \liminf_{n\to\infty} p_i^n\Big(a_i'\Big) \\
&\leq 1 - \sum_{a_i' \neq a_i} \psi\Big(\theta^{a_i'}\Big),
\end{aligned}
$$

where the first line follows from the fact that $\limsup_n(-z_n) = -\liminf_n z_n$ for any sequence of real numbers $(z_n)_n$, the second line from the superadditivity of the lower limit, and the third line from $\liminf_n p_i^n(a_i') \geq \psi(\theta^{a_i'})$, as shown above. ■

Proof of Corollary 1. Fix a small enough $\delta > 0$ that $\sum_{a_i' \neq a_i} \psi(\theta^{a_i'}) - 2\delta > 0$. Such a δ exists under Assumption 2. Since the information cost c_i is small enough, it follows that for each $a_i \in A_i$,

$$
p_i(a_i) < 1 - \sum_{a_i' \neq a_i} \psi\Big(\theta^{a_i'}\Big) + \delta < 1 - \delta < p_i(a_i \mid \theta^{a_i}),
$$

where the right inequality follows from Lemma 2 and the left inequality from Lemma 3. By acquiring information about a state, player i can increase his choice probability of the optimal action a_i conditional on the state θ^{a_i} by at least $\sum_{a'_i \neq a_i} \psi(\theta^{a'_i}) - 2\delta$. Since the information cost for this increase is arbitrarily small, player i will never choose a null-information strategy. ■

Proof of Lemma 4. To ease exposition, let $\eta_i = (a_j, \theta) \in A_j \times \Theta$. For example, player i's strategy is written as $\mathcal{P}_i = (p_i(\cdot \mid \eta_i))_{\eta_i}$.

A joint p_i is defined by player i's conditionals $\mathcal{P}_i$ and the marginal $p_{j\theta}$. Assume that the information cost c_i^n takes the following form:[13]

$$c_i^n(\mathcal{P}_i, p_{j\theta}) = \sum_{\eta_i} p_{j\theta}(\eta_i) g_i^n(p_i(\cdot \mid \eta_i), \eta_i).$$

Given such information costs, player i's payoff from conditionals $\mathcal{P}_i$ is equal to

$$\begin{aligned} &\sum_{\eta_i} p_{j\theta}(\eta_i) \sum_{a_i} p_i(a_i \mid \eta_i) u_i(a_i, \eta_i) - c_i^n(\mathcal{P}_i, p_{j\theta}) \\ &\quad = \sum_{\eta_i} p_{j\theta}(\eta_i) \Big\{ \sum_{a_i} p_i(a_i \mid \eta_i) u_i(a_i, \eta_i) - g_i^n(p_i(\cdot \mid \eta_i), \eta_i) \Big\}. \end{aligned}$$

Player i maximizes this function if and only if he solves the following subproblem: For each η_i with probability $p_{j\theta}(\eta_i) > 0$,

$$\max_{p_i(\cdot \mid \eta_i) \in \Delta(A_i)} \quad \sum_{a_i} p_i(a_i \mid \eta_i) u_i(a_i, \eta_i) - g_i^n(p_i(\cdot \mid \eta_i), \eta_i). \tag{10}$$

In the rest of this proof, I analyze how the objective (10) changes when player i changes his conditional choice probabilities from the target ones $\tau^n(\cdot \mid \eta_i)$ and then construct his information cost c_i^n such that any deviation from the target $\tau^n(\cdot \mid \eta_i)$ is penalized enough to prevent player i from doing so. Given the marginal $p_{j\theta}^n$, let $\mathcal{P}_i^n$ be player i's optimal strategy under the information cost c_i^n, and the unconditional $p_i^n(\cdot)$ is determined.

For each $a_i \in A_i$, each $\eta_i \in A_j \times \Theta$, and each $n \in \mathbb{N}$, let a function $d_{a_i}^n(\cdot, \eta_i) : [0, 1] \to \mathbb{R}_+$

[13] The Shannon information cost takes this form. This cost in the general model is:

$$\begin{aligned} &\sum_{\eta_i} \sum_{a_i} p(a_i, \eta_i)(\log p(a_i, \eta_i) - \log(p_i(a_i) p_{j\theta}(\eta_i))) \\ &\quad = \sum_{\eta_i} p_{j\theta}(\eta_i) \sum_{a_i} p_i(a_i \mid \eta_i)(\log p_i(a_i \mid \eta_i) - \log p_i(a_i)). \end{aligned}$$

Let $h(p_i(\cdot \mid \eta_i), \eta_i) = \sum_{a_i} p_i(a_i \mid \eta_i)(\log p_i(a_i \mid \eta_i) - \log p_i(a_i))$, and the Shannon information cost is equal to $\sum_{\eta_i} p_{j\theta}(\eta_i) h(p_i(\cdot \mid \eta_i), \eta_i)$.

be defined as follows:

$$d_{a_i}^n(x, \eta_i) := \begin{cases} \max\{x - \max\{\tau^n(a_i \mid \eta_i), p_i^n(a_i)\}, 0\} & \text{if } a_i = \text{BR}_i(\eta_i) \\ \max\{\min\{\tau^n(a_i \mid \eta_i), p_i^n(a_i)\} - x, 0\} & \text{if } a_i \neq \text{BR}_i(\eta_i). \end{cases}$$

Fix any $\rho > 0$, and suppose that player i takes, instead of the target $\tau^n(\cdot \mid \eta_i)$, a conditional $q_i(\cdot \mid \eta_i)$ such that:

$$\sum_{a_i} d_{a_i}^n(q_i(a_i \mid \eta_i), \eta_i) = \rho. \tag{11}$$

Then, the difference in the objective (10) is at most

$$\Delta u_i(\eta_i)\rho - \Big(g_i^n(q_i(\cdot \mid \eta_i), \eta_i) - g_i^n(\tau^n(\cdot \mid \eta_i), \eta_i)\Big), \tag{12}$$

where $\Delta u_i(\eta_i) := \max_{a_i} u_i(a_i, \eta_i) - \min_{a_i'} u_i(a_i', \eta_i)$. To induce player i to choose the target $\tau^n(\cdot \mid \eta_i)$, it suffices to construct the function g_i^n such that the value (12) is negative for any $q_i(\cdot \mid \eta_i)$. For each $a_i \in A_i$ and each $\eta_i \in A_j \times \Theta$, let

$$\delta_i^n(a_i \mid \eta_i) := \Big|\mathbf{1}_{\text{BR}_i(\eta_i)}(a_i) - \tau^n(a_i \mid \eta_i)\Big|.$$

Let the function g_i^n be defined such that for some parameter $\kappa > 0$,

$$g_i^n(q_i(\cdot \mid \eta_i), \eta_i) := \sum_{a_i} \big(\delta_i^n(a_i \mid \eta_i)\big)^{-\kappa} d_{a_i}(q_i(a_i \mid \eta_i), \eta_i),$$

where I write $(\delta_i^n(a_i \mid \eta_i))^{-\kappa} = 0$ if $\delta_i^n(a_i \mid \eta_i) = 0$ to ease exposition. Recall that the unconditional $p_i^n(a_i)$ is derived from the strategy $\mathcal{P}_i^n$ and the marginal $p_{j\theta}^n$. In general, it may or may not be greater than the target $\tau^n(a_i \mid \eta_i)$, but I will now show that whenever the costs are small enough, it is determined which is greater.

Lemma A.1. *Let $(\tau^n)_n$ be a sequence of target joints. Let $(c_i^n)_n$ be a sequence of information costs such that $\|c_i^n\| \to 0$ as $n \to \infty$. For any marginal $p_{j\theta}^n$, player i's optimal strategy $\mathcal{P}_i^n$ under the information cost c_i^n is such that for any large enough $n \in \mathbb{N}$*

$$\begin{aligned} p_i^n(a_i) \le \tau^n(a_i \mid a_j, \theta) & \quad \text{if } a_i = \text{BR}_i(a_j, \theta), \\ \tau^n(a_i \mid a_j, \theta) \le p_i^n(a_i) & \quad \text{if } a_i \neq \text{BR}_i(a_j, \theta). \end{aligned} \tag{13}$$

where $p_i^n(a_i)$ is the unconditional choice probability of action a_i.

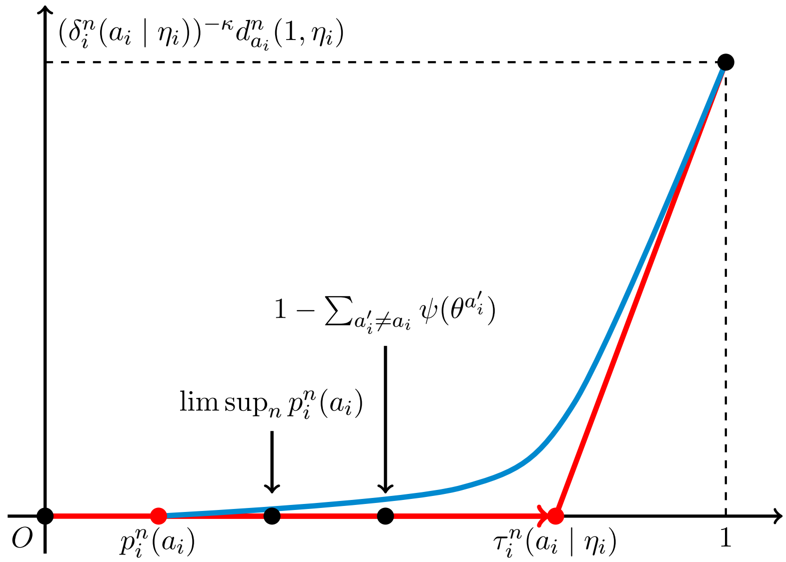

Figure 3: the Function g_i^n for Target $\tau_i^n(a_i \mid \eta_i)$ with Action $a_i = \mathrm{BR}_i(\eta_i)$

Proof. Under Assumption 2, there exists $\epsilon > 0$ such that for each $i \in I$ and each $\eta_i \in A_j \times \Theta$, the target $\tau^n(\cdot \mid \eta_i)$ is such that for each $a_i \in A_i$,

$$
\begin{aligned}
1 - \min_{a_i'} \psi\left(\theta^{a_i'}\right) + \epsilon \leq \tau^n(a_i \mid \eta_i) \qquad & \text{if } a_i = \mathrm{BR}_i(\eta_i), \\
\tau^n(a_i \mid \eta_i) \leq \min_{a_i'} \psi\left(\theta^{a_i'}\right) - \epsilon \qquad & \text{if } a_i \neq \mathrm{BR}_i(\eta_i).
\end{aligned}
$$

These inequalities with Lemma 3 show that for a small cost c_i, player i's unconditional $p_i^n(a_i)$ is such that:

$$
\begin{aligned}
\limsup_{n\to\infty} p_i^n(a_i) + \epsilon \leq \tau^n(a_i \mid \eta_i) \qquad & \text{if } a_i = \mathrm{BR}_i(\eta_i), \\
\tau^n(a_i \mid \eta_i) \leq \liminf_{n\to\infty} p_i^n(a_i) - \epsilon \qquad & \text{if } a_i \neq \mathrm{BR}_i(\eta_i).
\end{aligned}
$$

Since for any $\epsilon > 0$, there exists some $N \in \mathbb{N}$ such that for each $n \geq N$, $p_i^n(a_i) < \limsup_n p_i^n(a_i) + \epsilon$ and $p_i^n(a_i) > \liminf_n p_i^n(a_i) - \epsilon$, inequalities (13) follow. ■

From Lemma A.1, it follows that

$$
d^n_{a_i}(q_i(a_i \mid \eta_i), \eta_i) = \begin{cases} \max\{q_i(a_i \mid \eta_i) - \tau^n(a_i \mid \eta_i), 0\} & \text{if } a_i = \mathrm{BR}_i(\eta_i) \\ \max\{\tau^n(a_i \mid \eta_i) - q_i(a_i \mid \eta_i), 0\} & \text{if } a_i \neq \mathrm{BR}_i(\eta_i). \end{cases}
$$

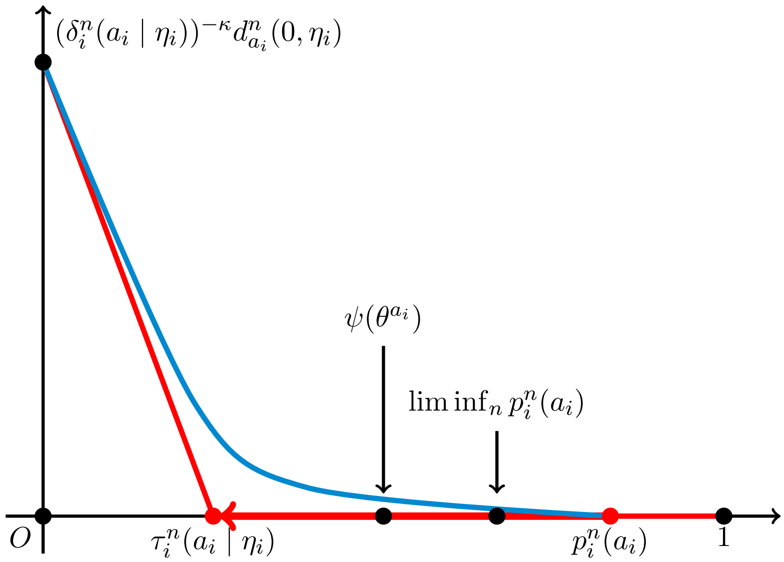

Figure 4: the Function g_i^n for Target $\tau_i^n(a_i \mid \eta_i)$ with Action $a_i \neq \mathrm{BR}_i(\eta_i)$

Figures 3 and 4 illustrate the construction described above.[14]

Since $d_{a_i}^n(\tau^n(a_i \mid \eta_i), \eta_i) = 0$ for each $a_i \in A_i$ and thus $g_i^n(\tau^n(\cdot \mid \eta_i), \eta_i) = 0$, it follows that

$$
\begin{aligned}
(12) &= \Delta u_i(\eta_i)\rho - g_i^n(q_i(\cdot \mid \eta_i), \eta_i) \\
&= \Delta u_i(\eta_i)\rho - \sum_{a_i} (\delta_i^n(a_i \mid \eta_i))^{-\kappa} d_{a_i}^n(q_i(a_i \mid \eta_i), \eta_i) \\
&\leq \Delta u_i(\eta_i)\rho - \min_{a_i : \delta_i^n(a_i|\eta_i) \neq 0} \left\{ \left(\delta_i^n(a_i \mid \eta_i)\right)^{-\kappa} \right\} \underbrace{\sum_{a_i} d_{a_i}^n(q_i(a_i \mid \eta_i), \eta_i)}_{= \rho \text{ by equation (11)}} \\
&= \left\{ \Delta u_i(\eta_i) - \min_{a_i : \delta_i^n(a_i|\eta_i) \neq 0} \left\{ \left(\delta_i^n(a_i \mid \eta_i)\right)^{-\kappa} \right\} \right\} \rho.
\end{aligned} \tag{14}
$$

For any large enough $\kappa > 0$, the last line is strictly negative for any $\rho > 0$. Hence, player i has no incentive to deviate from the target strategy $\mathcal{T}_i^n$ under the information cost c_i^n. That is, his optimal strategy $\mathcal{P}_i^n$ is equal to that under such costs.

Finally, I discuss the two requirements for information costs (Definitions 1 and 2). First, I consider the monotonicity. It is easy to modify the information cost constructed

[14]The information costs being constructed are illustrated in the red, but they do not have to be the lines with kinks. As is clear from the proof, the key to prevent a deviation is that the marginal information costs needed to change the conditional choice probabilities to the target levels are small (but not necessarily zero) but those needed to change more is large enough. Hence, smooth curves that are close to the red lines, such as the curves in the blue, also work.

above so that it satisfies the monotonicity. The idea is simple, as I will now state. Given a marginal $p_{j\theta}$, consider two strategies $\mathcal{P}_i := (p_i(\cdot \mid \eta_i))_{\eta_i}$ and $\mathcal{Q}_i := (q_i(\cdot \mid \eta_i))_{\eta_i}$, and let p, q be joints defined by $\mathcal{P}_i, \mathcal{Q}_i$ together with $p_{j\theta}$, respectively. Suppose that $\mathbb{E}_p[u_i] > \mathbb{E}_q[u_i]$ but $c_i(\mathcal{P}_i, p_{j\theta}) < c_i(\mathcal{Q}_i, p_{j\theta})$. Then, player i will never choose $\mathcal{Q}_i$. Without affecting his behavior, one can lower the cost $c_i(\mathcal{Q}_i, p_{j\theta})$ to $c_i(\mathcal{P}_i, p_{j\theta})$. Then, the monotonicity is satisfied (between the two). This idea can be easily generalized, and redefine the information cost as follows: For a marginal $p_{j\theta}$, let $c_i(\mathcal{Q}_i, p_{j\theta}) = \inf_{\mathcal{P}_i}\{c_i(\mathcal{P}_i, p_{j\theta}) : \mathbb{E}_p[u_i] \geq \mathbb{E}_q[u_i]\}$. Second, I show that the information cost for a null-information strategy is zero. For the null-information strategy, by definition, it must be that $p_i(a_i) = p_i(a_i \mid \eta_i)$ for each $a_i \in A_i$ and each $\eta_i \in A_j \times \Theta$. Hence,

$$\begin{aligned} d^n_{a_i}(p_i(a_i \mid \eta_i), \eta_i) &= d^n_{a_i}(p_i(a_i), \eta_i) \\ &= \begin{cases} \max\{p^n_i(a_i) - \max\{\tau^n(a_i \mid \eta_i), p^n_i(a_i)\}, 0\} & \text{if } a_i = \mathrm{BR}_i(\eta_i) \\ \max\{\min\{\tau^n(a_i \mid \eta_i), p^n_i(a_i)\} - p^n_i(a_i), 0\} & \text{if } a_i \neq \mathrm{BR}_i(\eta_i), \end{cases} \end{aligned}$$

both of which are zero, and thus the cost is zero. ■

Proof of Theorem 1. Let $(\tau^n)_n$ be a sequence of target joints such that for each θ, $\tau^n(\cdot \mid \theta) \to \mathbf{1}_{a_\theta}(\cdot)$ as $n \to \infty$, where such target joints exist because the conditionals $\tau^n(\cdot \mid a_j, \theta)$ also converge to the corresponding best response $\mathbf{1}_{\mathrm{BR}_i(a_j,\theta)}(\cdot)$ and thus satisfy the conditions in Definition 4. In fact, they can be constructed as follows. Fix a state θ, and define the set of ordered pairs $E_\theta = \{(a, a') \in A^2 : \exists\, i \in I,\, a'_i = \mathrm{BR}_i(a_j, \theta)\}$. The pair (A, E_θ) is a directed graph. By definition, any Nash equilibrium a^* of the game G_θ has no arrow with tail a^* (i.e., $(a^*, a) \notin E_\theta$ for each $a \neq a^*$). Suppose that the directed graph is acyclic.[15] Then, let $\tau^n(a_\theta) = 1 - \epsilon^n$ and $\tau^n(a^l_\theta) = (\epsilon^n)^l(1 - \epsilon^n)$ for all other Nash equilibria, labeled $a^1_\theta, a^2_\theta, \ldots, a^L_\theta$. Since a directed acyclic graph has a topological sort (i.e., a total order of its vertices such that for each arrow, the head comes before the tail), it is possible to distribute the remaining probability $(\epsilon^n)^{L+1}$ so that for each arrow, the ratio of the head's probability to the tail's probability is sufficiently large. Next, suppose that the directed graph is not acyclic. Then, let $\tau^n(a) = 0$ for each vertex a in a cycle. That is, for a cycle $(a^1, a^2), (a^2, a^3), \ldots, (a^L, a^1)$ let $\tau^n(a) = 0$ for all $a = a^1, a^2, \ldots, a^L$. Remove all arrows with either head or tail belonging to a cycle, and then the resulting restricted graph is acyclic, so that the construction in the first case is applicable.

[15] This condition is satisfied, for example, in the investment game of Section 2.

By Lemma 4, there exists a sequence of information costs $(c_i^n)_n$ such that $\|c_i^n\| \to 0$ as $n \to \infty$ and that under the information cost c_i^n, player i's unique optimal strategy is the target strategy $\mathcal{T}_i^n = (\tau^n(\cdot \mid a_j, \theta))_{a_j,\theta}$ given the marginal $\operatorname{marg}_{A_j \times \Theta} \tau^n$. By Lemma 1, the action distribution $\tau^n(\cdot \mid \theta)$ is a unique stationary distribution for each θ. Hence, the target joint τ^n has been shown to be the unique equilibrium of the RI game $\langle G, c^n \rangle$. ■

Proof of Theorem 2. Suppose, for a contradiction, that there exists a state $\tilde{\theta} \in \Theta$ and a non-Nash equilibrium action profile $\tilde{a} \in A$ of the complete-information game G_θ such that for any equilibrium p^n of the RI game $\langle G, c^n \rangle$,

$$\limsup_{n\to\infty} p^n\left(\tilde{a} \mid \tilde{\theta}\right) > 0.$$

Then, there exists a subsequence $(n_k)_{k\in\mathbb{N}}$ such that:

$$\lim_{k\to\infty} p^{n_k}\left(\tilde{a} \mid \tilde{\theta}\right) > 0. \tag{15}$$

Since action profile $\tilde{a}$ is not a Nash equilibrium at state $\tilde{\theta}$, there exists player i whose best response $\mathrm{BR}_i(\tilde{a}_j, \tilde{\theta})$ is such that for each $a_i \neq \mathrm{BR}_i(\tilde{a}_j, \tilde{\theta})$,

$$u_i\left(\mathrm{BR}_i\left(\tilde{a}_j, \tilde{\theta}\right), \tilde{a}_j, \tilde{\theta}\right) \geq u_i\left(a_i, \tilde{a}_j, \tilde{\theta}\right)$$

and in particular,

$$u_i\left(\mathrm{BR}_i\left(\tilde{a}_j, \tilde{\theta}\right), \tilde{a}_j, \tilde{\theta}\right) > u_i\left(\tilde{a}_i, \tilde{a}_j, \tilde{\theta}\right).$$

In this proof, I fix the equilibrium marginal $p_{j\theta}^{n_k}$, so that I write $c_i(\mathcal{P}_i^{n_k})$ for player i's information cost $c_i(\mathcal{P}_i^{n_k}, p_{j\theta}^{n_k})$. Player i's objective at the strategy $\mathcal{P}_i^{n_k}$ is:

$$U_i(\mathcal{P}_i^{n_k}) := \sum_{a_j,\theta} p_{j\theta}^{n_k}(a_j, \theta) \sum_{a_i} p_i^{n_k}(a_i \mid a_j, \theta) u_i(a_i, a_j, \theta) - c_i^{n_k}(\mathcal{P}_i^{n_k}).$$

Next, I define another strategy $\mathcal{Q}_i^{n_k} := (q_i^{n_k}(\cdot \mid a_j, \theta))_{a_j,\theta}$. Fix a small $\gamma > 0$, and consider the strategy $\mathcal{Q}_i^{n_k}$ such that:

(i) For each $k \in \mathbb{N}$ and each $(a_j, \theta) \in (A_j \times \Theta) \setminus \{(\tilde{a}_j, \tilde{\theta})\}$, $q_i^{n_k}(\cdot \mid a_j, \theta) = p_i^{n_k}(\cdot \mid a_j, \theta)$.

(ii) There exists $K_\gamma \in \mathbb{N}$ such that:

(a) For each $n \geq K_\gamma$, $q_i^{n_k}(a_i) \geq \gamma$ for each $a_i \in A_i$, where the unconditional $q_i^{n_k}(\cdot)$ is defined by the strategy $\mathcal{Q}_i^{n_k}$ and the marginal $p_{j\theta}^{n_k}$.

(b) For each $k \geq K_\gamma$, $q_i^{n_k}(\mathrm{BR}_i(\tilde{a}_j, \tilde{\theta}) \mid \tilde{a}_j, \tilde{\theta}) = 1$.

Note that $q_i^{n_k} \neq p_i^{n_k}$ for each $k \geq K_\gamma$ and that $c_i(\mathcal{Q}_i^{n_k}) \to 0$ as $k \to \infty$.

Then, I will show that

$$\liminf_{k \to \infty} \Big\{ U_i(\mathcal{Q}_i^{n_k}) - U_i(\mathcal{P}_i^{n_k}) \Big\} > 0,$$

which contradicts the optimality of the strategy $\mathcal{P}_i^{n_k}$.

To ease exposition, let, for each $k \in \mathbb{N}$ and each $(a_j, \theta) \in A_j \times \Theta$,

$$V_i^{n_k}(a_j, \theta) := \sum_{a_i} \Big(q_i^{n_k}(a_i \mid a_j, \theta) - p_i^{n_k}(a_i \mid a_j, \theta) \Big) u_i(a_i, a_j, \theta).$$

From Condition (i) of the strategy $\mathcal{Q}_i^{n_k}$, it follows that

$$\sum_{a_j, \theta} p_{j\theta}^{n_k}(a_j, \theta) V_i^{n_k}(a_j, \theta) = p_{j\theta}^{n_k}\Big(\tilde{a}_j, \tilde{\theta}\Big) V_i^{n_k}\Big(\tilde{a}_j, \tilde{\theta}\Big).$$

Then,

$$\begin{aligned} U_i(\mathcal{Q}_i^{n_k}) - U_i(\mathcal{P}_i^{n_k}) &= \sum_{a_j, \theta} p_{j\theta}^{n_k}(a_j, \theta) V_i^{n_k}(a_j, \theta) - \underbrace{\Big(c_i^{n_k}(\mathcal{Q}_i^{n_k}) - c_i^{n_k}(\mathcal{P}_i^{n_k}) \Big)}_{\leq c_i^{n_k}(\mathcal{Q}_i^{n_k}) \leq \|c_i^{n_k}\|}. \\ &\geq p_{j\theta}^{n_k}\Big(\tilde{a}_j, \tilde{\theta}\Big) V_i^{n_k}\Big(\tilde{a}_j, \tilde{\theta}\Big) - \|c_i^{n_k}\|. \end{aligned} \tag{16}$$

From Condition (ii)b, it follows that for each $k \geq K_\gamma$,

$$\begin{aligned} V_i^{n_k}\Big(\tilde{a}_j, \tilde{\theta}\Big) = \Big(1 - p_i^{n_k}\Big(\mathrm{BR}_i\Big(\tilde{a}_j, \tilde{\theta}\Big) \mid \tilde{a}_j, \tilde{\theta}\Big)\Big) u_i\Big(\mathrm{BR}_i\Big(\tilde{a}_j, \tilde{\theta}\Big), \tilde{a}_j, \tilde{\theta}\Big) \\ - \sum_{a_i \neq \mathrm{BR}_i(\tilde{a}_j, \tilde{\theta})} p_i^{n_k}\Big(a_i \mid \tilde{a}_j, \tilde{\theta}\Big) u_i\Big(a_i, \tilde{a}_j, \tilde{\theta}\Big). \end{aligned}$$

As noted above, player i's best response $\mathrm{BR}_i(\tilde{a}_j, \tilde{\theta})$ is such that for each $a_i \neq \mathrm{BR}_i(\tilde{a}_j, \tilde{\theta})$, $u_i(\mathrm{BR}_i(\tilde{a}_j, \tilde{\theta}), \tilde{a}_j, \tilde{\theta}) \geq u_i(a_i, \tilde{a}_j, \tilde{\theta})$ and that $u_i(\mathrm{BR}_i(\tilde{a}_j, \tilde{\theta}), \tilde{a}_j, \tilde{\theta}) > u_i(\tilde{a}_i, \tilde{a}_j, \tilde{\theta})$. It follows that

$$V_i^{n_k}\Big(\tilde{a}_j, \tilde{\theta}\Big) \geq p_i^{n_k}\Big(\tilde{a}_i \mid \tilde{a}_j, \tilde{\theta}\Big) \Big(u_i\Big(\mathrm{BR}_i(\tilde{a}_j, \tilde{\theta}), \tilde{a}_j, \tilde{\theta}\Big) - u_i\Big(\tilde{a}_i, \tilde{a}_j, \tilde{\theta}\Big)\Big).$$

Hence,

$$(16) \geq \underbrace{p_{j\theta}^{n_k}\left(\tilde{a}_j, \tilde{\theta}\right) p_i^{n_k}\left(\tilde{a}_i \mid \tilde{a}_j, \tilde{\theta}\right)}_{= p^{n_k}(\tilde{a} \mid \tilde{\theta})} \underbrace{\left(u_i\left(\mathrm{BR}_i(\tilde{a}_j, \tilde{\theta}), \tilde{a}_j, \tilde{\theta}\right) - u_i\left(\tilde{a}_i, \tilde{a}_j, \tilde{\theta}\right)\right)}_{> 0, \text{ for } \tilde{a}_i \neq \mathrm{BR}_i(\tilde{a}_j, \tilde{\theta}) \text{ by assumption}} - \|c_i^{n_k}\|.$$

Since $\lim_n \|c_i^n\| = 0$ by assumption, it follows that

$$\begin{aligned} &\liminf_{k\to\infty} \left\{ U_i(\mathcal{Q}_i^{n_k}) - U_i(\mathcal{P}_i^{n_k}) \right\} \\ &\quad \geq \liminf_{k\to\infty} (16) \\ &\quad \geq \underbrace{\lim_{k\to\infty} p^{n_k}\left(\tilde{a} \mid \tilde{\theta}\right)}_{> 0 \text{ by inequality (15)}} \underbrace{\left(u_i\left(\mathrm{BR}_i(\tilde{a}_j, \tilde{\theta}), \tilde{a}_j, \tilde{\theta}\right) - u_i\left(\tilde{a}_i, \tilde{a}_j, \tilde{\theta}\right)\right)}_{> 0} > 0, \end{aligned}$$

which contradicts the optimality of the strategy $\mathcal{P}_i^{n_k}$. ■

B Appendix: a Unified Perspective

In this appendix, I discuss how RI games are related to evolutionary games, quantal response equilibrium, mathematical models of statistical mechanics, and neural network models in machine learning. The unified perspective is *not* pedantic. It connects the apparently different models and deepens the understanding of an underlying mechanism of equilibrium selection, especially in coordination games. Moreover, the connections to statistical mechanics have potential applications of techniques developed in the field. To illustrate the power of these techniques, I will consider an issue raised by Sims (2010). In this appendix, **Shannon-RI** refers to rational inattention described by the Shannon information costs.

B.1 Shannon-RI Models and Statistical Mechanics

I consider a single-agent decision problem with the Shannon information costs. This model is studied by Matějka and McKay (2015). I discuss how it is connected to to statistical mechanics.[16] It is a preliminary step toward to rich connections discussed later.

[16] See also Steverson et al. (2017).

The Shannon-RI Problem. An agent chooses an action a from a finite set A. There is a finite set of payoff states Θ, and there is a full-support prior $\psi \in \Delta(\Theta)$. The agent's payoff is given by a function $v : A \times \Theta \to \mathbb{R}$. Analogously to the settings in Sections 2 and 3, he chooses conditionals $\mathcal{P} = (p(\cdot \mid \theta))_\theta$. Then, his problem is:

$$\max_{\mathcal{P}} \quad \sum_\theta \psi(\theta) \Bigg\{ \underbrace{\sum_a p(a \mid \theta) v(a, \theta)}_{\text{expected payoff}} - \underbrace{\lambda}_{\text{scale}} \underbrace{\sum_a p(a \mid \theta) \log \frac{p(a \mid \theta)}{p(a)}}_{\text{information costs}} \Bigg\}, \tag{17}$$

where $\lambda > 0$ is the scale factor. As before, let $p(a) = \sum_\theta \psi(\theta) p(a \mid \theta)$ be his unconditional of action a. Matějka and McKay (2015) solve this problem to show that the (unique) optimal conditionals $\mathcal{P}$ take the logit-like form:

$$p(a \mid \theta) = \frac{p(a) e^{\beta v(a,\theta)}}{\sum_{a'} p(a') e^{\beta v(a',\theta)}}, \tag{18}$$

where $\beta = \frac{1}{\lambda}$ is the reciprocal scale.

Statistical Mechanics. Let X be the finite set of possible microstates (unrelated to the state space Θ), and let $E(x)$ denote the energy level of a microstate x. Fix a constant $T > 0$, which is called temperature. An immediate implication of **the second law of thermodynamics** is that the equilibrium distribution of the microstates is given as the solution to the following problem:

$$\min_{q \in \Delta(X)} \quad \underbrace{\sum_x q(x) E(x)}_{\text{energy}} - \underbrace{T}_{\text{temperature}} \underbrace{\Big(- \sum_x q(x) \log q(x) \Big)}_{\text{entropy}}. \tag{19}$$

This objective is called the (Helmholtz) free energy, and it is known that the solution is given by the Boltzmann distribution—usually called the (multinomial) logit:

$$q(x) = \frac{e^{-\beta E(x)}}{\sum_{x'} e^{-\beta E(x')}}, \tag{20}$$

where $\beta = \frac{1}{T}$ is called the inverse temperature.[17] It is intentional to use the same notation β as in equation (18), since the temperature T will be identified as the scale λ.

[17]The Boltzmann constant is normalized to 1.

The Shannon-RI Problem Meets Statistical Mechanics. Problem (17) is equivalent to maximizing the expression in the bracket for each θ, while Problem (19) is equivalent to maximizing the negative free energy. Identifying payoff v as negative energy $-E$ and the scale λ as the temperature T, the two problems are equivalent except that Problem (17) has the denominator $p(a)$ in the log. To better understand this difference, replace $\log q(x)$ by $\log \frac{q(x)}{1/|X|}$ in Problem (19). Since this replacement ends up adding a constant to the objective, it yields the same solution (20). Indeed, if the unconditional, or the "a priori," choice probabilities are uniform (i.e., $p(a) = 1/|A|$ for each a) then the logit-*like* (18) reduces to the exact logit, thereby being consistent with the solution (20).[18] Hence, the essential difference between the two problems is that the a priori $p(\cdot)$ is endogenous in Problem (17) but exogenous in Problem (19).

B.2 Shannon-RI Games and Ising Models

Next, I study a strategic situation with multiple players. This situation is studied by Denti (2018). I analyze this model by connecting it to the Ising model in statistical mechanics.

The Shannon-RI Game. Consider the two-player investment game of Section 2. As will be discussed in Remark 3, it is straightforward to extend the analysis to general potential games. Player i's payoffs are summarized in Table 1 and represented by the function $u_i(a,\theta) = a_i a_j - \theta a_i$. Note that the investment game at each state θ is a potential game (Monderer and Shapley, 1996), and the potential function $v(\cdot,\theta) : A \to \mathbb{R}$ is such that $v(a,\theta) = a_1 a_2 - \theta \sum_i a_i$.

First, consider a centralized problem. There is a planner who maximizes, with respect to conditionals $\mathcal{P} = (p(\cdot \mid \theta))_\theta$, the expected potential minus the Shannon information costs:

$$\max_{\mathcal{P}} \quad \sum_\theta \psi(\theta)\Big\{\sum_a p(a \mid \theta)v(a,\theta) - \lambda \sum_a p(a \mid \theta)\log \frac{p(a \mid \theta)}{p_1(a_1)p_2(a_2)}\Big\}. \tag{21}$$

From Section B.1, it follows that the optimal conditional probabilities take the logit-like

[18] As this condition says that the agent chooses all actions with equal probabilities a priori—i.e., before conditioning on state θ—it is called the a priori homogeneity of actions (Matějka and McKay, 2015). In statistical mechanics, it is referred to as the postulate of equal a priori probability.

form:

$$p(a_1, a_2 \mid \theta) = \frac{p_1(a_1)p_2(a_2)e^{\beta v(a_1,a_2,\theta)}}{\sum_{a_1',a_2'} p_1(a_1')p_2(a_2')e^{\beta v(a_1',a_2',\theta)}}, \tag{22}$$

which is the unique solution to the centralized problem.

Second, consider a decentralized problem—i.e., the Shannon-RI game, in which each player maximizes his expected payoff minus his Shannon information costs. Given the marginal $p_{j\theta} \in \Delta(A_j \times \Theta)$, his problem is:

$$\max_{\mathcal{P}_i} \quad \sum_{a_j,\theta} p_{j\theta}(a_j, \theta)\Big\{\sum_{a_i} p_i(a_i \mid a_j, \theta)u_i(a_i, a_j, \theta) - \lambda H\Big\}, \tag{23}$$

with the Shannon information cost $H = \sum_{a_i} p(a_i \mid a_j, \theta)(\log p(a_i \mid a_j, \theta) - \log p_i(a_i))$. This H depends on the pair (a_j, θ), but by abuse of notation, I do not write the dependence explicitly as in Section 2.[19] Instead of discussing this game itself, I relate it to Problem (21). Probabilities (22) yield the following conditional choice probabilities:

$$p(a_i \mid a_j, \theta) = \frac{p_i(a_i)e^{\beta v(a_i,a_j,\theta)}}{p_i(1)e^{\beta v(1,a_j,\theta)} + p_i(0)e^{\beta v(0,a_j,\theta)}}. \tag{24}$$

Since $u_i(a, \theta) = v(a, \theta) + \theta a_j$ for each a and each θ and since the term θa_j is exogenous to player i, replacing v by u_i in equation (24) is equivalent to multiplying both numerator and denominator by $e^{\beta\theta a_j}$. Hence, the conditional choice probabilities (24) solve Problem (23). Hence, the unique equilibrium of the Shannon-RI game is given by the conditional choice probabilities (22), together with the prior ψ.[20]

The equivalence between the centralized problem and the decentralized problem is summarized as follows:

Proposition 1. *The unique equilibrium of the (decentralized) Shannon-RI game* (23) *is given as the (unique) solution* (22) *to the (centralized) Shannon-RI problem* (21).

The Ising Model. The (ferromagnetic) Ising model is a model of magnetism—roughly speaking, to explain why some material (e.g., iron) can be magnetized. I will review a simplified Ising model to make the connection to the Shannon-RI game clear.

[19] This game is exactly the same as the investment game with the Shannon information costs in Section 2.

[20] The uniqueness follows immediately from the Markov chain argument.

Suppose that there are two sites (or atoms), denoted $i = 1, 2$, and that there is a random variable $x_i \in \{0, 1\}$, called a spin, for each i.[21] A configuration $x = (x_1, x_2)$ specifies the spins. The energy of configuration x is given by $E(x) = -x_1 x_2 + h \sum_i x_i$ with external magnetic field $h \in \mathbb{R}$. Since a configuration corresponds to a microstate, the set of microstates is given by $X = \{0, 1\}^2$. In the Ising model, the distribution $q \in \Delta(X)$ is to minimize the following free energy with a fixed temperature $T > 0$:

$$\min_{q \in \Delta(X)} \quad \sum_{x_1, x_2} q(x_1, x_2) E(x_1, x_2) - T \sum_{x_1, x_2} q(x_1, x_2) \log q(x_1, x_2). \tag{25}$$

The solution is the Boltzmann distribution.

This Ising model assumes that the external magnetic field h is deterministic, but it is straightforward to generalize the Ising model to allow for random external magnetic fields.

The Shannon-RI Game Meets the Ising Model. The connection between the Shannon-RI game and the Ising model is as follows: player i is translated into site i, action a_i into spin x_i, action profile a into configuration x, the state θ into the external magnetic field h, the potential $v(\cdot, \theta)$ into the energy E, the scale λ into the temperature T. The only difference is, as in Section B.1, that the a priori $p_1(\cdot), p_2(\cdot)$ are endogenous in Problem (21) or (23) and exogenous in the Ising model (25).

Remark 3. The analysis can be easily extended to a general potential game. My argument for Proposition 1 does not depend on strategic complementarity of the investment game; it depends only on the existence of the potential $v(\cdot, \theta)$ for each θ. The same holds for the Ising model. The energy does not need to be specified as above, and it only needs to exist.

B.3 Shannon-RI Games and Machine Learning

As shown in Proposition 1, the unique equilibrium of the (decentralized) Shannon-RI game can be represented as the (unique) solution to the (centralized) Shannon-RI problem. This equivalence has been recognized and made use of in machine learning.[22]

[21] In statistical mechanics, it is usually assumed that a spin x_i takes value $+1$ (up) or -1 (down), which makes the Ising model symmetric with respect to 0. To make the connection to economics clearer, I assume that a spin x_i takes value 0 or 1, but this results in no essential difference.

[22] For example, Hertz et al. (1991) discuss machine-learning models, particularly neural network models, from a statistical mechanics perspective.

The Boltzmann Machine. The Boltzmann machine is a class of stochastic artificial neural network (Ackley et al., 1985; Hinton and Sejnowski, 1986). There are N sites, and site i has a bias $h_i \in \mathbb{R}$. These sites are connected by edges, and the edge between sites i and j is given a weight $w_{ij} \in \mathbb{R}$. Assume that $w_{ii} = 0$ (i.e., zero weight to itself) and that $w_{ij} = w_{ji}$ (i.e., symmetric weights) for all sites i, j. This network is characterized by the parameters $W = (w_{ij})_{ij}$ and $\mathbf{h} = (h_i)_i$. Each site i takes value $x_i \in \{0, 1\}$, and thus the network has the energy $E(x) = -\frac{1}{2}\sum_{i,j} w_{ij}x_ix_j + \sum_i h_ix_i$ for each $x = (x_i)_i$. It is now clear how the Boltzmann machine is related to the Ising model. For example, the energy is the same as that of the (two-site) Ising model when $N = 2$, $w_{12} = w_{21} = 1$, and $h_1 = h_2 = h$. Moreover, the Boltzmann machine assumes, as suggested by its name, that the (equilibrium) distribution is the Boltzmann distribution $\frac{e^{-\beta E(x)}}{\sum_{x'} e^{-\beta E(x')}}$.

Here arises a difficulty in computing the Boltzmann distribution given the parameters $(W, \mathbf{h})$. The Markov chain Monte Carlo (MCMC) simulation can be used to compute it. Given a spin configuration x, choose a site i randomly and flip spin x_i (from 0 to 1 or from 1 to 0) with probability $\frac{1}{1+e^{\beta\Delta_i}}$ while keeping all other spins x_{-i}, where $\beta > 0$ is the reciprocal of the parameter T and Δ_i is the energy change produced by the flip—i.e., $\Delta_i = E(0, x_{-i}) - E(1, x_{-i})$ if spin x_i flips from 0 to 1 and $\Delta_i = E(1, x_{-i}) - E(0, x_{-i})$ otherwise. Update the configuration based on the flip, and repeat this step again and again. The theoretical foundation of this MCMC simulation guarantees that after many steps, the resulting distribution of configurations will be arbitrarily close to the desired Boltzmann distribution. This result establishes the equivalence between the (centralized) Shannon-RI problem and the (decentralized) Shannon-RI game.

Simulated Annealing. As implied by equilibrium (22), the unique equilibrium of the Shannon-RI game will select action profile a that maximizes the potential at each state θ as the scale λ vanishes, and this observation is crucial to the equilibrium selection result of Denti (2018).

This result can be interpreted as simulated annealing (Kirkpatrick et al., 1983), which is a probabilistic technique to approximate the global optimum of a given function f. Take, as the function f, the potential $v(\cdot, \theta)$ given a state θ, or equivalently as the negative energy $-E$. The simulated annealing makes use of the fact that for any large β, the Boltzmann distribution is concentrated on the global optimum of the function; hence, the Boltzmann machine, together with the MCMC simulation, finds the potential maximizer with probability close to 1. This concentration of the Boltzmann distribution translates into the equilibrium action distribution at state θ being concentrated on the potential-maximizing action profile.

B.4 RI Games and Evolutionary Games

Now that I have connected the Shannon-RI game to statistical mechanics and machine learning, I connect RI games to evolutionary games with mutations (Kandori et al., 1993; Young, 1993; Bergin and Lipman, 1996). Such evolutionary games have selected the risk-dominance equilibrium of the two-player-two-action coordination game. The logic of the equilibrium selection is similar to simulated annealing (Kandori et al., 1993). As a result, the action profile, or the spin configuration, selected by these approaches coincide; in fact, the risk-dominance equilibrium maximizes the potential in the two-player-two-action coordination game.

The connection between the evolutionary game and statistical mechanics is pointed out by Blume (1993). He assumes that players' behavior follows the logit distribution (i.e., the Boltzmann distribution), and refers to the resulting evolutionary dynamics as the logit evolutionary dynamics. As discussed above, in the Shannon-RI game, the a priori $p_i(\cdot)$ is endogenous and thus not necessarily uniform, and thus the resulting behavior follows the logit-like, but possibly not exact logit, distribution.

B.5 Shannon-RI Games and Mean Field Approximation

In the Shannon-RI game of Section B.2, players' actions are correlated even conditional on a state θ—i.e., the spins are correlated in the Ising model. I study a variant game such that their actions are independent conditional on a state θ.

Yang (2015) studies a similar game in economics. The conditional independence is motivated by the analogy to global games, in which players receive noisy signals about a state θ but their signals (and thus actions) are conditionally independent (Carlsson and van Damme, 1993).

The Mean-Field Shannon-RI Game. Consider again the Shannon-RI game of Section B.2, but now I assume that players' choice of actions is independent conditional on a state θ. That is, I assume that player i's strategy is to choose conditionals $\check{\mathcal{P}}_i = (\check{p}_i(\cdot \mid \theta))_\theta$.

Given player j's strategy $\check{\mathcal{P}}_j$, player i's expected payoffs from action 1 is $\check{p}_j(1 \mid \theta) - \theta$, while that from action 0 is 0. Thus, his expected payoff from a strategy $\check{\mathcal{P}}_i$ is given by

$\check{p}_i(1 \mid \theta)(\check{p}_j(1 \mid \theta) - \theta)$. Accordingly, his problem is:

$$\max_{\check{\mathcal{P}}_i} \quad \sum_\theta \psi(\theta) \left\{ \check{p}_i(1 \mid \theta)(\check{p}_j(1 \mid \theta) - \theta) - \lambda \sum_{a_i} \check{p}_i(a_i \mid \theta) \log \frac{\check{p}_i(a_i \mid \theta)}{\check{p}_i(a_i)} \right\}. \tag{26}$$

This game is referred to as the **mean-field Shannon-RI game**. From the results (17) and (18), it follows that his optimal strategy is:

$$\check{p}_i(1 \mid \theta) = \frac{\check{p}_i(1) e^{\beta(\check{p}_j(1|\theta) - \theta)}}{\check{p}_i(1) e^{\beta(\check{p}_j(1|\theta) - \theta)} + \check{p}_i(0)}.$$

All equilibria of this game are symmetric.

Lemma B.1. *In the mean-field Shannon-RI game (26), all equilibria are symmetric:* $\check{p}_1(\cdot \mid \theta) = \check{p}_2(\cdot \mid \theta)$ *for each* $\theta \in \Theta$.

Proof. The first-order condition for Problem (26) is that for each θ,

$$m_j(\theta) - \lambda \left\{ \left(\log \frac{m_i(\theta)}{M_i} - \log \frac{1 - m_i(\theta)}{1 - M_i} \right) - \psi(\theta) \left(\frac{m_i(\theta)}{M_i} - \frac{1 - m_i(\theta)}{1 - M_i} \right) \right\} = 0.$$

This equation, as an implicit function, defines the best response r_θ as follows: $m_i(\theta) = r_\theta(m_j(\theta))$. By the implicit function theorem, it turns out that r_θ is increasing.

Suppose, for a contradiction, that there exists an asymmetric equilibrium. Then, there is some state θ with inequality $m_1(\theta) > m_2(\theta)$, but $m_2(\theta) = r_\theta(m_1(\theta)) > r_\theta(m_2(\theta)) = m_1(\theta)$, which is a contradiction. ■

By this lemma, I can focus on symmetric conditional choice probabilities $\check{p}(\cdot \mid \theta) := \check{p}_1(\cdot \mid \theta) = \check{p}_2(\cdot \mid \theta)$ for each θ, and I write the unconditional ones $\check{p}(\cdot) := \check{p}_1(\cdot) = \check{p}_2(\cdot)$. Substituting them into the above equation, I obtain the following equation:

$$\check{p}(1 \mid \theta) = \frac{\check{p}(1) e^{\beta(\check{p}(1|\theta) - \theta)}}{\check{p}(1) e^{\beta(\check{p}(1|\theta) - \theta)} + \check{p}(0)}. \tag{27}$$

Analogously to Lemmas 2 and 3, it follows that the equilibrium unconditional choice probabilities $\check{p}(0), \check{p}(1)$ are bounded away from 0 and 1 for any small λ (i.e., for any large β) under Assumption 2. Equation (27) has a unique solution for any large λ and multiple solutions for any small λ. This is illustrated in Figure 5, where the left- and right-hand sides are drawn as the diagonal and the curves, respectively. The argument so far is summarized as follows:

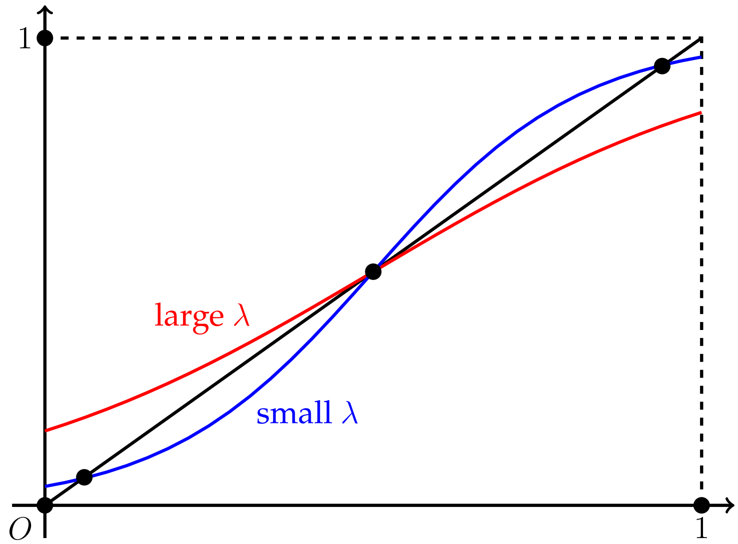

Figure 5: The Self-Consistency Equation

Proposition 2. *In the mean-field Shannon-RI game (26), all equilibrium strategies $\check{p} := \check{p}_1 = \check{p}_2$ satisfy equation (27). There are multiple equilibria for any small enough λ.*

The Mean-Field Ising Model. The Ising model of statistical mechanics is typically assumed to have many sites, and each site interacts with only some of the others. These complex interactions makes the model very difficult to solve, so that physicists have developed approximation techniques.

The **mean field approximation** is the simplest techniques among them. In the (two-site) Ising model of Section B.2, let $\check{q}_j$ be the distribution of spin x_j. Then, the mean of spin x_j is equal to $\check{q}_j(1) \times 1 + \check{q}_j(0) \times 0 = \check{q}_j(1)$. The mean field approximation ignores stochastic fluctuations and *replaces spin x_j by its mean $\check{q}_j(1)$*, thereby reducing the energy $E(x_i, x_j)$ into $E(x_i, \check{q}_j(1)) = -x_i(\check{q}_j(1) - h) + h\check{q}_j(1)$. The distribution $\check{q}_i$ is determined to minimize the free energy related to this energy:

$$\min_{\check{q}_i} \quad \sum_{x_i} \check{q}_i(x_i) E(x_i, \check{q}_j(1)) - T\Big(-\sum_{x_i} \check{q}_i(x_i) \log \check{q}_i(x_i)\Big). \tag{28}$$

From the results (19) and (20), it follows that the unique solution is:

$$\check{q}_i(1) = \frac{e^{\beta E(1, \check{q}_j(1))}}{e^{\beta E(1, \check{q}_j(1))} + e^{\beta E(1, \check{q}_j(0))}} = \frac{e^{\beta(\check{q}_j(1) - h)}}{e^{\beta(\check{q}_j(1) - h)} + 1}.$$

In statistical mechanics, the symmetry of spins x_1, x_2 is postulated, so that the outcome is described by the distribution $\check{q} := \check{q}_1 = \check{q}_2$ that satisfies the equation: $\check{q}(1) = \frac{e^{\beta(\check{q}(1)-h)}}{e^{\beta(\check{q}(1)-h)}+1}$.

The Mean-Field Shannon-RI Game Meets the Mean-Field Ising Model. The connection between the mean-field Shannon-RI game and the mean-field Ising model is clear. The only difference is, as in Section B.2, that the a priori $\check{p}(\cdot)$ is endogenous in Problem (26) and exogenous in Problem (28).

B.6 Mean-Field RI Games and QRE

I discuss the relationship between mean-field RI games and quantal response equilibrium—QRE for short (McKelvey and Palfrey, 1995; Goeree et al., 2016).

All equilibria of the mean-field Shannon-RI game (27) resemble the logit QRE. In particular, if the unconditional choice probabilities are such that $\check{p}(0) = \check{p}(1) = \frac{1}{2}$ then they exactly coincide with the logit QRE of the complete-information game G_θ. In general, the unconditional $\check{p}(\cdot)$ may not be uniform. In such non-uniform cases, the Shannon-RI equilibria are slightly different from the logit QRE. In what follows, I generalize this observation by using the result of Fosgerau et al. (2018) and then compare my main results (Theorems 1 and 2) with that of Haile et al. (2008).

Basics of QRE. Here I review the definition of the QRE notion. Recall the incomplete-information game G of Section 3. Fix any state $\theta_0 \in \Theta$, and consider the complete-information game $G_{\theta_0} = \langle I, (A_i)_i, (u_i(\cdot, \theta_0))_i \rangle$. Given player j's action distribution $p_j \in \Delta(A_j)$, let $u_i(a_i, p_j) := \sum_{a_j} p_j(a_j) u_i(a_i, a_j)$ be his expected payoff from action a_i.

The notion of QRE arises by introducting payoff noise. For each i, let $\hat{u}_i(a_i, p_j, \theta_0) := u_i(a_i, p_j, \theta_0) + \varepsilon_{ia_i}$, where payoff noise $\varepsilon_i = (\varepsilon_{ia_i})_{a_i}$ is drawn from a smooth joint density f_i with its support $\mathbb{R}^{|A_i|}$. Assume without loss of generality that noise ε_{ia_i} has mean zero for each a_i. Assume that noises are independent across players—i.e., noises $\varepsilon_i, \varepsilon_j$ are independent.

Given player j's action distribution p_j, player i is assumed to choose action a_i if and only if it brings the highest payoff: $\hat{u}_i(a_i, p_j, \theta_0) \geq \hat{u}_i(a'_i, p_j, \theta_0)$ for all actions a'_i. Let $R_{ia_i}(p_j) := \{\varepsilon_i : \hat{u}_i(a_i, p_j, \theta_0) \geq \hat{u}_i(a'_i, p_j, \theta_0)$ for all $a'_i\}$ be the set of realizations of payoff noises that induce player i to take action a_i. Then, the probability that he takes action a_i is $\int_{R_{ia_i}(p_j)} f_i(\varepsilon_i) d\varepsilon_i$. This defines player i's best-response action distribution to player j's p_j. The QRE requires that two action distributions be best responses to each other.

Recall that the Shannon-RI decision making leads to the logit-like choice probabilities (Matějka and McKay, 2015). Since the additive random utility model with the independent type I extreme-value errors leads to the logit choice probabilities (McFadden, 1984; Train, 2009), these two models are equivalent except that the prior-dependent shift—i.e., the unconditional choice probabilities $p(a)$ in equation (18). It is straightforward to translate this equivalence into that between the mean field Shannon-RI game and the logit QRE, in which players' payoff perturbations follow the independent type I extreme-value distribution.

The recent work by Fosgerau et al. (2018) generalizes the equivalence of these two decision-making models. They show that when information costs are modeled using a class of generalized entropies, the choice probabilities are observationally equivalent to some additive random utility model and vice versa. Hence, it is also straightforward to translate the generalized equivalence into that between the mean field RI game and the general QRE. This is a simple observation, but, to the best of my knowledge, it has not been point out in the literature.

The Main Result Revisited. Haile et al. (2008) relax the assumption that payoff noises are independent across players—i.e., they allow the noises to be correlated across players—and also allow the noises to be arbitrarily large. Then, they show that any action distribution can be played in a QRE. Such payoff noises correspond to large information costs in RI games. Indeed, any action profile, even non-rationalizable one, is played in an equilibrium if information costs are allowed to be arbitrarily large Denti (2018). These results are quite different from mine, since I assume that information costs are small, which correspond to small payoff noises. Accordingly, only strict Nash equilibria can be selected (Theorem 2).

B.7 Large Shannon-RI Games

So far I have considered how Shannon-RI models are related to other fields. A natural question to ask is: can economists use ideas, concepts, or techniques developed in these fields to solve new economic problems?

To illustrate the power of these techniques, I consider a "large" Shannon-RI game with many players. I will show that players are willing to acquire information only about a state θ at an equilibrium—i.e., the conditional independence of rational-inattention noisy behavior emerges endogenously.

Macroeconomic Implications. Sims (2010) points out an issue that arises in bringing rational inattention to equilibrium models is how to incorporate the correlation of agents' noisy behavior. An idea to model the potential correlation is that agents are allowed to choose correlations even conditional on a state (c.f., Afrouzi, 2018; Denti, 2018). Sims (2010) also points out that rational-inattention noise that is independent across agents will average out in macroeconomic behavior, whereas rational-inattention noise that is highly correlated will remain, thereby adding an additional source of macroeconomic randomness.

In this subsection, I will set up a stylized model of Shannon-RI investment game with many agents $N \to \infty$. Proposition 3 shows that their behavior will be independent conditional on a state irrespective of the scale of information costs, although they care about the distribution of the opponents' behavior. This result justifies the simplification assumption that ignores the correlation of agents' noisy behavior in macroeconomic setting.

The Large Shannon-RI Game. Consider an investment game with N players. As before, player i chooses an action a_i from the set $A_i = \{0, 1\}$, where action 1 is "Invest" and action 0 is "Not." As usual, let $A_{-i} = \prod_{j \neq i} A_j$ and $A = A_i \times A_{-i}$. There is a finite set of payoff states Θ, and there is a full-support common prior $\psi \in \Delta(\Theta)$.

Now I assume that player i's payoff depends on his own action a_i and the mean of the opponents' actions as well as on a state θ. To ease exposition, use $\frac{1}{N}\sum_{j \neq i} a_j$ as the mean of player $-i$'s actions, while recognizing that it should be divided by $N-1$, not by N. It does not affect the limiting results, as these two converge to the same value in the limit $N \to \infty$. Then, player i's payoff function is:

$$u_i(a_i, a_{-i}, \theta) = \left(\frac{1}{N}\sum_{j \neq i} a_j\right) a_i - \theta a_i = (\bar{a} - a_i)a_i - \theta a_i,$$

where $\bar{a} := \frac{1}{N}\sum_{j=1}^{N} a_j$. Then, his problem is:

$$\max_{\mathcal{P}_i} \quad \sum_{a_{-i},\theta} p_{-i\theta}(a_{-i}, \theta)\left\{\sum_{a_i} p_i(a_i \mid a_{-i}, \theta) u_i(a_i, a_{-i}, \theta) - \lambda H\right\},$$

with the Shannon information cost $H = \sum_{a_i} p(a_i \mid a_{-i}, \theta)(\log p(a_i \mid a_{-i}, \theta) - \log p_i(a_i))$. This H depends on the pair (a_j, θ), but by abuse of notation, I do not write the dependence explicitly as in Section 2.

The following proposition gives the full characterization of an equilibrium of the large Shannon-RI game with many players $N \to \infty$.

Proposition 3. *In the N-player Shannon-RI game, there exists a unique equilibrium, and it is symmetric: $\check{\check{p}}(\cdot \mid \theta) := \check{\check{p}}_i(\cdot \mid \theta)$ for each $i = 1, 2, \ldots, N$ and each $\theta \in \Theta$. In the limit $N \to \infty$, the equilibrium conditional choice probabilities are independent across players conditional on each state θ: for each $a_i \in A_i$, each $a_{-i} \in A_{-i}$, and each $\theta \in \Theta$*

$$\check{\check{p}}(a_i \mid a_{-i}, \theta) = \check{\check{p}}(a_i \mid \theta).$$

Moreover, the equilibrium conditional choice probabilities $\check{\check{p}}(\cdot \mid \theta)$ are uniquely determined as the solution to the following problem: for each $\theta \in \Theta$,

$$\max_{\check{\check{p}}(\cdot|\theta)} \quad -\frac{1}{2}\Big(\check{\check{p}}(1 \mid \theta)\Big)^2 + \lambda \log\Big(\check{\check{p}}(0) + \check{\check{p}}(1)e^{\beta(y-\theta)}\Big).$$

The first-order condition of this problem coincides with equation (27).

Toward the proof of Proposition 3, I review useful concepts and techniques and then prove the proposition. The concepts and techniques reviewed here are modified for use in the present context of economics.

Preliminaries. Consider the Shannon-RI problem (17). For each θ, the **state-wise partition function** $Z_\theta : \mathbb{R}_{++} \to \mathbb{R}$ is defined as the denominator of equation (18):

$$Z_\theta(\beta) := \sum_a p(a)e^{\beta v(a,\theta)}.$$

This function is named after the partition function in statistical mechanics. It may seem to be just the normalization constant in equation (18), but it contains important information about Problem (17). For example, the maximum value (17) is recovered from the state-wise partition functions, as shown below.

Lemma B.2. *Consider the state-wise partition function Z_θ defined above. Then,*

$$(17) = \frac{1}{\beta}\sum_\theta \psi(\theta) \log Z_\theta(\beta).$$

Proof. From the solution (18), it follows that $p(a \mid \theta) = p(a_i)e^{\beta v(a,\theta)}/Z_\theta(\beta)$. Then,

$$(17) = \sum_\theta \psi(\theta)\left\{\sum_{a_i} p(a \mid \theta)v(a,\theta) - \frac{1}{\beta}\sum_a p(a \mid \theta)\log\frac{e^{\beta v(a,\theta)}}{Z_\theta(\beta)}\right\}$$

$$= \sum_\theta \psi(\theta) \left\{ \frac{1}{\beta} \sum_a p(a \mid \theta) \log Z_\theta(\beta) \right\}$$
$$= \frac{1}{\beta} \sum_\theta \psi(\theta) \log Z_\theta(\beta),$$

which completes the proof. ■

Next, I review a technique called the Hubbard-Stratonovich transformation in statistical mechanics. This result follows from a simple observation of the Gaussian integral, but it turns out to be useful in analyzing large Shannon-RI game.

Lemma B.3. *Let $x_1, x_2, \ldots, x_N \in \{0, 1\}$ and $\bar{x} = \frac{1}{N}\sum_{i=1}^N x_i$. For each $\beta > 0$,*

$$e^{\frac{N\beta}{2}\bar{x}^2} = \sqrt{\frac{N\beta}{2\pi}} \int_{-\infty}^{\infty} e^{-\frac{N\beta}{2}t^2 + N\beta\bar{x}t} dt.$$

Proof. Since $-\frac{N\beta}{2}t^2 + N\beta\bar{x}t = -\frac{N\beta}{2}(t-\bar{x})^2 + \frac{N\beta}{2}\bar{x}^2$ by completing the square, it follows that

$$e^{\frac{N\beta}{2}\bar{x}^2} \underbrace{\int_{-\infty}^{\infty} e^{-\frac{N\beta}{2}(t-\bar{x})^2} dt}_{=\sqrt{2\pi/N\beta}} = \int_{-\infty}^{\infty} e^{-\frac{N\beta}{2}t^2 + N\beta\bar{x}t} dt,$$

where the integral on the left-hand side follows from the Gaussian integral. ■

Proof of Proposition 3. The N-player Shannon-RI game has the potential: for each θ,

$$V(a, \theta) = \frac{1}{2N}\left(\sum_{i,j} a_i a_j - \sum_i a_i^2\right) - \theta \sum_i a_i$$
$$= \frac{1}{2}N\bar{a}^2 - \left(\theta + \frac{1}{2N}\right) \sum_i a_i,$$

where I use the fact that $\sum_i a_i^2 = \sum_i a_i$.

Note that the equivalence result of Proposition 1 can be immediately extended to the N-player model. Hence, it suffices to study the centralized problem:

$$\max_{\mathcal{P}} \quad \sum_\theta \psi(\theta) \left\{ \sum_a p(a \mid \theta) V(a, \theta) - \lambda \sum_a p(a \mid \theta) \log \frac{p(a \mid \theta)}{\prod_i p_i(a_i)} \right\}. \tag{29}$$

Since the payoff structure is symmetric, there is a unique solution of this problem; hence, the equilibrium conditional choice probabilities are described as a distribution $p^N(\cdot \mid \theta) \in \Delta(\{0,1\})$: for each θ,

$$p^N(\cdot \mid \theta) := p_1(\cdot \mid \theta) = \cdots = p_N(\cdot \mid \theta),$$

where the superscript N denotes the number of players.

For each θ, the state-wise partition function $Z_\theta : \mathbb{R}_{++} \to \mathbb{R}$ is:

$$\begin{aligned} Z_\theta(\beta) &= \sum_a \Big(\prod_i p^N(a_i)\Big) e^{\beta V(a,\theta)} \\ &= \sum_a \Big(\prod_i p^N(a_i)\Big) e^{\frac{N\beta}{2}\bar{a}^2} e^{-\beta(\theta + \frac{1}{2N})\sum_i a_i}, \end{aligned} \tag{30}$$

where the summation $\sum_a$ takes over all action profiles—i.e., 2^N patterns. Since it involves many exponential terms, I transform expression (30) into another. By Lemma B.3,

$$e^{\frac{N\beta}{2}\bar{a}^2} = \sqrt{\frac{N\beta}{2\pi}} \int_{-\infty}^{\infty} e^{-\frac{N\beta}{2}t^2 + \beta t \sum_i a_i} dt.$$

Hence,

$$\begin{aligned} (30) &= \sqrt{\frac{N\beta}{2\pi}} \int_{-\infty}^{\infty} e^{-\frac{N\beta}{2}t^2} \sum_a \Big(\prod_i p^N(a_i)\Big) e^{\beta(t-\theta-\frac{1}{2N})\sum_i a_i} dt \\ &= \sqrt{\frac{N\beta}{2\pi}} \int_{-\infty}^{\infty} e^{-\frac{N\beta}{2}t^2} \sum_a \Big(\prod_i p^N(a_i) e^{\beta(t-\theta-\frac{1}{2N})a_i}\Big) dt \\ &= \sqrt{\frac{N\beta}{2\pi}} \int_{-\infty}^{\infty} e^{-\frac{N\beta}{2}t^2} \prod_i \Big(\sum_{a_i} p^N(a_i) e^{\beta(t-\theta-\frac{1}{2N})a_i}\Big) dt \\ &= \sqrt{\frac{N\beta}{2\pi}} \int_{-\infty}^{\infty} e^{-\frac{N\beta}{2}t^2} \Big(p^N(0) + p^N(1) e^{\beta(t-\theta-\frac{1}{2N})}\Big)^N dt \\ &= \sqrt{\frac{N\beta}{2\pi}} \int_{-\infty}^{\infty} \exp\Big\{ N\beta \underbrace{\left(-\frac{1}{2}t^2 + \frac{1}{\beta}\log\Big(p^N(0) + p^N(1)e^{\beta(t-\theta-\frac{1}{2N})}\Big)\right)}_{=:\, \phi_\theta(t)} \Big\} dt. \end{aligned} \tag{31}$$

To evaluate this integral, the saddle-point method can be used:

$$(31) = \sqrt{\frac{N\beta}{2\pi}} e^{N\beta \max_t\{\phi_\theta(t)\} + O(1)}.$$

By Lemma B.2,

$$\begin{aligned}(29) &= \frac{1}{\beta} \sum_\theta \psi(\theta) \log Z_\theta(\beta) \\ &= \sum_\theta \psi(\theta) \left\{ \frac{1}{\beta} \log \sqrt{\frac{N\beta}{2\pi}} + N \max_t \{\phi_\theta(t)\} + O(1) \right\}. \end{aligned} \tag{32}$$

Then, the value (32)$/N$ is such that as $N \to \infty$,

$$\frac{(32)}{N} \to \sum_\theta \psi(\theta) \max_{\breve{\check{p}}(\cdot|\theta)} \underbrace{\left\{ -\frac{1}{2}\Big(\breve{\check{p}}(1 \mid \theta)\Big)^2 + \frac{1}{\beta} \log\Big(\breve{\check{p}}(0 \mid \theta) + \breve{\check{p}}(1 \mid \theta) e^{\beta(\breve{\check{p}}(1|\theta) - \theta)}\Big) \right\}}_{=: \varphi_\theta(\breve{\check{p}}(\cdot \mid \theta))}.$$

The first-order condition for the state-by-state problem is as follows:

$$\breve{\check{p}}(1 \mid \theta) = \frac{\breve{\check{p}}(1) e^{\beta(\breve{\check{p}}(1|\theta) - \theta)}}{\breve{\check{p}}(1) e^{\beta(\breve{\check{p}}(1|\theta) - \theta)} + \breve{\check{p}}(0)},$$

which is exactly the same as equation (27).

The equilibrium conditional $\breve{\check{p}}(\cdot \mid \theta)$ minimizes the function φ_θ. Not all solutions to equation (27) constitute equilibria of the large Shannon-RI game, so that this proposition is different from Proposition 2. ■

References

Ackley, D. H., Hinton, G. E., and Sejnowski, T. J. (1985). A learning algorithm for Boltzmann machines. *Cognitive Science*, 9(1):147–169.

Afrouzi, H. (2018). Strategic inattention, inflation dynamics and the non-neutrality of money.

Angeletos, G.-M., Hellwig, C., and Pavan, A. (2006). Signaling in a global game: Coordination and policy traps. *Journal of Political economy*, 114(3):452–484.

Angeletos, G.-M., Hellwig, C., and Pavan, A. (2007). Dynamic global games of regime change: Learning, multiplicity, and the timing of attacks. *Econometrica*, 75(3):711–756.

Basu, P., Chatterjee, K., Hoshino, T., and Tamuz, O. (2018). Repeated coordination with private learning.

Bergemann, D. and Morris, S. (2016). Bayes correlated equilibrium and the comparison of information structures in games. *Theoretical Economics*, 11(2):487–522.

Bergin, J. and Lipman, B. L. (1996). Evolution with state-dependent mutations. *Econometrica*, 64(4):943–956.

Blume, L. E. (1993). The statistical mechanics of strategic interaction. *Games and Economic Behavior*, 5(3):387–424.

Caplin, A., Dean, M., and Leahy, J. (2017). Rationally inattentive behavior: Characterizing and generalizing Shannon entropy. *National Bureau of Economic Research*.

Carlsson, H. and van Damme, E. (1993). Global games and equilibrium selection. *Econometrica*, 61(5):989–1018.

Cover, T. M. and Thomas, J. A. (2006). *Elements of Information Theory*. Wiley-Interscience.

Dean, M. and Neligh, N. (2017). Experimental tests of rational inattention.

Denti, T. (2018). Unrestricted information acquisition.

Fosgerau, M., Melo, E., de Palma, A., and Shum, M. (2018). Discrete choice and rational inattention: A general equivalence result.

Goeree, J. K., Holt, C. A., and Palfrey, T. R. (2016). *Quantal Response Equilibrium: A Stochastic Theory of Games*. Princeton University Press.

Haile, P. A., Hortaçsu, A., and Kosenok, G. (2008). On the empirical content of qantal response equilibrium. *American Economic Review*, 98(1):180–200.

Hellwig, C. and Veldkamp, L. (2009). Knowing what others know: Coordination motives in information acquisition. *Review of Economic Studies*, 76(1):223–251.

Hertz, J., Krogh, A., and Palmer, R. G. (1991). *Introduction to the theory of neural computation.* CRC Press.

Hinton, G. E. and Sejnowski, T. J. (1986). Learning and relearning in Boltzmann machines. *Parallel Distributed Processing*, 1.

Hobson, A. (1969). A new theorem of information theory. *Journal of Statistical Physics*, 1(3):383–391.

Hoshino, T. (2018). Communication and coordination.

Kajii, A. and Morris, S. (1997). The robustness of equilibria to incomplete information. *Econometrica*, 65(6):1283–1309.

Kandori, M., Mailath, G. J., and Rob, R. (1993). Learning, mutation, and long run equilibria in games. *Econometrica*, 61(1):29–56.

Kirkpatrick, S., Gelatt, C. D., and Vecchi, M. P. (1983). Optimization by simulated annealing. *Science*, 220(4598):671–680.

Matějka, F. and McKay, A. (2015). Rational inattention to discrete choices: A new foundation for the multinomial logit model. *American Economic Review*, 105(1):272–98.

McFadden, D. L. (1984). Econometric analysis of qualitative response models. *Handbook of econometrics*, 2:1395–1457.

McKelvey, R. D. and Palfrey, T. R. (1995). Quantal response equilibria for normal form games. *Games and Economic Behavior*, 10(1):6–38.

Monderer, D. and Shapley, L. S. (1996). Potential games. *Games and Economic Behavior*, 14(1):124–143.

Morris, S. and Shin, H. S. (2003). Global games: Theory and applications. In *Advances in Economics and Econometrics: Theory and Applications, Eighth World Congress*, volume 1, pages 56–114.

Morris, S. and Yang, M. (2016). Coordination and continuous choice.

Myatt, D. P. and Wallace, C. (2011). Endogenous information acquisition in coordination games. *Review of Economic Studies*, 79(1):340–374.

Persico, N. (2000). Information acquisition in auctions. *Econometrica*, 68(1):135–148.

Rubinstein, A. (1989). The electronic mail game: Strategic behavior under “almost common knowledge”. *American Economic Review*, 79(3):385–391.

Simon, H. A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics*, 69(1):99–118.

Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics*, 50(3):665–690.

Sims, C. A. (2010). Rational inattention and monetary economics. In *Handbook of Monetary Economics*, volume 3, pages 155–181.

Steverson, K., Brandenburger, A., and Glimcher, P. W. (2017). Rational imprecision: Information-processing, neural, and choice-rule perspectives.

Train, K. E. (2009). *Discrete choice methods with simulation*. Cambridge University Press.

Weinstein, J. and Yildiz, M. (2007). A structure theorem for rationalizability with application to robust predictions of refinements. *Econometrica*, 75(2):365–400.

Yang, M. (2015). Coordination with flexible information acquisition. *Journal of Economic Theory*, 158:721–738.

Young, H. P. (1993). The evolution of conventions. *Econometrica*, 61(1):57–84.