

Robust Predictions in Games with Rational Inattention*

Tommaso Denti[†]
NYU Stern

Doron Ravid[‡]
University of Chicago

July 14, 2023

Abstract

We derive robust predictions in games involving flexible information acquisition, also known as *rational inattention* (Sims, 2003). These predictions remain accurate regardless of the exact specification of players' learning abilities. Compared to scenarios where information is predetermined, rational inattention reduces welfare and introduces additional constraints on behavior. We show these constraints generically do not bind; the two knowledge regimes are behaviorally indistinguishable in most environments. Yet, we demonstrate the welfare difference they generate is robust: optimal policy depends on whether one assumes information is given or acquired. We provide the necessary tools for policy analysis in this context.

1. Introduction

The economics discipline has long recognized that information matters for incentives. Many studies have also noted the reverse, namely, that incentives shape information. The theory of rational inattention, initiated by Sims (2003), is a case in point. This theory suggests that economic agents effectively use their limited ability to gather

*We thank Alex Bloedel, Benjamin Brooks, Modibo Camara, Marina Halac, Emir Kamenica, Ilia Krasikov, Elliot Lipnowski, Andrew McClellan, Stephen Morris, Luciano Pomatto, Daniel Rapoport, Phil Reny, and Joe Root for helpful discussions.

[†]td838@nyu.edu

[‡]dravid@uchicago.edu

and process information: faced with uncertainty, they flexibly acquire and use information in an optimal fashion, trading off costs and benefits. The rational inattention approach has been very successful, seeing a broad array of applications.¹

Despite their success, most rational inattention models come with a caveat: they require one to make non-trivial assumptions on the information acquisition environment, and these assumptions can be challenging to verify. One reason is that information costs and constraints are often affected by factors that are difficult to directly measure, such as time, effort, and cognitive resources. A few tests of rational inattention have been carried out in laboratory settings, and the results have been concerning. For example, Dean and Neligh (forthcoming) provide experimental evidence against the ubiquitous entropy-reduction cost (Matějka and McKay, 2015). Overall, these considerations raise the question of what predictions of rational inattention are robust to the exact specification of agents' learning abilities.

In this paper, we address these concerns by developing a framework for predictions that are valid across all information acquisition technologies that represent rational inattention. Using this framework, we identify the features that distinguish rational inattention from traditional models where information is given. We demonstrate that in most settings, there is no substantive difference in terms of behavioral outcomes, yet a robust distinction emerges when considering welfare consequences.

Our findings carry two key implications. Firstly, they indicate that studies that utilize rational inattention as a behavioral model need to consider and account for the subtle details of the subjects' attention limitations, either implicitly or explicitly; otherwise, the analysis reverts back to standard information economics. Secondly, our paper highlights a significant array of situations where mistakenly assuming information is given, when it is actually acquired, can lead to misleading welfare conclusions. We provide tools for identifying such situations and for conducting robust welfare analysis in economies with rationally inattentive agents.

Our model begins by considering a scenario in which strategic agents, referred to as players, are faced with the task of simultaneously selecting from a set of alternative actions. The utility of these actions is contingent upon an unknown state of nature. This scenario is commonly known as a *base game*. Our objective is to generate predictions regarding the players' behavior and welfare within this base game. Given the

¹See Maćkowiak, Matějka, and Wiederholt (2023) for a recent review.

inherent uncertainty surrounding the true payoff state, which directly influences the attractiveness of various actions, these predictions rely on two factors: the information that players are capable of acquiring and the associated costs of obtaining that information.

Standard rational inattention models obtain predictions by combining a base game with what we term an *information technology*, which specifies players' learning capabilities regarding the payoff state. In early studies, this technology assumes that information costs are proportional to entropy reduction, and that players can flexibly gather information, provided that it is independent of the other agents' information (e.g., Maćkowiak and Wiederholt, 2009; Yang, 2015). Recent papers allow for different cost functions and potentially correlated signals.²

Together, a base game and an information technology define an *information acquisition game*. This game begins with players covertly choosing what costly information to acquire. Subsequently, each player privately uses the acquired knowledge to take an action. Predictions are derived by solving for Nash equilibrium.

We depart from the standard approach, and do not pair the base game with a fixed information technology. Instead, we simultaneously find all the predictions one can obtain with a technology that is consistent with rational inattention. Specifically, we consider information technologies that satisfy three properties. First, no information comes at zero cost. Second, (strictly) more informative signals (in the sense of Blackwell, 1951, 1953) cost (strictly) more. And third, information choice is flexible. These properties are satisfied by virtually all applications of rational inattention. We are agnostic whether signals are independent or correlated across players.

Theorem 1 characterizes all behavioral and welfare implications of rational inattention in a given base game. The behavior of the players is summarized by the joint distribution of the actions they take and the payoff state drawn by nature. We call this distribution the *outcome* of the game. The welfare aspect is determined by the *value* that each player obtains from the game, which corresponds to their payoff net of any costs they incurred for acquiring information.

We show rational inattention can generate any outcome that satisfies two constraints: *obedience* and *separation*. Each constraint comes from a different optimality condition: obedience represents optimal behavior given a fixed signal structure,

²See, e.g., Morris and Yang (2022); Ravid, Roesler, and Szentes (2022); Denti (forthcoming); Hebert and La'O (forthcoming)

whereas separation stems from players' incentives to collect information. We follow the literature (Bergemann and Morris, 2013, 2016), and use the term *Bayes correlated equilibrium* (BCE) to refer to obedient outcomes. Thus, Theorem 1 says an outcome can be generated by rational inattention if and only if it is a separated BCE (sBCE).

Theorem 1 also establishes precise bounds on the net payoff that each player can achieve in a given equilibrium outcome. The upper bound corresponds to the scenario in which the player acquires information at essentially no cost, and we refer to this bound as the outcome's *gross value*. The lower bound represents the situation where the player gathers information that is so expensive that they would have been equally satisfied without acquiring any information at all. We name this lower bound the outcome's *uninformed value*. Theorem 1 demonstrates that, given a separated BCE, each player has the potential to achieve any payoff that surpasses their uninformed value but remains below their gross value.

Theorem 1 allows us to study whether rationally inattentive agents exhibit special behavioral properties, and the impact information costs have on welfare. For this purpose, we compare the predictions made by rational inattention to those generated by models where information is an exogenous variable. Under exogenous information, players do not choose what they know; instead, they are endowed with a fixed signal structure. Bergemann and Morris (2013, 2016) show that an outcome can be induced by some signal structure if and only if it is a BCE (i.e., it is obedient). Welfare is given by the outcome's gross value, since no learning costs are paid when information is predetermined.

Whereas the separation constraint can dramatically shrink the BCE set, such situations turn out to be rare. Specifically, we show in Theorem 2 that, for generic preferences, one can approximate every obedient outcome with outcomes that are both obedient and separated. To put it simply, by introducing small perturbations to the payoffs of a base game, we can transform the game into a scenario where nearly all BCEs become separated, even if the original set of non-separated BCEs was extensive. Thus, unless one commits to a highly structured economic setting or makes specialized assumptions on agents' information technologies, rational inattention is behaviorally indistinguishable from exogenous information.

Our analysis also reveals that rational inattention and exogenous information, de-

spite yielding the same behavioral predictions, can have distinct welfare implications. Intuitively, when information is exogenous, the value of information must go to the players. By contrast, in the worst case under rational inattention, players must pay this value to purchase their information. Theorem 3 explores the impact of this intuition on the set of attainable payoffs. The theorem demonstrates that, for generic preferences, rational inattention expands the set of achievable payoffs when compared to exogenous information. This expansion is not trivial, as it often results in a strict inclusion of additional payoffs within the set.

The welfare difference between rational inattention and exogenous information can have concrete repercussions for public policy. We give a proof of concept in Section 6, where we show how to use our results to conduct robust welfare in economies where agents are rationally inattentive. Specifically, we consider a utilitarian social planner who understands the economic environment (i.e., the base game), but does not know the source of players' information (i.e., the information technology). The planner takes a worst-case approach, as is common in robust mechanism design.³ First, we characterize the set of binary-action symmetric games in which this planner's welfare evaluations depend on whether she assumes players' information is given or acquired. We then apply this characterization to a regime change game, in which several investors choose whether or not to attack a distressed financial institution. We show that the planner may choose to bolster the institution's fundamentals if she thinks information is exogenous, but she should never do so if she believes the investors are rationally inattentive.

Since many settings of interest are non-generic (e.g., auctions), we conclude the paper by studying environments for which the separation constraint binds. We provide a tight characterization, and show that separation is an all-or-nothing refinement of BCE: the sBCE set is either dense or nowhere dense in the set of BCEs. Thus, whenever the separation constraint binds, it dramatically reduces the set of attainable outcomes.

Related literature. Our paper is related to several strands of literature. First, we contribute to the study of robustness in game-theoretic predictions. Especially pertinent is the work that uses the BCE solution concept to obtain robust predictions in games with incomplete information (e.g., Bergemann, Brooks, and Morris, 2017;

³See Carroll (2019) for a recent survey of this literature.

Brooks and Du, 2021). In particular, we show the robustness criterion of such studies extend to include endogenously determined information à la rational inattention, as long as the focus is only on the moments of the action-state distribution (e.g., the seller’s revenue in an auction), and the separation constraint is not binding. Even when separation binds, our genericity result means such constraint should be taken seriously only if the analyst is confident about the non-generic features of the economic environment. For studies that focus on welfare (e.g., Bergemann, Brooks, and Morris, 2015), our framework provides a road-map for understanding whether their conclusions continue to hold once one accounts for the cost of information.

To best of our knowledge, our paper is the first to conduct a robust welfare analysis under rational inattention. Several studies have considered the effects of costly information acquisition on public policy and efficiency (e.g., Colombo, Femminis, and Pavan, 2014; Ravid, Roesler, and Szentes, 2022; Angeletos and Sastry, 2023; Hebert and La’O, forthcoming). A distinctive feature of our work is that we evaluate welfare in a manner that is independent of the specific details of agents’ learning abilities. The tools we develop are portable and can help policy makers when estimating information costs is challenging or infeasible.

More broadly, we contribute to the literature on rational inattention in games (e.g., Yang, 2015; Ravid, 2020; Morris and Yang, 2022; Denti, forthcoming). Within this literature, the closely related work of Denti (2021) is the first to consider the question of robustness. Denti (2021) studies a two-player signaling game where the receiver needs to pay a cost to monitor the sender’s action. He shows that, when costs are strictly monotone, off-path beliefs play no role in equilibrium. He then characterizes the behavioral predictions that are consistent with some strictly monotone cost function. As part of this characterization, he obtains what is essentially the single-player version of our Theorem 1. Our Theorem 1 expands on Denti’s (2021) result by allowing for multiple players, and by studying welfare (Denti, 2021, does not discuss payoffs). In addition, our analysis of generic games and the comparison to exogenous information have no parallel in Denti (2021).

One can interpret some of our results as providing a test for rational inattention in a strategic settings: simply check whether the observed outcome satisfies obedience and separation. Moreover, our genericity result suggests that testing separation

requires a highly structured and controlled environment. Thus, our paper can be seen as adding to the growing literature on the testable implications of rational inattention (e.g., Caplin and Dean, 2015; Caplin, Dean, and Leahy, 2022; Denti, 2022; Lipnowski and Ravid, 2022). Within this literature, the most closely related paper is Caplin and Dean (2015).⁴ They develop a test for whether a single agent’s choices in multiple menus are consistent with costly information acquisition, but do not require costs to be strictly monotone in information. Their characterization includes obedience, as well as another condition called *no improving attention cycles* (NIAC), which restricts the agent’s behavior across decision problems. Since we consider a fixed base game, the NIAC restriction does not apply in our setting. By contrast, Caplin and Dean’s (2015) characterization does not require separation, because they allow for costs that are not strictly monotone. Hence, their characterization differs from ours in that it only considers a single agent, does not require costs to be strictly monotone, and accounts for multiple decision problems.

Our work also expands on the uses of correlated equilibrium and its cousins for spanning the set of predictions attainable across various ways of extending a base game (e.g., Aumann, 1974; Myerson, 1982; Forges, 1993; Bergemann and Morris, 2016; Doval and Ely, 2020). An early and closely related paper within this literature is Lipman and Srivastava (1990). They consider base games without payoff uncertainty, and ask which correlated equilibria can be attained by extending the game to allow players to acquire costly information about a common payoff-irrelevant state. They maintain two assumptions on players’ information technology: costs are (ordinally) symmetric across players, and players must use partitional information. They show an obedient outcome can be generated from an information technology satisfying their assumptions if and only if it satisfies a cyclical monotonicity condition across players that is similar to Caplin and Dean’s (2015) NIAC. Our work differs from theirs in that we allow for payoff uncertainty, asymmetric costs, and non-partitional information, and that we require costs to be strictly increasing in informativeness. In addition, Lipman and Srivastava (1990) have no analog of our payoff bounds.

⁴Another related paper is de Clippel and Rozen (2022). They consider a two-staged game where the first stage player chooses how much to obfuscate the state, and the second stage player chooses what to learn about that state whenever it is obfuscated. In their setting, the second player essentially faces a single agent decision problem. Using this fact, de Clippel and Rozen (2022) apply Caplin and Dean’s (2015) results to obtain testable prediction that are valid across all cost functions. They then test these predictions in a lab experiment.

2. Setup

We consider a game involving a finite number of players where the players face uncertain payoffs, call it the *base game*. Before playing this base game, the players can covertly acquire costly information to reduce the uncertainty they face. We call *information technology* the description of learning resources. Together, a base game and an information technology define an *information acquisition game*, our main object of study. We focus on two equilibrium predictions: the induced *outcome* of the base game, and the net *value* each player obtains from information acquisition. In this section, we precisely define all these objects and provide the associated notation.

Base Game. Let I be a finite set of players, with typical element i . Each player i has to choose an action a_i from a finite set A_i . As usual, we define $A_{-i} = \prod_{j \neq i} A_j$ and $A = A_i \times A_{-i}$. Accordingly, we use $a_{-i} = (a_j)_{j \neq i}$ to denote the action profile of all players other than i , and $a = (a_i, a_{-i})$ to denote the entire action profile. Throughout the paper we adopt the same notational conventions for all Cartesian products indexed by $I \setminus \{i\}$ and I .

Players are expected utility maximizers who care about each other's actions as well as an exogenous variable θ , which is drawn from a finite set Θ according to a full-support probability measure $\pi \in \Delta(\Theta)$. We denote by $u_i : A \times \Theta \rightarrow \mathbb{R}$ player i 's von Neumann-Morgenstern utility function.

Embracing an established terminology (see, e.g., Bergemann and Morris, 2016), we call a tuple $\mathcal{G} = (I, \Theta, \pi, (A_i, u_i)_{i \in I})$ a **base game**.

Information Technologies. Before taking an action in the base game, each player has the opportunity to acquire information about θ as well as other exogenous quantities of potential interests (e.g., sunspots, noisy public information). We succinctly represent them by a single variable z taking values in a finite set Z ; a Markov kernel $\zeta : \Theta \rightarrow \Delta(Z)$ details the conditional distribution of z given θ .

Following Blackwell (1951), we model the acquisition of information using experiments. An experiment for player i is a Markov kernel $\xi_i : Z \times \Theta \rightarrow \Delta(X_i)$, where X_i is a finite space of signal realizations privately observable by player i . The functions ξ_i details how the distribution of i 's signal x_i depends on θ and z .

By construction, the players' signals are conditionally independent given θ and

z . However, they may be correlated given θ only. Thus, our framework incorporates a form of correlated information acquisition, as in, among others, Hellwig and Veldkamp (2009), Myatt and Wallace (2012), Hebert and La'O (forthcoming), and Denti (forthcoming). Indeed, one can simply view z as a modeling device for situations in which players have access to information sources with correlated noise (e.g., newspapers with similar slants, consultants from the same firm).

The acquisition of information faces two kinds of frictions. First, each player is constrained in the kind of experiments she can use: player i can only choose experiments that lie in a given set \mathcal{E}_i . Second, experiments come at a cost, where $C_i : \mathcal{E}_i \rightarrow \mathbb{R}_+$ denotes i 's cost function. As a normalization, we assume the existence of an experiment that costs zero; that is, $C_i(\xi_i) = 0$ for some $\xi_i \in \mathcal{E}_i$.

We call a tuple $\mathcal{T} = (Z, \zeta, (X_i, \mathcal{E}_i, C_i)_{i \in I})$ an **information technology**.

Information Acquisition Games. Together, a base game \mathcal{G} and an information technology \mathcal{T} define an **information acquisition game**. The game begins with the realization of the state of nature $\omega = (z, \theta)$. Without observing the state, the players simultaneously and covertly choose experiments, and pay their costs. Then, each player privately observes the outcome of their own experiment and takes an action. We use $\sigma_i : X_i \rightarrow \Delta(A_i)$ to denote player i 's action plan in this game, and let Σ_i be the set of i 's action plans.

The solution concept we adopt is Nash equilibrium. A strategy for player i consists of an experiment $\xi_i \in \mathcal{E}_i$ and an action plan $\sigma_i \in \Sigma$. A strategy profile $(\xi_i^*, \sigma_i^*)_{i \in I}$ is an equilibrium if for all players i , (ξ_i^*, σ_i^*) maximizes

$$\left[\sum_{a, x, z, \theta} u_i(a, \theta) \sigma_i(a_i | x_i) \xi_i(x_i | z, \theta) \prod_{j \neq i} \sigma_j^*(a_j | x_j) \xi_j^*(x_j | z, \theta) \zeta(z | \theta) \pi(\theta) \right] - C_i(\xi_i).$$

over all $\xi_i \in \mathcal{E}_i$ and $\sigma_i \in \Sigma_i$. The objective function consists of two terms: the benefit of information (in square brackets) and the cost of information. As common in applications, benefit and cost are additively separable.

Equilibrium Predictions. We summarize the equilibria of information acquisition games using two statistics: the induced outcome of the base game, and the net value each player obtains from information acquisition.

The induced **outcome** of the base game is the joint distribution $p \in \Delta(A \times \Theta)$ of

the players' actions and the payoff-relevant state. Note that the marginal distribution of θ must be the prior π ; we denote by $\Delta_\pi(A \times \Theta)$ the set of probability measures over $A \times \Theta$ whose marginal on Θ is π .

The net **value** player i obtains from information acquisition is simply i 's expected payoff, counting both benefits and costs of information. We denote by $v = (v_i)_{i \in I} \in \mathbb{R}^I$ the vector of such values. Since v_i includes player i 's information acquisition costs, it may differ from the expectation of u_i under p .

3. Rational Inattention: A Definition

Our aim is to study the behavioral outcomes and welfare values that can be generated in equilibrium of information acquisition games as we fix the base game and range over all information technologies that represent rational inattention. Consistently with the literature, we interpret rational inattention as information technologies where the set of feasible experiments is flexible, and where it is costly to acquire more information.

To give a precise definition of “more information,” we build on the classic ranking of experiments due to Blackwell (1951, 1953). Given a pair of experiments ξ_i and ξ'_i , we say ξ_i **Blackwell dominates** ξ'_i (denoted $\xi_i \succsim \xi'_i$) if there exists a Markov kernel $g : X_i \rightarrow \Delta(X_i)$ such that for every $x_i \in X_i$, $\theta \in \Theta$, and $z \in Z$ with $\zeta(z|\theta) > 0$,

$$\xi'_i(x_i|z, \theta) = \sum_{x'_i \in X_i} g(x_i|x'_i) \xi(x'_i|z, \theta).$$

Thus, ξ_i Blackwell dominates ξ'_i if one can generate ξ'_i by garbling the output of ξ_i . As shown by Blackwell, $\xi_i \succsim \xi'_i$ if and only if player i is better off observing the output of ξ_i rather than the output of ξ'_i (holding fixed other players' behavior). In this sense, ξ_i is more informative than ξ'_i . We write $\xi_i \succ \xi'_i$ whenever $\xi_i \succsim \xi'_i$ and $\xi'_i \not\succeq \xi_i$.

We say a set of feasible experiments \mathcal{E}_i is **flexible** if, whenever a given amount of information is feasible, a lower amount of information is feasible as well: whenever $\xi_i \in \mathcal{E}_i$ and $\xi_i \succsim \xi'_i$, then $\xi'_i \in \mathcal{E}_i$. We say a cost function C_i is **monotone** if less informative experiments are cheaper to acquire: for all $\xi_i, \xi'_i \in \mathcal{E}_i$ such that $\xi_i \succsim \xi'_i$ (resp., $\xi_i \succ \xi'_i$), we have $C_i(\xi_i) \geq C_i(\xi'_i)$ (resp., $C_i(\xi_i) > C_i(\xi'_i)$). We say a technology \mathcal{T} represents **rational inattention** if for every player, the set of feasible experiments is flexible and the cost function is monotone.

Many authors regard flexibility as the key difference between rational inattention

and traditional information-acquisition models (see, e.g., Maćkowiak, Matějka, and Wiederholt, 2023, Section 2). Applications of rational inattention often assume all experiments are feasible: in the language of this paper, $\mathcal{E}_i = \Delta(X_i)^{Z \times \Theta}$. Our results are unchanged if we adopt this more extreme definition of flexibility. The reason is that there are no observable differences between experiments that are unfeasible and experiments that are excessively costly.

By pairing flexibility with monotonicity, we postulate that players can save on costs by only acquiring the information they actually use in making decisions. Matějka and McKay (2015) provide a standard example of monotone cost function: $C_i(\xi_i)$ equals the expected reduction in uncertainty about θ and z , as measured by Shannon entropy. More broadly, one could substitute Shannon’s entropy with any other strictly concave measure of uncertainty (Caplin, Dean, and Leahy, 2022). One could also consider increasing transformations of these costs (Denti, 2022; Zhong, 2022), or any differentiable cost function whose derivative is strictly convex (Lipnowski and Ravid, 2022).⁵

4. Robust Predictions

This section presents our first main result: a robust characterization of the outcomes and values that can be generated via rational inattention. To provide a benchmark, we first review the case of exogenous information.

The class of information acquisition games includes situations in which the players’ information is predetermined. One can obtain them by considering technologies in which every player has only one feasible experiment (whose cost is zero by our normalization). Each such technology can be identified with a tuple $\mathcal{S} = (Z, \zeta, (X_i, \xi_i)_{i \in I})$. Of course, an information acquisition game $(\mathcal{G}, \mathcal{S})$ is just a game of incomplete information à la Harsanyi, where \mathcal{S} is the predetermined information structure.

Among other results, Bergemann and Morris (2016) characterize the equilibrium outcomes that can arise in a game of incomplete information $(\mathcal{G}, \mathcal{S})$ as we fix the base game \mathcal{G} and range over all information structures \mathcal{S} : they call them Bayes correlated equilibria. A **Bayes correlated equilibrium** (BCE) of a base game \mathcal{G} is an outcome

⁵In each information acquisition game, prior beliefs are exogenously determined by π and ζ . Thus, the issue of experiment-based vs. posterior-based information costs that sometimes arise in applications of rational inattention (see, e.g., Ravid, 2020; Denti, Marinacci, and Rustichini, 2022) is irrelevant here.

$p \in \Delta_\pi(A \times \Theta)$ such that for all $i \in I$ and $a_i, b_i \in A_i$,

$$\sum_{a_{-i}, \theta} (u_i(a_i, a_{-i}, \theta) - u_i(b_i, a_{-i}, \theta)) p(a_i, a_{-i}, \theta) \geq 0. \quad (1)$$

Following standard terminology, we name (1) the **obedience constraint**.⁶ The standard way of viewing this constraint is through the lens of a mediator who generates p by observing the state and stochastically sending an action recommendation to each player. The players are willing to follow these recommendations if and only if the obedience constraint is satisfied. Even if our framework does not include an actual mediator, throughout the paper we find it helpful to use the mediator's metaphor as expositional device.

Calculating players' values under exogenous information is straightforward. If an outcome p arises under exogenous information in a base game \mathcal{G} , player i 's expected payoff is

$$\bar{v}_i(p) = \sum_{a, \theta} u_i(a, \theta) p(a, \theta).$$

We call $\bar{v}_i(p)$ the **gross value** for player i , since it ignores information costs. Let $\bar{v}(p) = (\bar{v}_i(p))_{i \in I}$ be the vector of gross values.

What changes when information is endogenous? Our analysis highlights two main differences between rational inattention and exogenous information.

The first difference is that outcomes generated by costly information acquisition must satisfy an additional constraint, which we name the *separation constraint*. To present this constraint, we require a few definitions. Given an outcome $p \in \Delta_\pi(A \times \Theta)$, a player $i \in I$, and an action $a_i \in A_i$, let $p(a_i) = \sum_{a_{-i}, \theta} p(a_i, a_{-i}, \theta)$ be the probability of player i taking action a_i under p , and let

$$\text{supp}_i(p) = \{a_i \in A_i : p(a_i) > 0\}$$

be the set of i 's actions that have positive probability. For each $a_i \in \text{supp}_i(p)$, let $p_{a_i} \in \Delta(A_{-i} \times \Theta)$ be the conditional distribution of the actions of the players other

⁶Our definition of BCE corresponds to the specialization of Bergemann and Morris's (2016) definition to the case in which players' original type spaces are degenerate. Forges (1993) refers to such outcomes as *universal Bayesian solutions*.

than i and the payoff-relevant state: for all $a_{-i} \in A_{-i}$ and $\theta \in \Theta$,

$$p_{a_i}(a_{-i}, \theta) = \frac{p(a_i, a_{-i}, \theta)}{p(a_i)}.$$

We also denote the set of player i 's best responses to p_{a_i} via

$$BR(p_{a_i}) = \operatorname{argmax}_{b_i \in A_i} \sum_{a_{-i}, \theta} u_i(b_i, a_{-i}, \theta) p_{a_i}(a_{-i}, \theta).$$

Finally, we say an outcome $p \in \Delta_\pi(A \times \Theta)$ satisfies the **separation constraint** if for all $i \in I$ and $a_i, b_i \in \operatorname{supp}_i(p)$,

$$p_{a_i} \neq p_{b_i} \quad \text{implies} \quad BR(p_{a_i}) \cap BR(p_{b_i}) = \emptyset. \quad (2)$$

In other terms, an outcome satisfies the separation constraint if distinct beliefs have separate best responses. We refer to a BCE that satisfies the separation constraint as a **separated BCE** (sBCE).

The second difference between rational inattention and exogenous information is in the set of payoffs players obtain from a given outcome. When information is given, each player i 's value is completely determined by the outcome p . By contrast, under rational inattention the same outcome can arise under multiple cost functions, and so is consistent with a set of values. As we demonstrate next, this set of values is convex, with an upper bound given by the gross value $\bar{v}_i(p)$. The lower bound is given by (what we call) the outcome's *uninformed value*, which is i 's maximal value if she has to act before receiving the mediator's recommendation. Formally, the **uninformed value** of an outcome p in a base game \mathcal{G} for player i is the quantity

$$\underline{v}_i(p) = \max_{b_i \in A_i} \sum_{a, \theta} u_i(b_i, a_{-i}, \theta) p(a, \theta).$$

Let $\underline{v}(p) = (\underline{v}_i(p))_{i \in I}$ be the vector of uninformed values.⁷

The next result summarizes our characterization of rational inattention in games:

⁷The uninformed value is related to the notion of *coarse correlated equilibrium* (see, e.g., Roughgarden, 2016, Definition 13.5). Coarse correlated equilibrium relaxes the obedience constraint, requiring instead that no player can benefit from deviating to a single action from all of the mediator's recommendation; that is, $\bar{v}_i(p) \geq \underline{v}_i(p)$ for every i .

Theorem 1. Fix a base game \mathcal{G} . A rational-inattention technology \mathcal{T} exists that induces the outcome-value pair (p, v) in an equilibrium of $(\mathcal{G}, \mathcal{T})$ if and only if

- (i) the outcome p is a separated BCE, and
- (ii) for every $i \in I$, $v_i = \underline{v}_i(p) = \bar{v}_i(p)$ or $v_i \in [\underline{v}_i(p), \bar{v}_i(p)]$.

Thus, an outcome-value pair is consistent with rational inattention if and only if two conditions hold. First, the outcome satisfies both the obedience constraint and the separation constraint. And second, each player's value is weakly above her uninformed value, but strictly below her gross value.

We now explain why the theorem's conditions are necessary. That every outcome generated by rational inattention must be obedient follows from the necessity of the obedience constraint under exogenous information. The reason is that any outcome one can attain when players choose their information must also be attainable if players were exogenously endowed with their chosen experiment. To understand why the separation constraint is necessary, consider a player i who takes with positive probability a pair of actions a_i and b_i such that $p_{a_i} \neq p_{b_i}$. As in BCE, we can interpret a_i and b_i as signals. With rational inattention, informative signals are costly. To save on information costs, the player could substitute a_i and b_i with a single action recommendation c_i . For this substitution not to be profitable, it must be that either $c_i \notin BR(p_{a_i})$ or $c_i \notin BR(p_{b_i})$. Since the choice of c_i is arbitrary, it must be that $BR(p_{a_i}) \cap BR(p_{b_i}) = \emptyset$.

Next, we explain the necessity of Theorem 1's payoff bounds. To understand the lower bound, suppose we have an information technology and an equilibrium that induces an outcome p . By assumption, player i always has the option of remaining uninformed at no cost. Therefore, i 's optimal payoff must be higher than $\underline{v}_i(p)$. As for the upper bound, recall that information costs are non-negative. Consequently, player i 's payoff from an equilibrium that induces p must be below her payoffs if her information was free; that is, $v_i \leq \bar{v}_i(p)$. Moreover, her information must come at a strictly positive cost whenever $\underline{v}_i(p) \neq \bar{v}_i(p)$: to generate p in this case player i must acquire *some* information. Hence, player i can actually attain her gross value from an outcome only if that value coincides with the uninformed value.

We now briefly review our proof that conditions (i) and (ii) of Theorem 1 are sufficient. The proof is constructive, and is based on a result by Denti (2021). Denti (2021) studies a two-player signaling game where the receiver pays a cost to monitor

the sender’s action. Among other results, Denti (2021) shows that, when the sender takes every action with positive probability, any receiver behavior satisfying (what we call in this paper) obedience and separation can be justified via rational inattention. To prove the “if” direction of Theorem 1, we extend Denti’s (2021) single-agent construction so that it applies to settings in which multiple agents simultaneously acquire information. In addition, we show that any payoff between the uninformed and the gross value can be generated in equilibrium—Denti (2021) focuses on equilibrium outcomes and does not discuss achievable payoffs.

Theorem 1 shows that any BCE that arises from rational inattention must satisfy the separation constraint. Next, we record that the separation constraint never eliminates all of a game’s Bayes correlated equilibria; that is, the set of separated BCEs is non-empty.

Corollary 1. *Every base game \mathcal{G} admits a separated BCE.*

The proof is straightforward (details omitted): A technology where all feasible experiments are free and uninformative is flexible and monotone. The corresponding game with information acquisition admits an equilibrium by standard arguments (information is de facto exogenous). It follows from Theorem 1 that the outcome of such an equilibrium is a separated BCE.⁸

Next, to help the reader familiarize themselves with the concepts we introduced in this section, we show Theorem 1’s implications in a simple coordination game.

Example 1. There are two players, $I = \{1, 2\}$, each of which chooses a binary action. The players get a payoff of 1 if their actions match, -1 otherwise. Specifically, $A_i = \{-1, 1\}$ and $u_i(a_1, a_2) = a_1 a_2$ for each $i \in I$. The payoff state is degenerate, hence omitted.

We first give an example of a non-separated BCE. Let $p \in \Delta(A)$ be the outcome such that assigns $(a_1, a_2) = (1, 1)$ a probability of $1/2$; each other pair of actions occurs with probability $1/6$. It is easy to see p is obedient: conditional on any action recommendation, the probability the other player takes the same action is at least $1/2$. However, this BCE is not separated: $p_{a_1=1}(a_2 = 1) = 3/4 \neq 1/2 = p_{a_1=-1}(a_2 = 1)$, but $a_1 = 1$ is a best response to both action recommendations. Thus,

⁸Another immediate consequence of this argument is that in the special case in which Θ is a singleton, all (pure or mixed) Nash equilibria of the base game are separated BCEs.

the recommendations $a_1 = 1$ and $a_1 = -1$ have a common best response, but lead to different beliefs.

For an example of a separated BCE, take $q \in \Delta(A)$ to be the uniform distribution over the two action profiles $(a_1, a_2) = (1, 1)$ and $(a_1, a_2) = (-1, -1)$. Because the players are perfectly coordinating their actions, taking the action that the mediator recommends is a strict best response. It follows q is a separated BCE.

We now demonstrate the difference between gross and uninformed value. Consider the sBCE q . Since the players perfectly coordinate their actions, the gross value from this BCE is $\bar{v}_1(q) = \bar{v}_2(q) = 1$. The uninformed value, however, is lower: a player who does not see the mediator's recommendation gets a payoff of 0 no matter what action she takes, given that the co-player takes both actions with equal probability. Therefore, $\underline{v}_1(q) = \underline{v}_2(q) = 0$.

We conclude the section with a brief discussion of what happens if we allow for non-flexible and non-monotone technologies. To accommodate such technologies, we need to adjust Theorem 1's statement in two ways. First, the separation constraint is no longer necessary. And second, players can attain their gross value even when it differs from their uninformed value. We refer the reader to Online Appendix D for the precise details.

5. Generic Environments

Theorem 1 shows rational inattention differs from exogenous information in two ways. First, rational inattention reduces the set of equilibrium outcomes: an additional separation constraint must be satisfied. Second, rational inattention expands the set of achievable payoffs for any given equilibrium outcome: every payoff between the uninformed value and the gross value can be obtained. In this section, we prove that only the second difference is meaningful in generic environments.

We adopt the following notion of genericity: We fix a finite set of players I , a finite set of payoff states Θ , a full-support prior $\pi \in \Delta(\Theta)$, and a finite set of actions A_i for each player $i \in I$. To obtain a base game, it remains to specify a profile of utility function $u = (u_i)_{i \in I}$. We identify u with an element of the Euclidean space $\mathbb{R}^{I \times A \times \Theta}$, and say a statement is true for **generic** u if the closure of the subset in $\mathbb{R}^{I \times A \times \Theta}$ for which it is false has Lebesgue measure zero. We denote by $BCE(u)$ and $sBCE(u)$ the sets of BCEs and separated BCEs for the base game corresponding to u .

As long as the sets Θ , I , and A are non-trivial, one can find many examples of $u \in \mathbb{R}^{I \times A \times \Theta}$ for which the separation constraint imposes sharp restrictions on behavior, so that most BCE are not separated—we discuss a few simple cases in Online Appendix E. However, the next theorem shows that all such examples are non-generic.

Theorem 2. *For generic u , the set $sBCE(u)$ is dense in the set $BCE(u)$.*

Thus, the environments in which rational inattention predict different outcomes than incomplete information are knife edge.⁹ An important caveat to this result is the notion of genericity we use: it is the most common for the static games we study in this paper, but also the most permissive. For example, according to this notion of genericity, many important economic applications—such as auctions—are non-generic. The notion of genericity is also not appropriate if the base game is not actually static but represents the strategic form of a primitive dynamic game. We discuss sBCE in non-generic environments in Section 7.

One might be tempted to think that Theorem 2 follows from indifferences being fragile, so to speak. This intuition works, but only for the single agent case. When I has one element, the set of *strict* BCEs—i.e., the set of outcomes p such that $BR(p_{a_i}) = \{a_i\}$ for every player i and $a_i \in \text{supp}_i(p)$ —is generically dense in the BCE set. Theorem 2 then follows from noting that every strict BCE is separated.¹⁰

The situation radically changes when there are at least two players. The reason is that, with two or more players, indifferences emerge in equilibrium in generic fashion: when I has more than one element, there exists an open set of games where *no* BCE is strict. For example, consider games in a neighborhood of Matching Pennies: they have a unique correlated equilibrium (the fully-mixed Nash equilibrium) in which both players are indifferent between both actions. In sum, to prove Theorem 2 beyond the single-agent case, one cannot hope to show that indifferences are somewhat fragile.

Our proof of Theorem 2 combines two independent lemmas. The first lemma shows that for any BCE p and utility profile u , there exists a perturbation of u that makes p separated. In other words, every BCE of a given game is the limit of separated BCEs of nearby games.

⁹It is easy to construct generic examples where the set of separated BCEs is not close. In particular, the statement “ $sBCE(u) = BCE(u)$ for generic u ” is false. It is also not true that “ $\text{cl}(sBCE(u)) = BCE(u)$ for all u .”

¹⁰See Online Appendix F for a detailed argument.

The second lemma shows that for generic games, any BCE that is a limit of separated BCE in nearby games is also a separated BCE when the game is held fixed. Specifically, the second lemma shows that the correspondence $u \mapsto \text{cl}(sBCE(u))$, which takes utilities to the closure of the sBCE set, generically is upper hemicontinuous (in fact, continuous). To prove it, we employ ideas from Blume and Zame (1994), who study the algebraic geometry of Nash, perfect, and sequential equilibria. In particular, we use the Tarsky-Siedenbergs Theorem to show that $u \mapsto \text{cl}(sBCE(u))$ has semi-algebraic graph, and so must be continuous for all utility profiles outside a closed low-dimensional manifold.

Theorem 2 implies that, in generic games, rational inattention and exogenous information are outcome equivalent. Next, we show this equivalence does not extend to players' welfare. In other words, even though the two knowledge regimes generically yield the same behavioral predictions, they can have distinct welfare implications.

For $u \in \mathbb{R}^{I \times A \times \Theta}$, let $V_R(u)$ be the closure of the set of attainable value vectors under rational inattention. By Theorem 1, these are the value vectors that lie between the uninformed value and the gross value of some limit of separated BCEs:¹¹

$$V_R(u) = \{v \in [\underline{v}(p, u), \bar{v}(p, u)] : p \in \text{cl}(sBCE(u))\}, \quad (3)$$

where $\underline{v}(p, u)$ and $\bar{v}(p, u)$ make explicit the dependence of uninformed and gross values on the players' utility functions.

We also denote by $V_I(u)$ the set of attainable value vectors under exogenous information. By Bergemann and Morris (2016), this is the set of gross value vectors attainable in some BCE,

$$V_I(u) = \{\bar{v}(p, u) : p \in BCE(u)\}. \quad (4)$$

For an arbitrary u , there is no simple relationship between $V_R(u)$ and $V_I(u)$: $sBCE(u)$ is a subset of $BCE(u)$, but $\bar{v}(p, u)$ is an element of $[\underline{v}(p, u), \bar{v}(p, u)]$, so one cannot easily conclude that $V_R(u)$ contains or is contained by $V_I(u)$, or neither.

In generic environments, the comparison is simpler:

¹¹For two vectors of real numbers $\lambda, \gamma \in \mathbb{R}^n$ such that $\lambda \leq \gamma$, we use the box, or ordered interval notation $[\lambda, \gamma] = \{\eta \in \mathbb{R}^n : \lambda \leq \eta \leq \gamma\}$ to denote the set of all vectors between λ and γ .

Theorem 3. *For generic u , $V_I(u) \subseteq V_R(u)$. In addition, if $|I| \geq 2$, $|\Theta| \geq 2$, and $|A_i| \geq 2$ for at least two distinct players i , then the set of u for which $V_I(u) \subset V_R(u)$ has non-empty interior. If instead $|I| = 1$, then $V_I(u) = V_R(u)$ for generic u .*

Thus, in generic environments, rational inattention expands the set of achievable payoffs, and it does so in a non-trivial way (i.e., with strict inclusion) for a class of environments of positive measure. As we demonstrate in Section 6, the difference between $V_I(u)$ and $V_R(u)$ has substantial implications for welfare analysis: if a utilitarian social planner mistakenly assumes that information is given rather than acquired, she may end up choosing a sub-optimal policy, even in situations in which the two knowledge regimes are outcome equivalent. Our results give the planner the tools to avoid such mistakes.

Theorem 3 distinguishes between single-agent and many-player settings. In generic single-agent settings, rational inattention and exogenous information are not only outcome equivalent (as Theorem 2 states), but also welfare equivalent: $BCE(u) = sBCE(u)$ and $V_I(u) = V_R(u)$ for generic u . In generic many-player environments, rational inattention and exogenous information are only outcome equivalent: $BCE(u) = sBCE(u)$ for generic u , but $V_I(u) \subset V_R(u)$ for a set of utility profiles u with non-empty interior.

We now clarify the difference between the case of a single agent and the case of multiple players. Suppose first there is only one agent, and let p^0 be a BCE in which the agent receives no information. Since the agent's payoffs depend only on the correlation between her action and the state, and because the state's marginal distribution is constant across all BCEs, the agent's uninformed value is also constant across all BCEs, and equals her expected utility under no information; that is, every BCE p has $\underline{v}(p, u) = \bar{v}(p^0, u)$. Moreover, since more information always helps the agent, the BCE p^0 minimizes the agent's gross value across all BCEs. In other words, the minimal value of $V_R(u)$ and $V_I(u)$ is the same, and equals $\bar{v}(p^0, u)$. The single-agent result then follows from noting that, for generic u , $V_R(u)$ and $V_I(u)$ are convex sets that have the same maximum value (i.e., complete information).

Suppose now there are multiple players. In this case, each player cares not only about her action and the state, but also about the actions of others. Consequently, it is possible that to minimize one player's uninformed value, a BCE must change

the distribution of another player’s actions. If such a change requires giving players valuable information, one can obtain a gap between the minimal uninformed value and the lowest gross value for a given player. The proof of Theorem 3 constructs a game with this property, and shows that, for this game, this property is preserved under small perturbations.

6. Application: Robust Welfare Analysis

In this section we demonstrate how one can use our results to conduct robust welfare analysis in an economy where agents exhibit rational inattention. The proofs for the results in this section are in the online appendix.

We consider an economy that consists of a fixed set of agents, I , who play an information acquisition game, $(\mathcal{G}, \mathcal{T})$. The structure of the game depends on the policy enacted by a utilitarian social planner. The planner has a good understanding of the policy’s material implications, i.e., she knows a given policy leads to a given \mathcal{G} . However, the planner is unsure about the accompanying information technology \mathcal{T} .

We focus on two cases regarding the source of the agents’ information. The planner either postulates that information is exogenously given, or that it is generated by rational inattention. In both cases, the planner identifies a policy with the corresponding base game \mathcal{G} and employs a robust criterion that evaluates it according to the worst-case utilitarian welfare across all relevant information technologies \mathcal{T} and ensuing equilibria.

For the exogenous information case, Bergemann and Morris’s (2016) results imply one can find the social value of a policy \mathcal{G} by minimizing the sum of the agents’ gross payoffs across all BCEs. Specifically, let $\bar{w}(p)$ be the utilitarian welfare of an outcome p assuming the players’ payoffs are given by their gross value,

$$\bar{w}(p) = \sum_i \bar{v}_i(p).$$

Then the planner’s value of a base game \mathcal{G} under exogenous information is

$$\bar{w} = \min_{p \in BCE} \bar{w}(p),$$

where BCE is the set of Bayes correlated equilibria of \mathcal{G} . Since the BCE set is

defined by linear inequalities, and $\bar{w}(p)$ is linear in p , one can compute \bar{w} via linear programming.

What about a planner who postulates that agents are rationally inattentive? Theorem 1 provides the answer: such a planner evaluates each base game according to the lowest sum of *uninformed* values that is attainable in some *separated* BCE. More precisely, for an outcome p , let $\underline{w}(p)$ be the utilitarian welfare implied by p in the base game \mathcal{G} if players' payoffs are given by their uninformed value,

$$\underline{w}(p) = \sum_i v_i(p).$$

Then Theorem 1 suggests that a planner who assumes information is endogenous would evaluate each game according to the minimum of $\underline{w}(p)$ across all separated BCEs p ,

$$\inf_{p \in sBCE} \underline{w}(p).$$

Recall, however, that Theorem 2 says that the separation constraint does not bind for generic games. As such, imposing the separation constraint is only appropriate if the planner is absolutely certain of the structure of the base game. Whereas such certainty might be justifiable in certain cases, here we take the perspective of a cautious planner who, in an economy where agents are rationally inattentive, evaluates \mathcal{G} according to the worst-case value of $\underline{w}(p)$ across all BCE,

$$\underline{w} = \min_{p \in BCE} \underline{w}(p).$$

Since the BCE set is defined by linear inequalities, and $\underline{w}(p)$ is convex in p , one can compute \underline{w} using convex programming.

In the rest of the section we study when the optimal policy under rational inattention differs from the best policy under exogenous information. That is, we are interested in situations where there are two policies \mathcal{G} and \mathcal{G}' such that \mathcal{G} is preferred to \mathcal{G}' under the hypothesis that information is exogenous, $\bar{w} > \bar{w}'$, while \mathcal{G}' is preferred to \mathcal{G} if the planner believes that rational inattention is a relevant feature of the economy, $\underline{w}' > \underline{w}$.

Clearly, the two knowledge regimes can generate different policy prescriptions only if one of the policies leads to a different worst-case value under rational inattention than it does under exogenous information. Whereas Theorem 3 guarantees the ex-

istence of a generic \mathcal{G} for which $\bar{w} > \underline{w}$, it does not tell us when such an inequality occurs for a fixed base game. Next, we zoom in on symmetric binary-action games to provide a characterization of the full class of policies \mathcal{G} for which $\bar{w} > \underline{w}$.

A base game $\mathcal{G} = (I, \Theta, \pi, (A_i, u_i)_{i \in I})$ has **binary actions** if for every player i , A_i contains two elements. It is **symmetric** if $A_i = A_j$ for all $i, j \in I$, and if for every permutation $\phi : I \rightarrow I$, player i , action profile a , and payoff state θ ,

$$u_i(a_\phi, \theta) = u_{\phi(i)}(a, \theta),$$

where $a_\phi = (a_{\phi(j)})_{j \in I}$ is the action profile such that each player j takes action $a_{\phi(j)}$. An outcome $p \in \Delta(A \times \Theta)$ of a symmetric base game is **symmetric** if $p(a, \theta) = p(a_\phi, \theta)$ for every permutation $\phi : I \rightarrow I$, action profile a , and payoff state θ . We denote the set of symmetric outcomes by $\Delta_\pi^{sy}(A \times \Theta)$, and the set of symmetric BCEs by BCE^{sy} .

The next result characterizes the binary-action symmetric games for which rational inattention yields strictly lower worst-case welfare than exogenous information.

Proposition 1. *Let \mathcal{G} be a symmetric, binary-action base game. Then, $\underline{w} < \bar{w}$ if and only if all $p^* \in \operatorname{argmin}_{p \in \Delta_\pi^{sy}(A \times \Theta)} \underline{w}(p)$ satisfy the following condition:*

$$a_i \in \operatorname{supp}_i(p^*) \text{ and } BR(p_{a_i}^*) = \{a_i\} \quad \text{for all } i \in I \text{ and } a_i \in A_i. \quad (5)$$

Thus, one can check whether $\underline{w} < \bar{w}$ by examining the minimizers of $\underline{w}(p)$ among all symmetric outcomes p , ignoring the players' obedience constraints. In particular, one needs to check whether all these minimizers recommend both actions to every player, and only send recommendations that induce unique best responses.¹²

An immediate corollary of Proposition 1 is that in a symmetric binary-action game, if $\underline{w} < \bar{w}$ then

$$\underline{w} = \min_{p \in \Delta_\pi^{sy}(A \times \Theta)} \underline{w}(p).$$

Consistently with Theorem 3, Proposition 1 also implies that $\underline{w} = \bar{w}$ if the base game has one player only, and such player has a binary action.¹³

¹²We note that, while symmetric games are non-generic, the strict inequality identified by Proposition 1 is robust. Formally, whenever the inequality $\underline{w} < \bar{w}$ holds for some binary-action symmetric game \mathcal{G} , the inequality also holds for all base games \mathcal{G}' in a neighborhood of \mathcal{G} . The result follows from the same logic we use in the proof of Lemma 7. As in Section 5, the distance between \mathcal{G} and \mathcal{G}' is the distance between the profiles of utility functions, keeping fixed I, Θ, π , and $(A_i)_{i \in I}$.

¹³For an explanation, suppose $I = \{i\}$ and $A_i = \{a_i, b_i\}$. Take $p^* \in BCE$ such that agent i have

We now briefly explain the proposition’s proof. The key observation is that in a binary action game, a BCE p has the property that $\bar{v}_i(p) > \underline{v}_i(p)$ for every player i if and only if p satisfies (5). To get intuition for the “if” direction, note that having two recommendations that lead to strict best responses means players get a strictly positive benefit from following them. This benefit creates a wedge between the gross value $\bar{v}_i(p)$, which accounts for the value of information, and the uninformed value $\underline{v}_i(p)$, which does not. For the converse direction, note that a violation of (5) means some player i has an action that is optimal across all of the mediator’s recommendations. As such, player i loses nothing by ignoring the mediator’s recommendations and taking that action. In other words, player i ’s gross value equals her uninformed value.

Armed with the above observation, we prove the proposition in two steps. The first step shows $\underline{w} < \bar{w}$ holds if and only if all optimal solutions of $\min_{p \in BCE^{sy}} \underline{w}(p)$ satisfy (5). This step follows from applying the above-mentioned observation to symmetric outcomes, and showing such outcomes are sufficient for minimizing utilitarian welfare. The second step shows that all optimal solutions of $\min_{p \in BCE^{sy}} \underline{w}(p)$ satisfy (5) if and only if all optimal solutions of $\min_{p \in \Delta_\pi^{sy}(A \times \Theta)} \underline{w}(p)$ satisfy (5). Loosely speaking, the reason is as follows: if the obedience constraint does not bind at the optimum—as (5) dictates—then it can be relaxed, and therefore minimizing over $p \in BCE^{sy}$ is the same as minimizing over $p \in \Delta_\pi^{sy}(A \times \Theta)$.

Proposition 1 enables us to find circumstances where the planner’s optimal policy depends on whether she believes information is exogenously given, or generated by rational inattention. We demonstrate this fact below in a regime change game.

Example 2. We consider a regime change game whereby a status quo is abandoned if a sufficiently large number of players take an action against it. Such games are well-studied and have been used to model a variety of social phenomena, including currency crises, bank runs, debt crises, and political revolts.¹⁴

In our application, there are n identical investors, $i \in I = \{1, \dots, n\}$, each of which decides whether to speculate against (i.e., attack) a distressed financial institution ($a_i = 1$) or not ($a_i = 0$). Speculating costs $k \in (0, 1)$. If enough investors speculate, the institution fails (i.e., the attack succeeds), generating a profit of 1 to

no information and does not randomize between a_i and b_i . Then, $p^* \in \operatorname{argmin}_{p \in \Delta_\pi(A \times \Theta)} \underline{v}_i(p)$ and $\operatorname{supp}_i(p^*)$ is a singleton, violating (5).

¹⁴See Morris and Shin (2003) for a review.

the speculators, and an externality of $-x$ to all passive investors, where $x \in (0, \infty)$. The payoff state $\theta \in \Theta \subseteq \{1, \dots, n\}$, determines the number of speculators required for the attack to succeed. We assume that $\min \Theta > 1$, and $\max \Theta < n - 1$, meaning no single investor can go against the will of all the others.¹⁵ We summarize these payoffs below:

	$\sum_j a_j \geq \theta$	$\sum_j a_j < \theta$
$u_i(1, a_{-i}, \theta)$	$1 - k$	$-k$
$u_i(0, a_{-i}, \theta)$	$-x$	0

It is easy to verify that, in this example, rational inattention and exogenous information are outcome equivalent, that is, the sBCE set is dense in the BCE set.¹⁶

Next we demonstrate that a planner who views information as exogenous may adopt different policies than a planner who thinks investors are rationally inattentive. In particular, we show that only a planner who takes information as given has an incentive to change the institution's fundamentals, that is, Θ and π .

We begin by finding conditions under which there is a difference between the worst-case welfare under rational inattention and exogenous information. By Proposition 1, answering this question requires us to minimize the sum of the uninformed values across all symmetric outcomes, ignoring obedience constraints. The following result characterizes the solutions to this optimization problem.

Claim 1. In the regime change game, an outcome p^* minimizes $\underline{w}(p)$ over all $p \in \Delta_\pi^{sy}(A \times \Theta)$ if and only if the following conditions hold: for all payoff states θ ,

$$p^* \left(\left\{ (a, \theta) : \theta - 1 \leq \sum_i a_i \leq \theta \right\} \right) = 0, \quad (6)$$

$$p^* \left(\left\{ (a, \theta) : \sum_i a_i > \theta \right\} \right) = \frac{k}{1+x}. \quad (7)$$

The two optimality conditions have simple interpretations. First, no investor is ever pivotal. And second, the attack succeeds with probability $k/(1+x)$. To obtain these conditions, we show p^* minimizes $\underline{w}(p)$ across all $p \in \Delta_\pi^{sy}(A \times \Theta)$ only if each

¹⁵A fortiori, $n > 3$ because $2 \leq \min \Theta \leq \max \Theta < n - 1$.

¹⁶Indeed, the regime change game admits a strict BCE p where all players take both actions with positive probability (e.g., a convex combinations of the pure Nash equilibria in which everyone speculates or no-one speculates). Thus, any BCE q can be approximated by a separated BCE $(1 - \epsilon)q + \epsilon p$ as $\epsilon \rightarrow 0$.

individual investor i is indifferent between never speculating and always speculating when other investors play according to p^* . Minimizing across all symmetric outcomes satisfying this indifference condition then delivers the result.¹⁷

Combining Proposition 1 and Claim 1, we deduce that $\underline{w} < \bar{w}$ if and only if (5) holds for every symmetric outcome p^* that satisfy (6) and (7). It is easy to find *one* p^* that meets all requirements: have all investors attack together with probability $k/(1+x)$ regardless of the state, and no one speculates otherwise. Calculating the sum of the uninformed values from such p^* immediately give the worst-case welfare under rational inattention,

$$\underline{w} = \underline{w}(p^*) = -\frac{nxk}{1+x}. \quad (8)$$

Hence, under rational inattention, the planner's value decreases in the size of the externality x and the cost of betting on the financial institution's demise k , but does not depend on π ; that is, the planner's value does not depend on the institution's fundamentals.

Is there a symmetric outcome p^* that satisfies (6) and (7), but violates (5)? That is, under what conditions on the primitives of game exogenous information and rational inattention generate the same worst-case welfare? The next result provides an answer when the payoff state is binary (see Claim 15 in the online appendix for a many-state generalization).

Claim 2. Suppose $\Theta = \{\underline{\theta}, \bar{\theta}\}$, where $\bar{\theta} \geq \underline{\theta}$. Then $\underline{w} = \bar{w}$ if and only if $\bar{\theta} - \underline{\theta} \geq 3$ and

$$1 - \frac{1}{3}(\bar{\theta} - \underline{\theta})\pi(\bar{\theta}) \leq \frac{k}{1+x} \leq \frac{1}{3}(\bar{\theta} - \underline{\theta})(1 - \pi(\bar{\theta})). \quad (9)$$

In other terms, with binary θ , worst-case welfare under rational inattention is the same as under exogenous information if and only if the state is sufficiently uncertain in the sense that $\bar{\theta} - \underline{\theta} \geq 3$, and the probability of $\bar{\theta}$ is not too extreme compared to $k/(1+x)$ —e.g., (9) fails if $\pi(\bar{\theta})$ goes to zero or one. In particular, the worst-case welfare under rational inattention is always strictly below the welfare under exogenous information when there's certainty about the institution's fundamentals (i.e., when θ is deterministic).

The proof of Claim 2 is rather detailed; here we provide only a rough intuition. Combining Proposition 1 and Claim 1, one can show the worst-case welfare under ex-

¹⁷Given the minmax nature of the optimization problem $\min_{p \in \Delta_{\pi^*}^{sy}(A \times \Theta)} \underline{w}(p)$, one might have conjectured the existence of an indifference condition.

ogenous information coincides with the worst-case welfare under rational inattention only if a symmetric outcome exists that satisfies two conditions. First, an investor who does not see the mediator’s recommendation is indifferent between attacking and not attacking. And second, for an investor who does see the mediator’s recommendation, either not speculating is a best response to a “speculate” recommendation, or vice versa. Appealing to Bayes rule, one can show that an outcome satisfies these conditions if and only if there is limited overlap between the event where many investors are attacking and the event in which the attack succeeds. Claim 2 follows from showing a sufficient disconnect between these two events is attainable if and only if there is enough uncertainty about θ . Intuitively, disconnecting the two events is easy when θ obtains both high and low values with large probability: one can have the number of speculators just come short of a successful attack when θ is high, and come just above the threshold when θ is low. The same cannot be done when θ is deterministic. In that case, successful attacks necessarily involve more speculating investors than failed ones.

A takeaway is that, unlike under rational inattention, a planner who views information as exogenous may adopt policies that change the institution’s fundamentals. For a concrete illustration, consider two policies \mathcal{G} and \mathcal{G}' that differ only in the institution’s fundamentals, that is, in the set of states Θ and their distribution $\pi \in \Delta(\Theta)$. Suppose \mathcal{G} satisfies the conditions of Claim 2, but \mathcal{G}' does not. Worst-case welfare under rational inattention is the same for \mathcal{G} and \mathcal{G}' : $\underline{w} = \underline{w}'$ by (8). By contrast, \mathcal{G} generates lower welfare under exogenous information: $\bar{w} < \bar{w}'$ by (8) and Claim 2. Consequently, a planner who views information as being exogenous would pay some amount to change the institution’s fundamentals, but a planner who believes investors are rationally inattentive should not do so.

7. Further results: Non-Generic Environments

Throughout the paper, we have focused on generic settings. However, many economic institutions of interests, such as auctions, are non-generic. We conclude our work with a discussion of rational inattention in non-generic environments. We characterize the (non-generic) settings in which the separation constraint has substantive impact, and show its force has an all-or-nothing flavor: if the sBCE set is not dense in the BCE set, it is in fact nowhere dense. The proofs are in the online appendix.

We begin by characterizing the environments where the separation constraint has bite. Our characterization is based on two definitions. The first definition represents a situation where, whenever the mediator recommends player i action b_i in a BCE of the game, the player would be equally happy to take action a_i . Myerson (1997) calls this scenario “jeopardization”: action a_i **jeopardizes** action b_i if, for every BCE p such that $b_i \in \text{supp}_i(p)$, $a_i \in BR(p_{b_i})$.¹⁸ We denote by $J(b_i)$ the set of actions that jeopardizes b_i .

Every action jeopardizes itself by the obedience constraint; hence, $J(b_i)$ is not empty. A sufficient condition for jeopardization is weak domination: if $u_i(a_i, a_{-i}, \theta) \geq u_i(b_i, a_{-i}, \theta)$ for all $a_{-i} \in A_{-i}$ and $\theta \in \Theta$, then a_i jeopardizes b_i . But the concept of jeopardization is broader than weak domination. For example, in Matching Pennies, Heads and Tails jeopardize each other, even if neither action is weakly dominant.

Next we introduce the second definition required for our characterization. This definition concerns a class of BCEs where the number of actions that induce different beliefs is minimal. To state this definition, say a BCE p is **maximally supported** if the support of every other BCE is contained by the support of p . A maximally-supported BCE p is **minimally mixed** if $q_{a_i} \neq q_{b_i}$ implies $p_{a_i} \neq p_{b_i}$ for every BCE q , $i \in I$, and $a_i, b_i \in \text{supp}_i(q)$.

For an interpretation of minimal mixing, consider a mediator who wants to implement a BCE p . When $p_{a_i} = p_{b_i}$, the mediator can replace the distinct recommendations of playing a_i and b_i with a single recommendation of mixing between the two actions with probabilities $p(a_i)/(p(a_i) + p(b_i))$ and $p(b_i)/(p(a_i) + p(b_i))$. A BCE p is minimally mixed if a mediator has the least amount of opportunities to implement p recommending mixed actions. Whereas minimally mixed BCEs seem esoteric at first, they are in fact, ubiquitous: the set of minimally mixed BCEs is open and dense in the BCE set (see Lemma 13 in the online appendix).

Our next result uses the concepts of jeopardization and minimally mixed BCEs to characterize when the BCE set and the closure of the sBCE set coincide.

Proposition 2. *The following statements are equivalent:*

- (i) *The sBCE set is dense in the BCE set.*
- (ii) *A minimally mixed sBCE exists.*

¹⁸Myerson defines jeopardization for games without payoff uncertainty; here we give the obvious extension to games where Θ is not a singleton.

(iii) For every BCE p , $i \in I$, $a_i, b_i \in \text{supp}_i(p)$, if $p_{a_i} \neq p_{b_i}$ then $J(a_i) \cap J(b_i) = \emptyset$.

The result shows how jeopardization and minimal mixing can be used in applications to study sBCE. To verify that the sBCE set is dense in the BCE set, it is enough to produce a minimally mixed sBCE. To verify that the sBCE set is not dense in the BCE set, it is enough to produce a BCE in which two actions induce distinct beliefs and share a common jeopardizing action. As shown by Myerson (1997), the jeopardizing actions can be easily computed from the dual of the system of linear inequalities that defines BCE.

Next, we build on Proposition 2 and obtain that sBCE is an all-or-nothing refinement of BCE.

Proposition 3. *The sBCE set is either dense or nowhere dense in the BCE set.*

For a rough explanation, consider first the case in which a minimally mixed sBCE exists. Then, by Proposition 2, the sBCE set is dense in the BCE set. Consider now the case in which a minimally mixed sBCE does not exist. By Proposition 2, the sBCE set is not dense in the BCE set. To reach the stronger conclusion that the sBCE set is nowhere dense in the BCE set, we use the fact that the set of minimally mixed BCE is open and dense in the BCE set.

Thus, whereas the separation constraint does not bind in most circumstances, whenever it does bind, it significantly restricts the set of attainable outcomes.

Appendix

A. Proof of Theorem 1

The proof of the theorem proceeds as follows. First, we recall a single-agent lemma due Denti (2021). Next, with the help of the lemma, we prove the “if” and “only if” statements of the theorem.

A.1. A single-agent lemma

We take the perspective of an individual i who has to choose an action $a_i \in A_i$ whose utility $w_i(a_i, \omega)$ depends on an uncertain state of nature $\omega \in \Omega$. Both A_i and

Ω are finite. Let $\rho \in \Delta(\Omega)$ be the prior distribution of the state; ρ may not have full support. Before choosing an action, the decision maker can run an experiment $\xi_i : \Omega \rightarrow \Delta(X_i)$ at a cost $C_i(\xi_i) \in \mathbb{R}_+$. Let \mathcal{E}_i be the set of feasible experiments. The signal space X_i is finite. Overall, the decision maker faces the following information acquisition problem:

$$\max_{\xi_i \in \mathcal{E}_i, \sigma_i \in \Sigma_i} \left[\sum_{\omega, x_i, a_i} w_i(a_i, \omega) \sigma_i(a_i | x_i) \xi_i(x_i | \omega) \rho(\omega) \right] - C_i(\xi_i) \quad (10)$$

where Σ_i is the set of all action plans $\sigma_i : X_i \rightarrow \Delta(A_i)$.

In accordance with the terminology used in the main text, $\xi_i \succsim \xi'_i$ if there is a Markov kernel $g : X_i \rightarrow \Delta(X_i)$ such that for every $x_i \in X_i$ and $\omega \in \Omega$ with $\rho(\omega) > 0$,

$$\xi'_i(x_i | \omega) = \sum_{x'_i \in X_i} g(x_i | x'_i) \xi(x'_i | \omega).$$

We say that \mathcal{E}_i is **flexible** if, whenever $\xi_i \succsim \xi'_i$ and $\xi_i \in \mathcal{E}_i$, then $\xi'_i \in \mathcal{E}_i$. We also say that $C_i : \mathcal{E}_i \rightarrow \mathbb{R}_+$ is **monotone** if, whenever $\xi_i, \xi'_i \in \mathcal{E}_i$ are such that $\xi_i \succsim \xi'_i$ (resp., $\xi_i \succ \xi'_i$), then $C_i(\xi_i) \geq C_i(\xi'_i)$ (resp., $C_i(\xi_i) > C_i(\xi'_i)$).

Among other results, Denti (2021) characterizes the pairs (ξ_i, σ_i) that are optimal solutions of (10) for some flexible \mathcal{E}_i and monotone C_i . To describe the characterization, let $\mu(x_i) = \sum_{\omega} \xi_i(x_i | \omega) \rho(\omega)$ be the ex-ante probability that ξ_i generates x_i . For every x_i such that $\mu(x_i) > 0$, we denote by $\mu_{x_i} \in \Delta(\Omega)$ the posterior distribution of the state ω . Finally, let $BR(\mu_{x_i}) = \arg \max_{a_i \in A_i} \sum_{\omega} w_i(a_i, \omega) \mu_{x_i}(\omega)$ be the set of best responses to μ_{x_i} .

Lemma 1 (Denti, 2021). *A flexible \mathcal{E}_i and a monotone C_i exist such that the pair (ξ_i, σ_i) is an optimal solution of (10) if and only if the following conditions hold:*

- (i) *For all x_i with $\mu(x_i) > 0$, $\sigma_i(BR(\mu_{x_i}) | x_i) = 1$.*
- (ii) *For all x_i and x'_i with $\mu(x_i)\mu(x'_i) > 0$, $\mu_{x_i} \neq \mu_{x'_i}$ implies $BR(\mu_{x_i}) \cap BR(\mu_{x'_i}) = \emptyset$.*

In addition, one can choose C_i so that

$$C_i(\xi_i) = \sum_{\omega, x_i, a_i} w_i(a_i, \omega) \sigma_i(a_i | x_i) \xi_i(x_i | \omega) \rho(\omega) - \max_{a_i \in A_i} \sum_{\omega} w_i(a_i, \omega) \rho(\omega).$$

A.2. Proof of the “if” statement of Theorem 1.

Let (p, v) be an outcome-value pair that satisfies Theorem 1-(i) and Theorem 1-(ii). First, we consider the case in which $v = \underline{v}(p)$:

Lemma 2. *There exist a rational-inattention technology $\mathcal{T} = (Z, \zeta, (X_i, \mathcal{E}_i, C_i)_{i \in I})$ and an equilibrium (ξ, σ) of $(\mathcal{G}, \mathcal{T})$ whose outcome-value pair is $(p, \underline{v}(p))$.*

Proof. For every player i and pair of actions $a_i, b_i \in \text{supp}_i(p)$, let $a_i \sim_i b_i$ be $p_{a_i} = p_{b_i}$. Note that \sim_i is an equivalence relation on $\text{supp}_i(p)$. Let Z_i be the corresponding set of equivalence classes.

We take $Z = \prod_{i \in I} Z_i$. To simplify the exposition, we write $a \in z$ if $a_i \in z_i$ for all $i \in I$. Given this convention, we define $\zeta : \Theta \rightarrow Z$ by $\zeta(z|\theta) = \sum_{a \in z} p(a, \theta) / \pi(\theta)$. For every player i , we take $X_i = Z_i$, and we define $\xi_i : Z \times \Theta \rightarrow \Delta(X_i)$ by $\xi_i(x_i|z, \theta) = 1$ if $x_i = z_i$, and $\xi_i(x_i|z, \theta) = 0$ otherwise. We also define $\sigma_i : X_i \rightarrow \Delta(A_i)$ by $\sigma_i(a_i|x_i) = p(a_i) / \sum_{b_i \in x_i} p(b_i)$ if $a_i \in x_i$, and $\sigma_i(a_i|x_i) = 0$ otherwise.

Claim 3. The outcome induced by (ξ, σ) is p , that is, for all $a \in A$ and $\theta \in \Theta$

$$p(a, \theta) = \sum_z \left[\prod_{i \in I} \sigma_i(a_i|z_i) \right] \zeta(z|\theta) \pi(\theta). \quad (11)$$

To prove the claim, we need an intermediate result.

Claim 4. For all $z \in Z$, $a \in z$, and $\theta \in \Theta$,

$$p(a, \theta) = \left[\prod_{i \in I} \sigma_i(a_i|z_i) \right] \sum_{b \in z} p(b, \theta).$$

Proof. It suffices to show that for all $i \in I$,

$$p(a, \theta) = \sigma_i(a_i|z_i) \sum_{b_i \in z_i} p(b_i, a_{-i}, \theta). \quad (12)$$

The desired result then follows from reasoning by induction on the number of players.

So, fix some player i . By Bayes rule, $p(a, \theta) = p(a_i) p_{a_i}(a_{-i}, \theta)$. Recall that $b_i \in z_i$

if and only if $p_{b_i} = p_{a_i}$; in addition, $\sum_{b_i \in z_i} \sigma_i(b_i|z_i) = 1$. Therefore,

$$\begin{aligned} p(a_i, a_{-i}, \theta) &= p(a_i)p_{a_i}(a_{-i}, \theta) = p(a_i) \sum_{b_i \in z_i} \sigma_i(b_i|z_i) p_{a_i}(a_{-i}, \theta) \\ &= p(a_i) \sum_{b_i \in z_i} \sigma_i(b_i|z_i) p_{b_i}(a_{-i}, \theta). \end{aligned}$$

Substituting in the definition of σ_i , one obtains (12). \square

Proof of Claim 3. Fix $a \in A$ and $\theta \in \Theta$. If $p(a_i) = 0$ for some player i , then $\sigma(a_i|x_i) = 0$ for all $x_i \in X_i$. Thus, both sides of (11) are equal to zero.

Suppose now that $p(a_i) > 0$ for all players i . Take $z^* \in Z$ such that $a \in z^*$. Using Claim 4 we get that

$$p(a, \theta) = \left[\prod_{i \in I} \sigma_i(a_i|z_i^*) \right] \sum_{b \in z^*} p(b, \theta).$$

The equality (11) follows from noting that $\prod_{i \in I} \sigma_i(a_i|z_i) > 0$ if and only if $z = z^*$. \square

Next, for each player i , we use Lemma 1 to construct \mathcal{E}_i and C_i . Given $\Omega = Z \times \Theta$, we define $\rho \in \Delta(\Omega)$ and $w_i : A_i \times \Omega \rightarrow \mathbb{R}$ by

$$\begin{aligned} \rho(z, \theta) &= \zeta(z|\theta)\pi(\theta), \\ w_i(a_i, z, \theta) &= \sum_{a_{-i}} u_i(a_i, a_{-i}, \theta) \left[\prod_{j \neq i} \sigma_j(a_j|z_j) \right]. \end{aligned}$$

As in Section A.1, let $\mu(x_i)$ be the ex-ante probability that ξ_i generates x_i , and let $\mu_{x_i} \in \Delta(\Omega)$ be i 's posterior beliefs about $\omega = (z, \theta)$ when x_i is observed. Note that μ_{x_i} is different from p_{a_i} , which is i 's posterior beliefs about (a_{-i}, θ) when the mediator recommends a_i . Recall that $BR(\mu_{x_i}) = \arg \max_{b_i} \sum_{\omega} w_i(b_i, \omega) \mu_{x_i}(\omega)$, while $BR(p_{a_i}) = \arg \max_{b_i} \sum_{\theta, a_{-i}} u_i(b_i, a_{-i}, \theta) p_{a_i}(a_{-i}, \theta)$.

To be able to use Lemma 1, we need the following result.

Claim 5. For all $x_i \in X_i$ and $a_i \in x_i$, $BR(p_{a_i}) = BR(\mu_{x_i})$.

Proof. It is enough to show that for all $a_{-i} \in A_{-i}$ and $\theta \in \Theta$,

$$p_{a_i}(a_{-i}, \theta) = \sum_z \left[\prod_{j \neq i} \sigma_j(a_j|z_j) \right] \mu_{x_i}(z, \theta). \quad (13)$$

Towards this goal, we first observe that for all $a_{-i} \in A_{-i}$ and $\theta \in \Theta$,

$$\begin{aligned}
p(a_i, a_{-i}, \theta) &= \sum_{x'_i, x_{-i}} \left[\sigma_i(a_i|x'_i) \prod_{j \neq i} \sigma_j(a_j|x_j) \right] \zeta(x'_i, x_{-i}|\theta) \pi(\theta) \\
&= \sigma_i(a_i|x_i) \sum_{x_{-i}} \left[\prod_{j \neq i} \sigma_j(a_j|x_j) \right] \zeta(x_i, x_{-i}|\theta) \pi(\theta) \\
&= \frac{p(a_i)}{\sum_{b_i \in x_i} p(b_i)} \sum_{x_{-i}} \left[\prod_{j \neq i} \sigma_j(a_j|x_j) \right] \zeta(x_i, x_{-i}|\theta) \pi(\theta), \tag{14}
\end{aligned}$$

where the first equality holds by Claim 3, the second equality because $\sigma_i(a_i|x'_i) > 0$ if and only if $x'_i = x_i$, and the third equality by definition of $\sigma_i(a_i|x_i)$. Summing both sides of (14) over a_{-i} and θ , we obtain that

$$p(a_i) = \frac{p(a_i)\mu(x_i)}{\sum_{b_i \in x_i} p(b_i)}.$$

Thus, $\mu(x_i) = \sum_{b_i \in x_i} p(b_i)$. We plug this finding in (14) and obtain that for all $a_{-i} \in A_{-i}$ and $\theta \in \Theta$,

$$p(a_i, a_{-i}, \theta) = \frac{p(a_i)}{\mu(x_i)} \sum_{x_{-i}} \left[\prod_{j \neq i} \sigma_j(a_j|x_j) \right] \zeta(x_i, x_{-i}|\theta) \pi(\theta).$$

The equality (13) follows from observing that $\mu_{x_i}(z, \theta) = \zeta(z|\theta)\pi(\theta)/\mu(x_i)$ if $z_i = x_i$, and $\mu_{x_i}(z, \theta) = 0$ otherwise. \square

To verify Lemma 1-(i), take $x_i \in X_i$ and $a_i \in A_i$ such that $\sigma_i(a_i|x_i) > 0$. By the construction of the action plan, $a_i \in x_i$. By the obedience constraint for p , $a_i \in BR(p_{a_i})$. By Claim 5, $BR(p_{a_i}) = BR(\mu_{x_i})$. Therefore, $a_i \in BR(\mu_{x_i})$. We deduce that Lemma 1-(i) holds.

To verify Lemma 1-(ii), take $x_i, x'_i \in X_i$ such that $\mu_{x_i} \neq \mu_{x'_i}$. Let $a_i \in x_i$ and $b_i \in x'_i$. Since $\mu_{x_i} \neq \mu_{x'_i}$, we must have $x_i \neq x'_i$ and, therefore, $p_{a_i} \neq p_{b_i}$. By the separation constraint for p , $BR(p_{a_i}) \cap BR(p_{b_i}) = \emptyset$. By Claim 5, $BR(p_{a_i}) = BR(\mu_{x_i})$ and $BR(p_{b_i}) = BR(\mu_{x'_i})$. We deduce $BR(\mu_{x_i}) = BR(\mu_{x'_i}) = \emptyset$: Lemma 1-(ii) holds.

In sum, for every player i , we can invoke Lemma 1 to find a flexible \mathcal{E}_i and a monotone C_i such that (ξ_i, σ_i) is a best reply to (ξ_{-i}, σ_{-i}) . This shows that $\mathcal{T} = (Z, \zeta, (X_i, \mathcal{E}_i, C_i)_{i \in I})$ is a rational-inattention technology such that (ξ, σ) is

an equilibrium $(\mathcal{G}, \mathcal{T})$ with outcome p . In addition, as in the last part of Lemma 1, we can choose costs so that, for every player i , $C_i(\xi_i) = \bar{v}_i(p) - \underline{v}_i(p)$, that is, $\bar{v}_i(p) - C_i(\xi_i) = \underline{v}_i(p)$. \square

Now we consider the general case in which v may differ from $\underline{v}(p)$. As a starting point, take $\mathcal{T} = (Z, \zeta, (X_i, \mathcal{E}_i, C_i)_{i \in I})$ and (ξ, σ) as in Lemma 2. Without loss of generality, we assume that $\mathcal{E}_i = \{\xi'_i : \xi_i \succsim \xi'_i\}$ for every player i .

For each $i \in I$, take $\lambda_i \in (0, 1]$ and define $C_{\lambda_i} : \mathcal{E}_i \rightarrow \mathbb{R}_+$ by $C_{\lambda_i}(\xi'_i) = \lambda_i C_i(\xi'_i)$. Since C_i is monotone, C_{λ_i} is monotone. Given $\lambda = (\lambda_i)_{i \in I}$, consider the rational-inattention technology $\mathcal{T}_\lambda = (Z, \zeta, (X_i, \mathcal{E}_i, C_{\lambda_i})_{i \in I})$.

Lemma 3. *The strategy profile (ξ, σ) is an equilibrium of $(\mathcal{G}, \mathcal{T}_\lambda)$.*

Proof. For a player i , consider an alternative strategy (ξ'_i, σ'_i) . With a slight abuse of notation, we denote by $u_i(\xi'_i, \sigma'_i, \xi_{-i}, \sigma_{-i})$ the expected utility player i obtains by deviating to (ξ'_i, σ'_i) . Since (ξ, σ) is an equilibrium of $(\mathcal{G}, \mathcal{T})$,

$$\bar{v}_i(p) - C_i(\xi_i) \geq u_i(\xi'_i, \sigma'_i, \xi_{-i}, \sigma_{-i}) - C_i(\xi'_i).$$

In addition, since $\xi_i \succeq \xi'_i$ (recall: $\mathcal{E}_i = \{\xi''_i : \xi_i \succsim \xi''_i\}$), $\bar{v}_i(p) \geq u_i(\xi'_i, \sigma'_i, \xi_{-i}, \sigma_{-i})$. It follows that

$$\bar{v}_i(p) - \lambda_i C_i(\xi_i) \geq u_i(\xi'_i, \sigma'_i, \xi_{-i}, \sigma_{-i}) - \lambda_i C_i(\xi'_i).$$

We deduce the deviation is not profitable. \square

Thus, (ξ, σ) is an equilibrium of $(\mathcal{G}, \mathcal{T}_\lambda)$ whose outcome is p . In addition,

$$\bar{v}_i(p) - C_{\lambda_i}(\xi_i) = \bar{v}_i(p) - \lambda_i C_i(\xi_i) = \bar{v}_i(p) - \lambda_i(\bar{v}_i(p) - \underline{v}_i(p)) = (1 - \lambda_i)\bar{v}_i + \lambda_i \underline{v}_i(p).$$

Hence, by appropriately choosing the vector λ , we can be sure that v is the value of (ξ, σ) . This concludes the proof of the “if” direction of Theorem 1.

A.3. Proof of the “only if” statement of Theorem 1.

Let (p, v) be the outcome-value pair induced by an equilibrium (ξ, σ) of an information acquisition game $(\mathcal{G}, \mathcal{T})$ where \mathcal{T} represents rational inattention.

First we show that p is a separated BCE. For every player i , by the definition of equilibrium, the equilibrium strategy (ξ_i, σ_i) maximizes

$$\left[\sum_{a,x,z,\theta} u_i(a, \theta) \sigma'_i(a_i|x_i) \xi'_i(x_i|z, \theta) \prod_{j \neq i} \sigma_j(a_j|x_j) \xi_j(x_j|z, \theta) \zeta(z|\theta) \pi(\theta) \right] - C_i(\xi_i)$$

over all $\xi'_i \in \mathcal{E}_i$ and $\sigma'_i \in \Sigma_i$. Since \mathcal{E}_i is flexible and C_i is monotone, we can apply Lemma 1 with appropriate definitions for Ω , ρ , and w_i . Specifically, we take $\Omega = Z \times \Theta$, and we define $\rho \in \Delta(\Omega)$ and $w_i : A_i \times \Omega \rightarrow \mathbb{R}$ by

$$\begin{aligned} \rho(z, \theta) &= \zeta(z|\theta) \pi(\theta), \\ w_i(a_i, z, \theta) &= \sum_{x_{-i}, a_{-i}} u_i(a, \theta) \left[\prod_{j \neq i} \sigma_j(a_j|x_j) \xi_j(x_j|z, \theta) \right]. \end{aligned}$$

From Lemma 1-(i), we obtain $\sigma_i(BR(\mu_{x_i})|\mu_{x_i}) = 1$ for all x_i with $\mu(x_i) > 0$. From Lemma 1-(ii), we get that for all x_i and x'_i such that $\mu(x_i) > 0$ and $\mu(x'_i) > 0$, if $\mu_{x_i} \neq \mu_{x'_i}$ then $BR(\mu_{x_i}) \neq BR(\mu_{x'_i})$.

Next, we relate $BR(\mu_{x_i})$ and $BR(p_{a_i})$. Towards this goal, for each player i and action a_i , let X_{a_i} be the set of positive-probability signals that makes player i take action a_i : $X_{a_i} = \{x_i : \mu(x_i) > 0 \text{ and } \sigma(a_i|x_i) > 0\}$.

Lemma 4. *For all $i \in I$ and $a_i \in \text{supp}_i(p)$, $BR(p_{a_i}) = \bigcap_{x_i \in X_{a_i}} BR(\mu_{x_i})$.*

Proof. First, take $b_i \in BR(p_{a_i})$: we wish to show that $b_i \in BR(\mu_{x_i})$ for all $x_i \in X_{a_i}$. By Lemma 1-(i), $a_i \in BR(\mu_{x_i})$ for all $x_i \in X_{a_i}$. Hence,

$$\sum_{z,\theta} w_i(a_i, z, \theta) \mu_{x_i}(z, \theta) \geq \sum_{z,\theta} w_i(b_i, z, \theta) \mu_{x_i}(z, \theta). \quad (15)$$

For each $x_i \in X_{a_i}$, define $\mu(x_i|a_i) = \sigma_i(a_i|x_i) \nu(x_i) / p(a_i)$. Simple algebra shows that

$$p_{a_i}(a_{-i}, \theta) = \sum_{x_i \in X_{a_i}} \mu(x_i|a_i) \sum_{x_{-i}, z} \left[\prod_{j \neq i} \sigma_j(a_j|x_j) \xi_j(x_j|z, \theta) \right] \mu_{x_i}(z, \theta). \quad (16)$$

Therefore,

$$\begin{aligned} \sum_{x_i \in X_{a_i}} \mu(x_i|a_i) \left(\sum_{z, \theta} w_i(a_i, z, \theta) \mu_{x_i}(z, \theta) \right) &= \sum_{a_{-i}, \theta} u_i(a_i, a_{-i}, \theta) p_{a_i}(a_{-i}, \theta), \\ \sum_{x_i \in X_{a_i}} \mu(x_i|a_i) \left(\sum_{z, \theta} w_i(b_i, z, \theta) \mu_{x_i}(z, \theta) \right) &= \sum_{a_{-i}, \theta} u_i(b_i, a_{-i}, \theta) p_{a_i}(a_{-i}, \theta). \end{aligned}$$

As a consequence, since $b_i \in BR(p_{a_i})$, we have that

$$\sum_{x_i \in X_{a_i}} \mu(x_i|a_i) \left(\sum_{z, \theta} w_i(a_i, z, \theta) \mu_{x_i}(z, \theta) \right) \leq \sum_{x_i \in X_{a_i}} \mu(x_i|a_i) \left(\sum_{z, \theta} w_i(b_i, z, \theta) \mu_{x_i}(z, \theta) \right).$$

It follows that (15) holds with equality for all $x_i \in X_{a_i}$. Therefore, since $a_i \in BR(\mu_{x_i})$ for all $x_i \in X_{a_i}$, we deduce that $b_i \in BR(\mu_{x_i})$ for all $x_i \in X_{a_i}$.

Next, take $b_i \in A_i$ such that $b_i \in BR(\nu_{x_i})$ for all $x_i \in X_{a_i}$: we wish to show that $b_i \in BR(p_{a_i})$. By hypothesis, for all $x_i \in X_{a_i}$ and $c_i \in A_i$,

$$\sum_{z, \theta} w_i(c_i, z, \theta) \mu_{x_i}(z, \theta) \leq \sum_{z, \theta} w_i(b_i, z, \theta) \mu_{x_i}(z, \theta).$$

Averaging across inequalities over $x_i \in X_{a_i}$, we obtain that for all $c_i \in A_i$,

$$\sum_{x_i \in X_{a_i}} \mu(x_i|a_i) \left(\sum_{z, \theta} w_i(c_i, z, \theta) \mu_{x_i}(z, \theta) \right) \leq \sum_{x_i \in X_{a_i}} \mu(x_i|a_i) \left(\sum_{z, \theta} w_i(b_i, z, \theta) \mu_{x_i}(z, \theta) \right).$$

Using (16), we deduce that

$$\sum_{a_{-i}, \theta} u_i(c_i, a_{-i}, \theta) p_{a_i}(a_{-i}, \theta) \leq \sum_{a_{-i}, \theta} u_i(b_i, a_{-i}, \theta) p_{a_i}(a_{-i}, \theta).$$

We conclude that $b_i \in BR(p_{a_i})$. □

We are ready to show that p is a separated BCE. For the obedience constraint, take $i \in I$ and $a_i \in \text{supp}_i(p)$. By Lemma 1-(i), $a_i \in BR(\mu_{x_i})$ for all $x_i \in X_{a_i}$. By Lemma 4, $a_i \in BR(p_{a_i})$. For the separation constraint, take $i \in I$ and $a_i, b_i \in \text{supp}_i(p)$ such that $p_{a_i} \neq p_{b_i}$. Since $p_{a_i} \neq p_{b_i}$, there are $x_{a_i} \in X_{a_i}$ and $x_{b_i} \in X_{b_i}$ such that $\mu_{x_{a_i}} \neq \mu_{x_{b_i}}$ (this should be obvious, but for confirmation, see (16)). By Lemma 1-(ii), $BR(\mu_{x_{a_i}}) \cap BR(\mu_{x_{b_i}}) = \emptyset$. By Lemma 4, $BR(p_{a_i}) \cap BR(p_{b_i}) = \emptyset$.

To complete the proof of the “only if” direction of Theorem 1, we need to show that for every player i , either $v_i = \underline{v}_i(p) = \bar{v}_i(p)$ or $v_i \in [\underline{v}_i(p), \bar{v}_i(p)]$.

Fix a player i . Since $C_i(\xi_i) \geq 0$, it must be $v_i \leq \bar{v}_i(p)$. In addition, uninformative experiments have zero cost by hypothesis. Thus, since (ξ_i, σ_i) is a best response to (ξ_{-i}, σ_{-i}) , we have that for any uninformative experiment ξ'_i ,

$$v_i \geq \max_{\sigma'_i} \sum_{a,x,z,\theta} u_i(a, \theta) \left[\sigma'_i(a_i|x_i) \xi_i(x_i|z, \theta) \prod_{j \neq i} \sigma_j(a_j|x_j) \xi_j(x_j|z, \theta) \right] \zeta(z|\theta) \pi(\theta) = \underline{v}_i(p).$$

Overall, we conclude that $v_i \in [\underline{v}_i(p), \bar{v}_i(p)]$

If $\underline{v}_i(p) = \bar{v}_i(p)$, then $v_i \in [\underline{v}_i(p), \bar{v}_i(p)]$ implies $v_i = \underline{v}_i(p) = \bar{v}_i(p)$. Suppose instead that $\underline{v}_i(p) < \bar{v}_i(p)$. If

$$\sum_{a,x,z,\theta} u_i(a, \theta) \left[\prod_{j \in I} \sigma_j(a_j|x_j) \xi_j(x_j|z, \theta) \right] \zeta(z|\theta) \pi(\theta) < \bar{v}_i(p),$$

then obviously $v_i < \bar{v}_i(p)$ because $C_i(\xi_i) \geq 0$. If, on the other hand,

$$\sum_{a,x,z,\theta} u_i(a, \theta) \left[\prod_{j \in I} \sigma_j(a_j|x_j) \xi_j(x_j|z, \theta) \right] \zeta(z|\theta) \pi(\theta) = \bar{v}_i(p),$$

then ξ_i cannot be uninformative because $\underline{v}_i(p) < \bar{v}_i(p)$. By monotonicity, $C_i(\xi_i) > 0$, which implies that $v_i < \bar{v}_i(p)$. In sum, if $\underline{v}_i(p) < \bar{v}_i(p)$, then $v_i \in [\underline{v}_i(p), \bar{v}_i(p)]$.

B. Proof of Theorem 2

Let $\|u\|$ be the Euclidean norm of $u \in \mathbb{R}^{I \times A \times \Theta}$. We begin with the following lemma.

Lemma 5. *For every $u \in \mathbb{R}^{I \times A \times \Theta}$, $p \in BCE(u)$, and $\epsilon > 0$, there exists $u' \in \mathbb{R}^{I \times A \times \Theta}$ such that $\|u - u'\| \leq \epsilon$ and $p \in sBCE(u')$.*

Proof. For each player $i \in I$, we consider a set $P_i \subseteq \Delta(A_{-i} \times \Theta)$ given by $P_i = \{p_{a_i} : a_i \in \text{supp}_i(p)\}$. Let n_i be the cardinality of P_i (of course, it could be that n_i is smaller than the cardinality of $\text{supp}_i(p)$). Reasoning inductively, we can find an enumeration p_1, \dots, p_{n_i} of the elements of P_i such that, for every $m_i \in \{1, \dots, n_i\}$, p_{m_i} is an extreme point of the convex hull of $\{p_1, \dots, p_{m_i}\}$.

By an hyperplane separation theorem (e.g., Rockafellar, 1970, Corollary 11.4.2) for every $m_i \in \{2, \dots, n_i\}$ we can find a function $f_{m_i} : A_{-i} \times \Theta \rightarrow \mathbb{R}$ such that

$$\sum_{a_{-i}, \theta} f_{m_i}(a_{-i}, \theta) p_{m_i}(a_{-i}, \theta) > 0 \geq \max_{l_i \in \{1, \dots, m_i - 1\}} \sum_{a_{-i}, \theta} f_{m_i}(a_{-i}, \theta) p_{l_i}(a_{-i}, \theta). \quad (17)$$

For $m_i = 1$, we define $f_1(a_{-i}, \theta) = 1$ for all $a_{-i} \in A_{-i}$ and $\theta \in \Theta$.

For every $l_i \in \{1, \dots, n_i - 1\}$, we choose $t_{l_i} \in (0, 1]$ such that for every $m_i \in \{l_i + 1, \dots, n_i\}$,

$$\sum_{a_{-i}, \theta} f_{m_i}(a_{-i}, \theta) p_{m_i}(a_{-i}, \theta) > t_{l_i} \sum_{a_{-i}} f_{l_i}(a_{-i}) p_{m_i}(a_{-i}). \quad (18)$$

We can choose such a t_{l_i} because the left-hand side of (18) is positive—see (17). For $l_i = n_i$, we simply define $t_{n_i} = 1$.

For every $l_i \in \{1, \dots, n_i\}$, we define $s_{l_i} = \prod_{m_i=l_i}^{n_i} t_{m_i}$. Using (18), simple algebra shows that for every $l_i \in \{1, \dots, n_i - 1\}$ and $m_i \in \{l_i + 1, \dots, n_i\}$,

$$s_{m_i} \sum_{a_{-i}, \theta} f_{m_i}(a_{-i}, \theta) p_{m_i}(a_{-i}, \theta) > s_{l_i} \sum_{a_{-i}, \theta} f_{l_i}(a_{-i}, \theta) p_{m_i}(a_{-i}, \theta). \quad (19)$$

For $a_i \in \text{supp}_i(p)$, we define $g_{a_i} : A_{-i} \times \Theta \rightarrow \mathbb{R}$ by $g_{a_i}(a_{-i}, \theta) = s_{m_i} \cdot f_{m_i}(a_{-i}, \theta)$ where m_i is such that $p_{a_i} = p_{m_i}$. For $a_i \notin \text{supp}_i(p)$, we define $g_{a_i} = 0$.

Claim 6. For all $a_i \in \text{supp}_i(p)$ and $b_i \notin \{c_i \in \text{supp}_i(p) : p_{c_i} = p_{a_i}\}$,

$$\sum_{a_{-i}, \theta} g_{a_i}(a_{-i}, \theta) p_{a_i}(a_{-i}, \theta) > \sum_{a_{-i}, \theta} g_{b_i}(a_{-i}, \theta) p_{a_i}(a_{-i}, \theta). \quad (20)$$

Proof. Pick $m_i \in \{1, \dots, n_i\}$ such that $p_{a_i} = p_{m_i}$. From the left-hand side of (17) and the fact that $s_{m_i} > 0$, we obtain that

$$\sum_{a_{-i}, \theta} g_{a_i}(a_{-i}, \theta) p_{a_i}(a_{-i}, \theta) = s_{m_i} \sum_{a_{-i}, \theta} f_{m_i}(a_{-i}, \theta) p_{m_i}(a_{-i}, \theta) > 0. \quad (21)$$

Hence, for $b_i \notin \text{supp}_i(p)$, we have

$$\sum_{a_{-i}, \theta} g_{a_i}(a_{-i}, \theta) p_{a_i}(a_{-i}, \theta) > 0 = \sum_{a_{-i}, \theta} g_{b_i}(a_{-i}, \theta) p_{a_i}(a_{-i}, \theta),$$

where the equality follows from $g_{b_i} = 0$.

Assume now that $b_i \in \text{supp}_i(p)$. Choose l_i such that $p_{b_i} = p_{l_i}$. Since $p_{a_i} \neq p_{b_i}$, it must be that $m_i \neq l_i$. Suppose that $l_i > m_i$. It follows from the right-hand side of (17)—in (17) the roles of l_i and m_i are inverted—that $0 \geq \sum_{a_{-i}, \theta} f_{l_i}(a_{-i}, \theta) p_{m_i}(a_{-i}, \theta)$. Thus, given that $s_{l_i} > 0$, we deduce that

$$0 \geq \sum_{a_{-i}, \theta} g_{b_i}(a_{-i}, \theta) p_{a_i}(a_{-i}, \theta) = s_{l_i} \sum_{a_{-i}, \theta} f_{l_i}(a_{-i}, \theta) p_{m_i}(a_{-i}, \theta).$$

We obtain that $\sum_{a_{-i}, \theta} g_{a_i}(a_{-i}, \theta) p_{a_i}(a_{-i}, \theta) > 0 \geq \sum_{a_{-i}, \theta} g_{b_i}(a_{-i}, \theta) p_{a_i}(a_{-i}, \theta)$, where we use again (21). For the case $l_i < m_i$, the condition $\sum_{a_{-i}, \theta} g_{a_i}(a_{-i}, \theta) p_{a_i}(a_{-i}, \theta) > \sum_{a_{-i}, \theta} g_{b_i}(a_{-i}, \theta) p_{a_i}(a_{-i}, \theta)$ is equivalent to (19). We conclude that (20) holds. \square

To complete the proof of the lemma, for every $\delta > 0$ we define $u' = (u'_i)_{i \in I}$ by $u'_i(a, \theta) = u_i(a, \theta) + \delta g_{a_i}(a_{-i}, \theta)$. By choosing δ sufficiently small, we can be make sure that $\|u - u'\| \leq \epsilon$. Since $p \in BCE(u)$, it follows from (20) that for all $i \in I$ and $a_i \in \text{supp}_i(p)$, $a_i \in BR'(p_{a_i}) \subseteq \{b_i \in \text{supp}_i(p) : p_{a_i} = p_{b_i}\}$ where $BR'(p_{a_i})$ is the set of i 's best response to a belief p_{a_i} given utility function u'_i . Thus, $p \in sBCE(u')$. \square

A subset of a Euclidean space is **semi-algebraic** if it is defined by finite systems of polynomial inequalities. A correspondence between Euclidean spaces is semi-algebraic if its graph is semi-algebraic. The background knowledge on semi-algebraic sets that we use in this proof can be gathered from Blume and Zame (1994, Section 2).

To state the next result, let $\text{cl}(sBCE(u))$ be the closure of the sBCE set.

Lemma 6. *The correspondences $u \mapsto BCE(u)$, $u \mapsto sBCE(u)$, and $u \mapsto \text{cl}(sBCE(u))$ are semi-algebraic.*

Proof. The BCE correspondence is semi-algebraic: for all $u \in \mathbb{R}^{I \times A \times \Theta}$ and $p \in \mathbb{R}^{A \times \Theta}$, $p \in BCE(u)$ if and only if the pair (u, p) is a solution to the following finite system of polynomial inequalities: $p(a, \theta) \geq 0$ for all $a \in A$ and $\theta \in \Theta$; $\sum_a p(a, \theta) = \pi(\theta)$ for all $\theta \in \Theta$; and $\sum_{a_{-i}, \theta} (u(a, \theta) - u(b_i, a_{-i}, \theta)) p(a, \theta) \geq 0$ for all $i \in I$ and $a_i, b_i \in A_i$.

The sBCE correspondence is also semi-algebraic. To prove it, for $u \in \mathbb{R}^{I \times A \times \Theta}$, $p \in \mathbb{R}^{A \times \Theta}$, $i \in I$, and $a_i, b_i, c_i \in A_i$, we denote by $F(u, p, a_i, b_i, c_i)$ the quantity

$$\sum_{a_{-i}, \theta} (u(a_i, a_{-i}, \theta) - u(c_i, a_{-i}, \theta)) p(a_i, a_{-i}, \theta) + (u(b_i, a_{-i}, \theta) - u(c_i, a_{-i}, \theta)) p(b_i, a_{-i}, \theta).$$

We observe that $p \in sBCE(u)$ if and only if $p \in BCE(u)$ and for every $i \in I$ there is $T_i \subseteq A_i \times A_i$ such that the pair (u, p) is a solution of the following finite system of polynomial inequalities:

$$\sum_{a_{-i}, \theta} (p(a_i, a_{-i}, \theta)p(b_i) - p(b_i, a_{-i}, \theta)p(a_i))^2 = 0 \quad \text{for all } (a_i, b_i) \in T_i,$$

$$F(u, p, a_i, b_i, c_i) > 0 \quad \text{for all } (a_i, b_i) \notin T_i \text{ and } c_i \in A_i.$$

Thus, $p \in sBCE(u)$ if and only if it is the solution of one of finitely many systems of polynomial inequalities; we conclude that the sBCE correspondence is semi-algebraic.

The correspondence $u \mapsto \text{cl}(sBCE(u))$ is also semi-algebraic. Indeed, we have $p \in \text{cl}(sBCE(u))$ if and only if for every $\epsilon > 0$ there exists $q \in \mathbb{R}^{A \times \Theta}$ such that $\|p - q\| \leq \epsilon$ and $q \in sBCE(u)$. Thus, since the sBCE correspondence is semi-algebraic, the graph of the correspondence $u \mapsto \text{cl}(sBCE(u))$ is defined by a first-order formula and therefore semi-algebraic by the Tarski-Seidenberg theorem (Blume and Zame, 1994, page 787). \square

We are ready to complete the proof of the theorem. By Lemma 6, the correspondence $u \mapsto \text{cl}(sBCE(u))$ is semi-algebraic. Hence, there is an open subsets U of $\mathbb{R}^{I \times A \times \Theta}$ such that the complement of U has Lebesgue measure zero, and $u \mapsto \text{cl}(sBCE(u))$ is continuous on U (Blume and Zame, 1994, page 786).

We claim that for all $u \in U$, $BCE(u) = \text{cl}(sBCE(u))$. To prove the claim, take any $u \in U$. Since $BCE(u)$ is closed and $sBCE(u) \subseteq BCE(u)$, $\text{cl}(sBCE(u)) \subseteq BCE(u)$. To verify the other inclusion, we use Lemma 5 to find a sequence of games $(u^n)_{n=1}^\infty$ such that $u^n \rightarrow u$ and, for every n , $p \in sBCE(u^n) \subseteq \text{cl}(sBCE(u^n))$. Since $\text{cl}(sBCE(u))$ is continuous at u , we have $p \in \text{cl}(sBCE(u))$; see Aliprantis and Border (2006, Theorem 17.16). Hence, $BCE(u) \subseteq \text{cl}(sBCE(u))$. We deduce that $BCE(u) = \text{cl}(sBCE(u))$, as desired.

C. Proof of Theorem 3

The first part of the theorem immediately follows from Theorem 2, together with (3) and (4). To prove the theorem's last part, note that, when $|I| = 1$, $\min V_I = \min V_R = \bar{v}(p^0)$, where p^0 is a BCE that provides the player with no information (i.e., a BCE where the player's action does not vary with the state). The result then follows from

noting that for a generic game, both V_I and V_R are convex sets that share the same maximal value.

Next, we assume that $|I| \geq 2$, $|\Theta| \geq 2$, and $|A_i| \geq 2$ for at least two distinct player i , and show there is an open set $U \subseteq \mathbb{R}^{I \times A \times \Theta}$ such that $V_I(u) \subset V_R(u)$ for all $u \in U$. To ease the exposition, we actually assume (without loss of generality) that $|A_i| \geq 2$ for all players i .

To prove $V_I(u) \subset V_R(u)$, we will determine that

$$\min_{v \in V_R(u)} \sum_i v_i < \min_{v \in V_I(u)} \sum_i v_i. \quad (22)$$

Consistently with the notation of Section 6, we write $\underline{w}(p, u) = \sum_i v_i(p, u)$ and $\bar{w}(p, u) = \sum_i \bar{v}_i(p, u)$. We also denote by $\underline{w}(u)$ the minimum of $\underline{w}(p, u)$ over all $p \in BCE(u)$, and by $\bar{w}(u)$ the minimum of $\bar{w}(p, u)$ over all $p \in BCE(u)$. Note that

$$\min_{v \in V_R(u)} \sum_i v_i = \min_{p \in \text{cl}(sBCE(u))} \underline{w}(p, u) \geq \underline{w}(u), \quad \text{and} \quad \min_{v \in V_I(u)} \sum_i v_i = \bar{w}(u).$$

The next lemma gives a sufficient condition for the existence of an open set $U \subseteq \mathbb{R}^{I \times A \times \Theta}$ such that (22) holds for all $u \in U$.

Lemma 7. *Suppose $u^* \in \mathbb{R}^{I \times A \times \Theta}$ and $p^* \in \Delta_\pi(A \times \Theta)$ satisfy the following properties:*

- (i) *Each player i takes at least two actions at p^* : $|\text{supp}_i(p^*)| \geq 2$.*
- (ii) *p^* is a strict BCE: $BR(p_{a_i}) = \{a_i\}$ for all $i \in I$ and $a_i \in \text{supp}_i(p)$.*
- (iii) *p^* is the unique minimizer of $\underline{w}(p, u^*)$ over all $p \in BCE(u^*)$.*

Then there is a neighborhood U of u^ such that $\underline{w}(p, u) < \bar{w}(u)$ for all $u \in U$ and $p \in \text{cl}(sBCE(u))$.*

Proof. Take u^* and p^* that satisfy (i)-(iii). First, we verify that

$$\underline{w}(p^*, u^*) < \bar{w}(u^*). \quad (23)$$

Take $p \in BCE(u^*)$ such that $\bar{w}(u^*) = \bar{w}(p, u^*)$. If $p \neq p^*$, then $\bar{w}(p, u^*) \geq \underline{w}(p, u^*) > \underline{w}(p^*, u^*)$, where the strict inequality holds by (iii); thus, $\bar{w}(u^*) > \underline{w}(p^*, u^*)$. If instead $p = p^*$, then $\bar{w}(p^*, u^*) > \underline{w}(p^*, u^*)$ by (i) and (ii); thus $\bar{w}(u^*) > \underline{w}(p^*, u^*)$. Overall, we conclude that (23) holds, as desired.

The rest of the proof proceed by contradiction. To attain this contradiction, suppose there is a sequence $(u^n)_{n=1}^\infty$ converging to u^* such that

$$\min_{p \in \text{cl}(sBCE(u^n))} \underline{w}(p, u^n) = \bar{w}(u^n) \quad \text{for all } n. \quad (24)$$

By (ii), $p^* \in sBCE(u^n)$ for all n sufficiently large. Thus,

$$\underline{w}(p^*, u^n) \geq \min_{p \in \text{cl}(sBCE(u^n))} \underline{w}(p, u^n) \quad \text{for all } n \text{ large enough.} \quad (25)$$

Combining (24) and (25), we obtain that

$$\underline{w}(p^*, u^n) \geq \bar{w}(u^n) \quad \text{for all } n \text{ large enough.} \quad (26)$$

By standard arguments, $\underline{w}(p^*, u)$ is continuous in u . In addition, since the correspondence $u \mapsto BCE(u)$ is upper hemicontinuous, $\bar{w}(u) = \min_{p \in BCE(u)} \bar{w}(p, u)$ is lower semicontinuous in u (e.g., Aliprantis and Border, 2006, Lemma 17.3). It follows from (26) that $\underline{w}(p^*, u^*) = \liminf_{n \rightarrow \infty} \underline{w}(p^*, u^n) \geq \liminf_{n \rightarrow \infty} \bar{w}(u^n) \geq \bar{w}(u^*)$. Hence, $\underline{w}(p^*, u^*) \geq \bar{w}(u^*)$, which contradicts (23). \square

To complete the proof of the theorem, we construct a utility profile u^* and an outcome p^* that satisfy the conditions of Lemma 7. This lemma then delivers a neighborhood U of u^* such that (22) holds for all $u \in U$, which in turn means that the set of u such that $V_I(u) \subset V_R(u)$ has non-empty interior.

We now construct u^* and p^* . Let n be cardinality of I ; by hypothesis, $n \geq 2$. For every player i , we order the set of actions from 0 to m_i (where $m_i + 1$ is the cardinality of A_i): $A_i = \{0, \dots, m_i\}$. By hypothesis, A_i contains at least two distinct elements, thus $m_i \geq 1$. We also consider a partition $\Theta = \Theta_l \cup \Theta_h$ of the set of payoff states such that both Θ_l and Θ_h are nonempty; this is feasible because, by hypothesis, Θ contains at least two elements.

For player i , we define u_i^* as follows:

$$u_i^*(a, \theta) = \begin{cases} 0 & \text{if } a_i = 0, \\ \frac{1}{\pi(\Theta_l)} \left(-1 + \frac{1}{n-1} \sum_{j \neq i} a_j \right) & \text{if } a_i = 1 \text{ and } \theta \in \Theta_l, \\ \frac{1}{\pi(\Theta_h)} \left(2 - \frac{1}{n-1} \sum_{j \neq i} a_j \right) & \text{if } a_i = 1 \text{ and } \theta \in \Theta_h, \\ -1 & \text{if } a_i > 1. \end{cases}$$

Thus, action 0 is a safe action. Action 1 has a payoff that depends both on the state and on the average action of the opponents. For states in Θ_l , action 1 generates a *negative* baseline payoff of -1 , but there is also a *positive* externality from the actions of others; these payoffs are scaled by $1/\pi(\Theta_l)$. For states in Θ_h , action 1 generates a *positive* baseline payoff of 2 , but there is also a *negative* externality from the actions of others; these payoffs are scaled by $1/\pi(\Theta_h)$. Any action outside $\{0, 1\}$ is strictly dominated by 0.

Let p^* be the outcome such that all players take action 0 when $\theta \in \Theta_l$, and all players take action 1 when $\theta \in \Theta_h$. Clearly, (i) and (ii) of Lemma 7 hold.

All that remains is to verify (iii). To do so, note first that p^* is the unique minimizer of $\sum_i \sum_{a,\theta} u_i^*(1, a_{-i}, \theta)p(a, \theta)$ over all $p \in \Delta_\pi(\{0, 1\}^I \times \Theta)$. Since any action outside $\{0, 1\}$ is strictly dominated, we deduce that p^* is the unique minimizer of $\sum_i \sum_{a,\theta} u_i^*(1, a_{-i}, \theta)p(a, \theta)$ over all $p \in BCE(u^*)$. Moreover, simple algebra shows that for all players i ,

$$\max_{b_i} \sum_{a,\theta} u_i^*(b_i, a_{-i}, \theta)p^*(a, \theta) = \sum_{a,\theta} u_i^*(1, a_{-i}, \theta)p^*(a, \theta).$$

In turn, this implies that $\underline{w}(p^*, u^*) = \sum_i \sum_{a,\theta} u_i^*(1, a_{-i}, \theta)p^*(a, \theta)$. Therefore, every $p \in BCE(u^*) \setminus \{p^*\}$ has

$$\begin{aligned} \underline{w}(p, u^*) &= \sum_i \max_{b_i} \sum_{a,\theta} u_i^*(b_i, a_{-i}, \theta)p(a, \theta) \geq \sum_i \sum_{a,\theta} u_i^*(1, a_{-i}, \theta)p(a, \theta) \\ &> \sum_i \sum_{a,\theta} u_i^*(1, a_{-i}, \theta)p^*(a, \theta) = \underline{w}(p^*, u^*). \end{aligned}$$

we conclude that $\underline{w}(u^*) = \underline{w}(p^*, u^*)$ if and only if $p = p^*$, that is, (iii) of Lemma 7 holds. The proof is now complete.

References

- C. D. Aliprantis and K. C. Border. *Infinite Dimensional Analysis*. Springer, third edition, 2006.
- G.-M. Angeletos and K. A. Sastry. Inattentive economies. 2023. URL <https://www.karthiksastry.com/>.

- R. J. Aumann. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1:67–96, 1974.
- D. Bergemann and S. Morris. Robust predictions in games with incomplete information. *Econometrica*, 81:1251–1308, 2013.
- D. Bergemann and S. Morris. Bayes correlated equilibrium and the comparison of information structures in games. *Theoretical Economics*, 11:487–522, 2016.
- D. Bergemann, B. Brooks, and S. Morris. The limits of price discrimination. *American Economic Review*, 105:921–957, 2015.
- D. Bergemann, B. Brooks, and S. Morris. First-price auctions with general information structures: Implications for bidding and revenue. *Econometrica*, 85:107–143, 2017.
- D. Blackwell. Comparison of experiments. In *Proceedings of the 2nd Berkeley Symposium on Mathematical Statistics and Probability*, pages 93–102. University of California Press, 1951.
- D. Blackwell. Equivalent comparisons of experiments. *Annals of Mathematical Statistics*, 24:265–272, 1953.
- L. E. Blume and W. R. Zame. The algebraic geometry of perfect and sequential equilibrium. *Econometrica*, 62:783–794, 1994.
- B. Brooks and S. Du. Optimal auction design with common values: An informationally robust approach. *Econometrica*, 89:1313–1360, 2021.
- A. Caplin and M. Dean. Revealed preference, rational inattention, and costly information acquisition. *American Economic Review*, 105:2183–2203, 2015.
- A. Caplin, M. Dean, and J. Leahy. Rationally inattentive behavior: Characterizing and generalizing Shannon entropy. *Journal of Political Economy*, 130:1676–1715, 2022.
- G. Carroll. Robustness in mechanism design and contracting. *Annual Review of Economics*, 11:139–166, 2019.

- L. Colombo, G. Femminis, and A. Pavan. Information acquisition and welfare. *Review of Economic Studies*, 81:1438–1483, 2014.
- G. de Clippel and K. Rozen. Communication, perception, and strategic obfuscation. 2022. URL <https://sites.google.com/brown.edu/kareenrozen>.
- M. Dean and N. L. Neligh. Experimental tests of rational inattention. *Journal of Political Economy*, forthcoming.
- T. Denti. Costly monitoring in signaling games. 2021. URL <https://sites.google.com/site/tommasojdenti>.
- T. Denti. Posterior separable cost of information. *American Economic Review*, 112:3215–59, 2022.
- T. Denti. Unrestricted information acquisition. *Theoretical Economics*, forthcoming.
- T. Denti, M. Marinacci, and A. Rustichini. Experimental cost of information. *American Economic Review*, 112:3106–23, 2022.
- L. Doval and J. Ely. Sequential information design. *Econometrica*, 88:2575–2608, 2020.
- F. Forges. Five legitimate definitions of correlated equilibrium in games with incomplete information. *Theory and Decision*, 35:277–310, 1993.
- B. Hebert and J. La’O. Information acquisition, efficiency, and non-fundamental volatility. *Journal of Political Economy*, forthcoming.
- C. Hellwig and L. Veldkamp. Knowing what others know: Coordination motives in information acquisition. *Review of Economic Studies*, 76:223–251, 2009.
- B. L. Lipman and S. Srivastava. Computation as a correlation device. *Games and Economic Behavior*, 2:154–172, 1990.
- E. Lipnowski and D. Ravid. Predicting choice from information costs. 2022. URL <https://arxiv.org/abs/2205.10434>.
- B. Maćkowiak and M. Wiederholt. Optimal sticky prices under rational inattention. *American Economic Review*, 99:769–803, 2009.

- B. Maćkowiak, F. Matějka, and M. Wiederholt. Rational inattention: A review. *Journal of Economic Literature*, 61:226–273, 2023.
- F. Matějka and A. McKay. Rational inattention to discrete choices: A new foundation for the multinomial logit model. *American Economic Review*, 105:272–298, 2015.
- S. Morris and H. S. Shin. Global games: theory and applications. In *Proceedings of the 8th World Congress of the Econometric Society*, pages 56–114. Cambridge University Press, 2003.
- S. Morris and M. Yang. Coordination and continuous stochastic choice. *The Review of Economic Studies*, 89:2687–2722, 2022.
- D. P. Myatt and C. Wallace. Endogenous information acquisition in coordination games. *Review of Economic Studies*, 79:340–374, 2012.
- R. B. Myerson. Optimal coordination mechanisms in generalized principal–agent problems. *Journal of Mathematical Economics*, 10:67–81, 1982.
- R. B. Myerson. Dual reduction and elementary games. *Games and Economic Behavior*, 21:183–202, 1997.
- D. Ravid. Ultimatum Bargaining with Rational Inattention. *American Economic Review*, 110:2948–63, 2020.
- D. Ravid, A.-K. Roesler, and B. Szentes. Learning before trading: on the inefficiency of ignoring free information. *Journal of Political Economy*, 130:346–387, 2022.
- R. T. Rockafellar. *Convex analysis*. Princeton University Press, 1970.
- T. Roughgarden. *Twenty lectures on algorithmic game theory*. Cambridge University Press, 2016.
- C. A. Sims. Implications of rational inattention. *Journal of Monetary Economics*, 50:665–690, 2003.
- M. Yang. Coordination with flexible information acquisition. *Journal of Economic Theory*, 158:721–738, 2015.
- W. Zhong. Optimal dynamic information acquisition. *Econometrica*, 90:1537–1582, 2022.