

Robust Predictions in Games with Incomplete Information*

Dirk Bergemann[†]

Stephen Morris[‡]

May 24, 2011

Abstract

We analyze games of incomplete information and offer equilibrium predictions which are valid for all possible private information structures that the agents may have. The predictions about the joint equilibrium distributions which are robust to the specification of the private information structure, rely on an epistemic result which establishes a relationship between the set of Bayes correlated and Bayes Nash equilibria.

We completely characterize the set of equilibria in a class of games with quadratic payoffs in terms of restrictions on the first and second moments of the equilibrium action-state distribution.

Finally, we reverse the perspective and investigate the identification problem under concerns for robustness to private information. We show how the presence of private information leads to partial rather than complete identification of the structural parameters of the game.

JEL CLASSIFICATION: C72, C73, D43, D83.

KEYWORDS: Incomplete Information, Correlated Equilibrium, Robustness to Private Information, Moments Restrictions, Identification, Informations Bounds.

*We acknowledge financial support through NSF Grant SES 0851200. We benefitted from comments of Steve Berry, Vincent Crawford, Phil Haile, Marc Henry, Arthur Lewbel, Larry Samuelson, and Elie Tamer, and research assistance from Brian Baisa. We would like to thank seminar audiences at Boston College, the Collegio Carlo Alberto, Ecole Polytechnique, European University Institute, HEC, Microsoft Research, Northwestern University, the Paris School of Economics, Stanford University and the University of Colorado for stimulating conversations; and we thank David McAdams for his discussion of this work at the 2011 North American Winter Meetings of the Econometric Society in Denver.

[†]Department of Economics, Yale University, New Haven, CT 06520, U.S.A., dirk.bergemann@yale.edu.

[‡]Department of Economics, Princeton University, Princeton, NJ 08544, U.S.A. smorris@princeton.edu

1 Introduction

In games of incomplete information, the private information of each agent typically induces posterior beliefs about some payoff relevant state, and a posterior belief about the beliefs of the other agents. In turn, the private information of the agent, the type in the language of Bayesian games, influences the optimal strategies of the agents, and ultimately the equilibrium distribution over actions and states. The posterior belief about the payoff relevant state represents the knowledge about the *payoff environment* that the player is facing, whereas the posterior belief about the beliefs of the other agents represents the knowledge about the *belief environment* that the player is facing. The objective of this paper is to obtain equilibrium predictions for a given payoff environment which are independent of - and in that sense robust to - the specification of the belief environment.

We define the payoff environment as the complete description of the agents' preferences and common prior over payoff relevant variables. We define the belief environment as the complete description of the common prior type space over and above the information contained in the common prior distribution of the payoff relevant variable. The belief environment then describes a potentially rich type space which is only subject to the constraint that the marginal distribution over the fundamental variable coincides with the common prior over payoff relevant variables. A pair of payoff environment and belief environment form a standard Bayesian game. Yet, for a given payoff environment, there are many belief environments, and each distinct belief environment may lead to distinct equilibrium distribution over outcomes, namely actions and fundamentals.

The objective of the paper is to describe the equilibrium implications of the "payoff environment" for all possible "belief environments" relative to the given payoff environment. Consequently, we refer to the (partial) characterization of the equilibrium outcomes that are independent of the belief environment as robust predictions. We examine these issues in a tractable class of games with a continuum of players, symmetric payoff functions, and linear best response functions. A possible route towards a comprehensive description of the equilibrium implications stemming from the payoff environment alone, would be an exhaustive analysis of all Bayes Nash equilibria of all belief environments which are associated with a given payoff environment. Here we shall not pursue this direct approach. Instead we shall use a related equilibrium notion, namely the notion of Bayes correlated equilibrium to obtain a comprehensive characterization. We begin with an epistemic result that establishes the equivalence between the class of Bayes Nash equilibrium distribution for all possible belief environments and the class of Bayes correlated equilibrium. This result is a natural extension of a seminal result by Aumann (1987). In games with complete information about the payoff environment, he establishes the equivalence between the set of Bayes Nash equilibria and the correlated equilibria. We present the epistemic result for the class of games with a continuum of agent and symmetric payoff functions, and show that the insights of Aumann (1987) generalizes naturally to this class of games with incomplete information. In a companion paper, Bergemann and Morris (2011), we report the definition of Bayes correlated equilibrium and the relevant epistemic results in general finite player and finite action incomplete information games and how they relate to the results of Forges (1993).

Subsequently we use the epistemic result to provide a complete characterization of the Bayes correlated equilibria in the class of games with quadratic payoffs. With quadratic games, the best response function of each agent is a linear function and in consequence the conditional expectations of the agents are linked through linear conditions

which in turn permits an explicit construction of the equilibrium sets. The class of quadratic games has featured prominently in many recent contributions to games of incomplete information, for example the analysis of rational expectations in competitive markets by Guesnerie (1992), the analysis of the beauty contest by Morris and Shin (2002) and the equilibrium use of information by Angeletos and Pavan (2007). We offer a characterization of the equilibrium outcomes in terms of the moments of the equilibrium distributions. In the class of quadratic games, we show that the expected mean is constant across all equilibria and provide sharp inequalities on the variance-covariance of the joint outcome state distributions. If the underlying uncertainty about the payoff relevant variable and the equilibrium distribution itself is normally distributed then the characterization of the equilibrium is completely given by the first and second moment. If the distribution of uncertainty or the equilibrium distribution itself is not normally distributed, then the characterization of first and second moments remains valid, but of course is not a complete characterization in the sense that the determination of the higher moments is incomplete.

With the given characterization of the Bayes correlated equilibrium, we then ask what information structures, or type spaces, would generate in a Bayes Nash equilibrium a specific Bayes correlated equilibrium. We show that a bivariate information structure, consisting of a public signal and a private signal about the payoff variable, is sufficient to generate the entire set of Bayes correlated equilibria.

The initial equivalence result between Bayes correlated and Bayes Nash equilibrium relied on very weak assumptions about the belief environment of the agents. In particular, we allowed for the possibility that the agents may have very little or in fact zero additional information beyond the common prior about the payoff variable. Yet, in some circumstances the agents may be commonly known to have some given prior information, or background information. Consequently, we then analyze how a lower bound on either the public or the private information of the agents, can be used to further refine the robust predictions and impose additional moment restrictions on the equilibrium distribution.

The payoff environment is specified by the (ex-post) observable outcomes, the actions and the payoff variable. By contrast, the elements of the belief environment, the beliefs of the agents, the beliefs over the beliefs of the agents, etc. are rarely directly observed or inferred from the revealed choices of the agents. The absence of the observability (via revealed preference) of the belief environment then constitutes a separate reason to be skeptical towards an analysis which relies on very specific and detailed assumptions about the belief environment. (In separate work, Bergemann, Morris, and Takahashi (2010) ask what can be learned about agents' possibly interdependent preferences by observing how they behave in strategic environments. As they are interested in identifying when two types are strategically distinguishable in the sense that they are guaranteed to behave differently in some finite game, their framework is different from the current one, as here we consider a given game rather than quantifying over all games.)

Finally, we reverse the perspective of our analysis and consider the issue of identification rather than the issue of prediction. In other words, we are asking whether the observable data, namely actions and payoff-relevant variable, can identify the structural parameters of the payoff functions, and thus of the game, without stringent assumptions on the belief environment. The question of identification is to ask whether the observable data imposes restrictions on the unobservable structural parameters of the game given the equilibrium hypothesis. Similarly to the problem of *robust* equilibrium prediction, the question of *robust* identification then is whether the restriction are common to

all possible belief environments given a specific payoff environment. In the context of the quadratic payoffs that we study, we find that the robust identification does allow us to identify the sign of some interaction parameters, but will leave other parameters, in particular whether the agents are playing a game of strategic substitutes or complements, unidentified, even in terms of the sign of the interaction. The identification results here, in particular the contrast between Bayes Nash equilibrium and Bayes correlated equilibrium, are related to, but distinct from the results presented in Aradillas-Lopez and Tamer (2008). In their analysis of an entry game with incomplete information, they document the loss in identification power that arises with a more permissive solution concept, i.e. level k -rationalizability. As we compare Bayes Nash and correlated equilibrium, we show that the lack of identification is not necessarily due to the lack of a common prior, as associated with rationalizability, but rather the richness of the possible private information structures (but all with a common prior).

In recent years, the concern for a robust equilibrium analysis in games of incomplete information has been articulated in many ways. In mechanism design, where the rules of the games can be chosen to have favorable robustness properties, a number of positive results have been obtained. Dasgupta and Maskin (2000), Bergemann and Välimäki (2002) and Bergemann and Morris (2005), among others, show that the efficient social allocation can be implemented in an ex-post equilibrium and hence in Bayes Nash equilibrium for all type spaces, with or without a common prior.¹ But in “given” rather than “designed” games, such strong robustness results seem out of reach for most classes of games. In particular, many Bayesian games simply do not have ex post or dominant strategy equilibria. In the absence of such global robustness results, a natural first step is then to investigate the robustness of the Bayes Nash equilibrium to a small perturbation of the information structure. For example, Kajii and Morris (1997) consider a Nash equilibrium of a complete information game and say that the Nash equilibrium is robust to incomplete information if every incomplete information game with payoffs almost always given by the complete information game has an equilibrium which generates behavior close to the Nash equilibrium. In this paper, we take a different approach and use the dichotomy between the payoff environment and the belief environment to analyze the equilibrium behavior in a given payoff environment while allowing for any arbitrary, but common prior, type space, as long as it is consistent with the given common prior of the payoff type space. In the current contribution, we trace out the Bayes Nash equilibria associated with all possible information structures. A related literature seeks to identify the best information structure consistent with the given common prior over payoff types. For example, Bergemann and Pesendorfer (2007) characterizes the revenue-maximizing information structure in an auction with many bidders. Similarly, in a class of sender-receiver games, Kamenica and Gentzkow (2010) derive the sender-optimal information structure.

Chwe (2006) discusses the role of statistical information in single-agent and multi-agent decision problems. In a series of related settings, he argues that the correlation between the revealed choice of an agent, referred to as incentive compatibility, and a random variable, not controlled by the agent, allows us an analyst to infer the nature of the payoff interaction between the agent’s choice and the random variable. For example, in 2×2 games, observing the signed covariance in the actions is sufficient to identify pure strategy Nash equilibria. In the context of a two player game with quadratic payoffs and complete information, he shows that the sign of the covariance of the agents’

¹Jehiel and Moldovanu (2001) and Jehiel, Moldovanu, Meyer-Ter-Vehn, and Zame (2006) demonstrate the limits of these results by considering multi-dimensional payoff types.

actions can be predicted across all correlated equilibria.

The remainder of the paper is organized as follows. Section 2 defines the relevant solution concepts and establishes the epistemic result which relates the set of Bayes Nash equilibria to the set of Bayes correlated equilibria. Beginning with Section 3, we confine our attention to a class of quadratic games with normally distributed uncertainty about the payoff relevant state. Section 4 reviews the standard approach to games with incomplete information and analyses the Bayes Nash equilibria under a bivariate belief environment in which each agent receives a private and a public signal about the payoff relevant state. Section 5 begins with the analysis of the Bayes correlated equilibrium and we present a complete description of the equilibrium set in terms of moment restrictions on the joint equilibrium distribution. We then establish the link between the set of Bayes correlated equilibria and the set Bayes Nash equilibrium under the bivariate belief environment. In Section 6 we analyze how prior information about the belief environment can further restrict the equilibrium predictions. In Section 7, we turn from prediction to the issue of identification. We ask how much we can learn from the observable actions and payoff relevant states about the structural parameters of the game. Here we consider both the case of observable individual actions as well as observable aggregate actions. With only aggregate actions observable, we consider the possibility of robust identification within the context of the classic demand and supply identification problem. Section 8 discusses some possible extensions and offers concluding remarks. The Appendix collects some of the proofs from the main body of the text.

2 Set-Up

We will define the solution concept of Bayes correlated equilibrium and relate it to robust predictions in a class of continuum player symmetric payoffs games, and then apply it to the linear best response normal games studied in the main analysis of this paper. In a companion paper, Bergemann and Morris (2011), we first develop the solution concept and relation to robust predictions in finite player, finite action games with asymmetric payoffs and discuss how the results can be adapted first to symmetric payoffs and then to continua of players and actions.

Payoff Environment There is a continuum of players and an individual player is indexed by $i \in [0, 1]$. Each player chooses an action $a \in \mathbb{R}$. There will then be a realized population action distribution $h \in \Delta(\mathbb{R})$. There is a payoff relevant state $\theta \in \Theta$. All players have the same payoff function $u : \mathbb{R} \times \Delta(\mathbb{R}) \times \Theta \rightarrow \mathbb{R}$, where $u(a, h, \theta)$ is a player's payoff if she chooses action a , the population action distribution is h and the state is θ . There is a prior distribution $\psi \in \Delta(\Theta)$. A payoff environment is thus parameterized by (u, ψ) . We may also refer to (u, ψ) as the "basic game" as $\psi \in \Delta(\Theta)$ only specifies the common prior distribution over the payoff relevant state $\theta \in \Theta$ whereas it does not specify the private information the agents may have access to.

Belief Environment Each player will observe a signal (or realize a type) $t \in T$. In each state of the world $\theta \in \Theta$, there will be a realized distribution of signals $g \in \Delta(T)$ drawn according to a distribution $k \in \Delta(\Delta(T))$. Let $\pi : \Theta \rightarrow \Delta(\Delta(T))$ give the distribution over signal distributions. Thus the belief environment, or alternatively an information system, is parameterized by (T, π) .

Bayes Correlated Equilibrium We will be interested in probability distributions $\mu \in \Delta(\Delta(\mathbb{R}) \times \Theta)$ with the interpretation that μ is the joint distribution of the population action distribution h and the state θ . For any such μ , we write $\hat{\mu}$ for the induced probability distribution on $\mathbb{R} \times \Delta(\mathbb{R}) \times \Theta$ if $(h, \theta) \in \Delta(\mathbb{R}) \times \Theta$ are drawn according to μ and there is then a conditionally independent draw of $a \in \mathbb{R}$ according to h . For each $a \in \mathbb{R}$, we write $\hat{\mu}(\cdot|a)$ for the probability on $\Delta(\mathbb{R}) \times \Theta$ conditional on a (we will write as if it is uniquely defined).

Definition 1 (Bayes Correlated Equilibrium)

A probability distribution $\mu \in \Delta(\Delta(\mathbb{R}) \times \Theta)$ is a Bayes correlated equilibrium (BCE) of (u, ψ) if

$$\mathbb{E}_{\hat{\mu}(\cdot|a)} u(a, h, \theta) \geq \mathbb{E}_{\hat{\mu}(\cdot|a')} u(a', h, \theta)$$

for each a and a' ; and

$$\text{marg}_{\Theta} \mu = \psi.$$

The current notion of Bayes correlated equilibrium is closely related to the notion of a “universal Bayesian approach”, as defined by Forges (1993). Forges (1993) is titled and identifies “five legitimate definitions of correlated equilibrium in games with incomplete information” and Forges (2006) describes a mistake in Forges (1993) that leads to a sixth definition. There are however a few differences between our formulation of games and solution concepts and those of Forges: (i) we directly define solution concepts as subsets of action type state distributions, while she characterizes the set of equilibrium payoffs satisfying a set of restrictions which implicitly define the solution concept; (ii) we work with a payoff environment”, $G = \left((A_i, u_i)_{i=1}^I, \psi \right)$, describing prior and payoffs and an “information environment” $S = \left((T_i)_{i=1}^I, \pi \right)$, while she distinguishes between the “decision problem with incomplete information,” $(A_i, u_i)_{i=1}^I$ and includes the prior on payoff relevant states in her description of the “information scheme”; and (iii) we include the distribution of payoff relevant states Θ in our solution concept, while she integrates out payoff relevant states. Thus, in her notion of “universal Bayesian approach” a prior “information scheme” (in our language, prior on Θ and belief environment) is not taken as given. But if one looks at the joint distributions on actions and states consistent with the assumptions that she makes in her universal Bayesian approach, they are - in our language - Bayes correlated equilibria of the basic game G (i.e., the Bayesian game with the null information system). To be more precise, if one fixes $(A_i, u_i)_{i=1}^I$, and looks at the vector of ex ante payoffs the players might receive in the Bayes Correlated Equilibria of the basic game $G = \left((A_i, u_i)_{i=1}^I, \psi \right)$, for any $\psi \in \Delta(\Theta)$, one gets Forges’ universal Bayesian solutions.

In our definition of Bayes correlated equilibrium, the types T_i are implicit in the sense that the probability distribution μ will induce a belief over actions and beliefs of the other players. Thus, our definition extends the notion of a correlated equilibrium in Aumann (1987) to an environment with uncertain payoffs, represented by the state of the world θ . We introduce the language of types later on when we consider games in which the players are known to have private information about the state of world, which is encoded in the types.

In recent work, Lehrer, Rosenberg, and Shmaya (2006), present conditions under which one information structure leads to the same payoff set than another information structure with respect to a number of game-theoretic solution concepts, including the Bayes correlated equilibrium. Their characterization uses extensions of the garbling condition

of Blackwell (1953) to many agents. In Lehrer, Rosenberg, and Shmaya (2010), the analysis is refined to games of common interest and necessary and sufficient conditions are presented that characterize when one information structure generates more favorable outcomes than another information structure. They show that an information structure \mathcal{I} is better than information structure \mathcal{I}' , if \mathcal{I}' is a some kind of garbling of \mathcal{I} , where the exact nature of the garbling condition depends on the solution concept under consideration. Gossner (2000) compares information structures by analyzing the correlated equilibrium distributions they can induce. In his construction, the actions of each agent have to be measurable with respect to the signals that each agent receives, in other words the actions cannot display any correlation beyond the correlation contained in the information structure (but of course they can contain less information in some correlated equilibrium). We discuss the relationships between these different equilibrium notions and the work of ?? in detail in a companion paper, Bergemann and Morris (2011). In particular, we discuss the impact of private information on the size of the equilibrium set.

Bayes Nash Equilibrium The payoff environment (u, ψ) and the belief environment (T, π) together define a game of incomplete information $((u, \psi), (T, \pi))$. A strategy in the game is then defined by $\sigma : T \rightarrow \Delta(\mathbb{R})$. The interpretation is that $\sigma(t)$ is the realized distribution of actions among those players observing signal t (i.e., we are "assuming the law of large numbers" on the continuum). A distribution of signals $g \in \Delta(T)$ and $\sigma \in \Sigma$ induce a probability distribution $g \circ \sigma \in \Delta(\mathbb{R})$. The prior $\psi \in \Delta(\Theta)$ and signal distribution $\pi : \Theta \rightarrow \Delta(T)$ induce a probability distribution $\psi \circ \pi \in \Delta(\Delta(T) \times \Theta)$. As before, write $\widehat{\psi \circ \pi}$ for the probability distribution on $T \times \Delta(T) \times \Theta$ if $(g, \theta) \in \Delta(T) \times \Theta$ are drawn according to $\psi \circ \pi$ and there is then a conditionally independent draw of $t \in T$ according to the realized $g \in \Delta(T)$. For each $t \in T$, we write $\widehat{\psi \circ \pi}(\cdot | t)$ for the probability on $\Delta(T) \times \Theta$ conditional on t (we will write as if it is uniquely defined).

Definition 2 (Bayes Nash Equilibrium)

A symmetric strategy $\sigma \in \Sigma$ is a Bayes Nash equilibrium (BNE) of $((u, \psi), (T, \pi))$ if

$$\mathbb{E}_{\widehat{\psi \circ \pi}(\cdot | t)} u(a, g \circ \sigma, \theta) \geq \mathbb{E}_{\widehat{\psi \circ \pi}(\cdot | t)} u(a', g \circ \sigma, \theta),$$

for all $t \in T$, a in the support of $\sigma(\cdot | t)$ and $a' \in \mathbb{R}$.

Let $\psi \circ \pi \circ \sigma$ be the probability distribution on $\Delta(\Delta(\mathbb{R}) \times \Theta)$ induced if $(g, \theta) \in \Delta(T) \times \Theta$ are drawn according to $\psi \circ \pi$ and $h \in \Delta(\mathbb{R})$ is set equal to $g \circ \sigma$.

Definition 3 (Bayes Nash Equilibrium Distribution)

A probability distribution $\mu \in \Delta(\Delta(\mathbb{R}) \times \Theta)$ is a BNE action state distribution of $((u, \psi), (T, \pi))$ if there exists a BNE σ of $((u, \psi), (T, \pi))$ such that $\mu = \psi \circ \pi \circ \sigma$.

Epistemic Result We are now in a position to relate the Bayes correlated equilibria with the Bayes Nash equilibria.

Proposition 1

A probability distribution $\mu \in \Delta(\Delta(\mathbb{R}) \times \Theta)$ is a Bayes correlated equilibrium of (u, ψ) if and only if it is a BNE action state distribution $((u, \psi), (T, \pi))$ for some information system (T, π) .

Proof. Suppose that μ is a BCE of (u, ψ) . Let $T = \mathbb{R}$, let $\pi : \Theta \rightarrow \Delta(T)$ be set equal to the conditional probability $\mu : \Theta \rightarrow \Delta(\mathbb{R})$ and let σ be the "truth-telling" strategy with type a choosing action a with probability 1. Now

$$\mathbb{E}_{\widehat{\psi \circ \pi}(\cdot|a)} u(a', g \circ \sigma, \theta) = \mathbb{E}_{\widehat{\mu}(\cdot|a)} u(a', h, \theta)$$

by construction and the BCE equilibrium conditions imply the BNE equilibrium conditions.

Suppose that σ is a BNE of $((u, \psi), (T, \pi))$ and so

$$\mathbb{E}_{\widehat{\psi \circ \pi}(\cdot|t)} u(a, g \circ \sigma, \theta) \geq \mathbb{E}_{\widehat{\psi \circ \pi}(\cdot|t)} u(a', g \circ \sigma, \theta) \quad (1)$$

for all $t \in T$, a in the support of $\sigma(\cdot|t)$ and $a' \in \mathbb{R}$. Now $\mathbb{E}_{\widehat{\psi \circ \pi}(\cdot|t)} u(a', g \circ \sigma, \theta)$ is a function of t . The expectation of this expectation conditional on a being drawn under strategy σ is

$$\mathbb{E}_{\widehat{\psi \circ \pi \circ \sigma}(\cdot|a)} u(a', g \circ \sigma, \theta)$$

and thus taking the expectation of both sides of (1) establishes that $\psi \circ \pi \circ \sigma$ is a BCE. ■

Aumann (1987) establishes the relation between Nash equilibria and correlated equilibria in games with complete information. In the companion paper, Bergemann and Morris (2011), we establish the relevant epistemic results for canonical game theoretic environments in more detail.

3 Quadratic Payoff Environment

For the remainder of this paper, we consider a quadratic and symmetric model of interaction. There is a continuum of agents, $i \in [0, 1]$. The individual action is denoted by $a_i \in \mathbb{R}$ and the average action is denoted by $A \in \mathbb{R}$:

$$A \triangleq \int_i a_i di.$$

The payoff of agent i is denoted by $u_i(a_i, A, \theta)$ and depends on the individual action a_i , the average action A and a state of the world $\theta \in \mathbb{R}$. The payoffs are quadratic and symmetric across agents and given by:

$$u_i(a_i, A, \theta) \triangleq \begin{pmatrix} \lambda_a \\ \lambda_A \\ \lambda_\theta \end{pmatrix}' \begin{pmatrix} a_i \\ A \\ \theta \end{pmatrix} + \begin{pmatrix} a_i \\ A \\ \theta \end{pmatrix}' \begin{pmatrix} \gamma_a & \gamma_{aA} & \gamma_{a\theta} \\ \gamma_{aA} & \gamma_A & \gamma_{A\theta} \\ \gamma_{a\theta} & \gamma_{A\theta} & \gamma_\theta \end{pmatrix} \begin{pmatrix} a_i \\ A \\ \theta \end{pmatrix}. \quad (2)$$

The vector $\lambda = (\lambda_a, \lambda_A, \lambda_\theta)$ represents the linear returns and matrix $\Gamma = \{\gamma_{ij}\}$ represents the interaction structure of the game, together λ and Γ completely describes the payoffs of the agents. The entries in the interaction matrix Γ are uniformly denoted by γ . The interaction parameters with a single subscript, namely $\gamma_a, \gamma_A, \gamma_\theta$, are the diagonal entries in the interaction matrix. We assume that the payoffs are concave in the own action:

$$\gamma_a < 0, \quad (3)$$

and that the interaction of the individual action and the average action (the "indirect effect") is bounded by the own action (the "direct effect"):

$$-\gamma_{aA}/\gamma_a < 1 \Leftrightarrow \gamma_a + \gamma_{aA} < 0. \quad (4)$$

The concavity and the moderate interaction jointly imply that the complete information game has a unique and bounded Nash equilibrium. We assume that $\gamma_{a\theta} \neq 0$ to have the fundamental matter. The game displays strategic complementarity if $\gamma_{aA} > 0$ and strategic substitutes if $\gamma_{aA} < 0$.

The state of the world θ is distributed normally with

$$\theta \sim N(\mu_\theta, \sigma_\theta^2). \quad (5)$$

The quadratic environment encompasses a wide class of interesting economic environment. The following two applications are prominent examples and we shall return to them throughout the paper to illustrate some of the results.

Example 1 (Beauty Contest) *In Morris and Shin (2002), a continuum of agents, $i \in [0, 1]$, have to choose an action under incomplete information about the state of the world θ . Each agent i has a payoff function given by:*

$$u_i(a_i, A, \theta) = -(1-r)(a_i - \theta)^2 - r(a_i - A)^2.$$

The weight r reflects concern for the average action taken A in the population. Morris and Shin (2002) then analyze the Bayes Nash equilibrium in which each agent i has access to a private (idiosyncratic) signal and a public (common) signal of the world. In terms of our model, the beauty contest model set $\gamma_a = -1$, $\gamma_{aA} = r$ and $\gamma_{a\theta} = (1-r)$. Angeletos and Pavan (2007) generalize the analysis of Bayes Nash equilibrium under this bivariate information structure for the general class of quadratic environments defined above by (2).

Example 2 (Competitive Market) *Guesnerie (1992) offer an analysis of the stability of the competitive equilibrium by considering a continuum of producers with a quadratic cost of production and a linear inverse demand function. In terms of our variables, the cost of production is given by $c(a) = \gamma_a a_i^2/2$ and the inverse demand function is given by $p(A) = \gamma_{a\theta}\theta - \gamma_{aA}A$, where the unknown state of the world θ represents the intercept of the inverse demand.*

4 Bayes Nash Equilibrium

We initially report the standard approach to analyze games of incomplete information. Namely, we start with the game of incomplete information, which includes the basic game and a specific type space. Here, the type space consists of a two-dimensional signal that each agent receives. In the first dimension, the signal is privately observed and idiosyncratic to the agent, whereas in the second dimension, the signal is publicly observed and common to all the agents. In either dimension, the signal is normally distributed and centered around the true state of the world θ . In this class of normally distributed signals, a specific type space is determined by the variance of the noise along each dimension of the signal. For a given variance-covariance matrix, and hence for a given type space, we then analyze the Bayes Nash equilibrium/a of the basic game. Now, the type space, as parametrized by the variance of the noise, naturally belongs to a class of possible private information environments and hence type spaces, namely the class of normally distributed bivariate signal structures. In the process of the analysis, we shall observe that the equilibrium behavior across this class of normally distributed information environments displays common features. We shall then proceed to analyze the basic game with the notion of Bayes Correlated equilibrium and establish which

predictions are robust across all of the private information environments, independent of the specific bivariate and normal type space to be considered now.

Accordingly, we consider the following bivariate normal information structure. Each agent i is observing a private and a public noisy signal of the true state of the world θ . The private signal x_i , observed only by agent i , is defined by:

$$x_i = \theta + \varepsilon_i, \quad (6)$$

and the public signal, common and commonly observed by all the agents is defined by:

$$y = \theta + \varepsilon. \quad (7)$$

The random variables ε_i and ε are normally distributed with zero mean and variance given by σ_x^2 and σ_y^2 , respectively; moreover ε_i and ε are independently distributed, with respect to each other and the state θ . This model of bivariate normally distributed signals appears frequently in games of incomplete information, see Morris and Shin (2002) and Angeletos and Pavan (2007) among many others. It is at times convenient to express the variance of the random variables in terms of the precision:

$$\tau_x \triangleq \sigma_x^{-2}, \quad \tau_y \triangleq \sigma_y^{-2}, \quad (8)$$

and we refer to the vector τ with

$$\tau \triangleq (\tau_x, \tau_y), \quad (9)$$

as the information structure of the game.

A special case of the noisy environment is the environment with zero noise. In this environment, the complete information environment, each agent observes the state of the world θ without noise. We begin the equilibrium analysis with the complete information environment. Given the payoff environment, the best response of agent i is given by

$$a_i = -\frac{\lambda_a + \gamma_{a\theta}\theta + \gamma_{aA}A}{\gamma_a}. \quad (10)$$

The best response reflects the, possibly conflicting, objectives that agent i faces. The quadratic payoff function induces each agent to solve a prediction-like problem in which he wishes to match with his action, with the state θ and the average action A . The interaction parameters, the indirect effects $\gamma_{a\theta}$ and γ_{aA} , determine the weight that each component receives in the deliberation of the agent, and the direct effect γ_a , determines the overall responsiveness to state θ and average action A . To wit, if there is zero strategic interaction, or $\gamma_{aA} = 0$, then each agent faces a pure prediction problem. Now, it follows that the resulting Nash equilibrium strategy is given by:

$$a^*(\theta) \triangleq -\frac{\lambda_a}{\gamma_a + \gamma_{aA}} - \frac{\gamma_{a\theta}}{\gamma_a + \gamma_{aA}}\theta. \quad (11)$$

Given the earlier assumptions on individual and aggregate concavity of the payoff function, namely $\gamma_a < 0$ and $\gamma_{aA} + \gamma_a < 0$, it follows that the symmetric strategy $a^*(\theta)$, given any realization θ , is the unique Nash equilibrium of the game with complete information. In fact, Neyman (1997) established that $a^*(\theta)$ is also the unique correlated equilibrium of the game.

Next, we analyze the game with incomplete information, where each agent receives a bivariate noisy signal (x_i, y) . In particular, we shall compare how responsive the strategy of each agent is to the underlying state of the world

relative to the responsiveness in the game with complete information. To this end, we shall refer to the terms in equilibrium strategy (11), $-\lambda_a/(\gamma_a + \gamma_{aA})$ and $-\gamma_{a\theta}/(\gamma_a + \gamma_{aA})$, as the *equilibrium intercept* and the *equilibrium slope*, respectively.

In the game with incomplete information, agent i receives a pair of signals, x_i and y , generated by the information structure (6) and (7). The prediction problem now becomes more difficult for the agent. First, he does not observe the state θ , but rather he receives some noisy signals, x_i and y , of θ . Second, since he does not observe the other agents' signal either, he can only form an expectation about their action, but again has to rely on the signals x_i and y to form the conditional expectation. The best response function of agent i then requires that action a is justified by the conditional expectation, given x_i and y :

$$a_i = -\frac{\lambda_a + \gamma_{a\theta}\mathbb{E}[\theta | x_i, y] + \gamma_{aA}\mathbb{E}[A | x_i, y]}{\gamma_a}. \quad (12)$$

In this linear quadratic environment with normal distributions, we conjecture that the equilibrium strategy is given by a function linear in the signals x_i and y :

$$a(x_i, y) = \alpha_0 + \alpha_x x_i + \alpha_y y. \quad (13)$$

The equilibrium is then identified by the linear coefficients $\alpha_0, \alpha_x, \alpha_y$, which we expect to depend on the interaction matrix Γ and the information structure τ .

Proposition 2 (Linear Bayesian Nash Equilibrium)

The unique Bayes Nash equilibrium is a linear equilibrium:

$$a(x, y) = \alpha_0^* + \alpha_x^* x + \alpha_y^* y \quad (14)$$

with the coefficients given by:

$$\alpha_0^* = -\frac{\lambda_a + \gamma_{a\theta}\mu_\theta}{\gamma_a + \gamma_{aA}} - \frac{\gamma_a}{\gamma_a + \gamma_{aA}} \frac{\sigma_\theta^{-2}\gamma_{a\theta}\mu_\theta}{\gamma_{aA}\sigma_x^{-2} + \gamma_a\sigma^{-2}}, \quad (15)$$

$$\alpha_x^* = -\frac{\gamma_{a\theta}\sigma_x^{-2}}{\gamma_{aA}\sigma_x^{-2} + \gamma_a\sigma^{-2}}, \quad (16)$$

$$\alpha_y^* = -\frac{\gamma_a}{\gamma_a + \gamma_{aA}} \frac{\gamma_{a\theta}\sigma_y^{-2}}{\gamma_{aA}\sigma_x^{-2} + \gamma_a\sigma^{-2}}. \quad (17)$$

The derivation of the linear equilibrium strategy already appeared in many contexts, e.g., in Morris and Shin (2002) for the beauty contest model, and for the present general environment, in Angeletos and Pavan (2007). Interestingly, the Bayes-Nash equilibrium shares the uniqueness property with the Nash equilibrium, its complete information counterpart. We observe that the linear coefficient α_x^* and α_y^* display the following relationship:

$$\frac{\alpha_y^*}{\alpha_x^*} = \frac{\sigma_x^2}{\sigma_y^2} \frac{\gamma_a}{\gamma_a + \gamma_{aA}}. \quad (18)$$

Thus, if there is zero strategic interaction, or $\gamma_{aA} = 0$, then the signals x_i and y receive weights proportional to the precision of the signals. The fact that x_i is a private signal and y is a public signal does not matter in the absence of

strategic interaction, all that matters is the ability of the signal to predict the state of the world. By contrast, if there is strategic interaction, $\gamma_{aA} \neq 0$, then the relative weights also reflect the informativeness of the signal with respect to the average action. Thus if the game displays strategic complements, $\gamma_{aA} > 0$, then the public signal y receives a larger weight. The commonality of the public signal across agents means that their decision is responding to the public signal at the same rate, and hence in equilibrium the public signal is more informative about the average action than the private signal. By contrast, if the game displays strategic complementarity, $\gamma_{aA} < 0$, then each agent would like to move away from the average, and hence places a smaller weight on the public signal y , even though it still contains information about the underlying state of the world.

Now, if we compare the equilibrium strategies under complete and incomplete information, (11) and (13), we find that in the incomplete information environment, each agent still responds to the state of the world θ , but his response to θ is noisy as both x_i and y are noisy realizations of θ , but centered around θ : $x_i = \theta + \varepsilon_i$ and $y = \theta + \varepsilon$. Now, given that the best response, and hence the equilibrium strategy, of each agent is linear in the expectation of θ , the variation in the action is “explained” by the variation in the true state, or more generally in the expectation of the true state. Thus the “flows” of the action have to be balanced by the “flows” of the underlying state θ . But across all of the information structures, the distribution $\psi(\theta)$ of the state θ remains the same, which suggests that the expected flows have to stay constant across the information structures.

Proposition 3 (Attenuation)

The mean action in equilibrium is:

$$\mathbb{E}[a] = \alpha_0^* + \alpha_x^* \mu_\theta + \alpha_y^* \mu_\theta = -\frac{\lambda_a + \gamma_{a\theta} \mu_\theta}{\gamma_a + \gamma_{aA}},$$

and the sum of the weights, $\alpha_x^ + \alpha_y^*$, is:*

$$\alpha_x^* + \alpha_y^* = -\frac{\gamma_{a\theta}}{\gamma_a + \gamma_{aA}} \left(1 - \frac{\gamma_a \sigma_\theta^{-2}}{\gamma_a \sigma^{-2} + \gamma_{aA} \sigma_x^{-2}} \right) \leq -\frac{\gamma_{a\theta}}{\gamma_a + \gamma_{aA}}.$$

Thus, the average action in equilibrium, $\mathbb{E}[a]$, is in fact independent of the information structure τ . In addition, we find that the linear coefficients of the equilibrium strategy under incomplete information are (weakly) less responsive to the true state θ than under complete information. In particular, the sum of the weights is strictly increasing in the precision of the noisy signals x_i and y . The equilibrium response to the state of the world θ is diluted by the noisy signals, that is the response is attenuated. But as the expected flows have to be balanced, the residual is always picked up by the intercept of the equilibrium response.

Now, if we ask how the joint distribution of the Bayes Nash equilibrium vary with the information structure, then Proposition 3 established that it is sufficient to consider the higher moments of the equilibrium distribution. But given the normality of the equilibrium distribution, it follows that it is sufficient to consider the second moments, that is the variance-covariance matrix. The variance-covariance matrix of the equilibrium joint distribution over individual actions a_i, a_j , and state θ is given by:

$$\Sigma_{a_i, a_j, \theta} = \begin{bmatrix} \sigma_a^2 & \rho_a \sigma_a^2 & \rho_{a\theta} \sigma_a \sigma_\theta \\ \rho_a \sigma_a^2 & \sigma_a^2 & \rho_{a\theta} \sigma_a \sigma_\theta \\ \rho_{a\theta} \sigma_a \sigma_\theta & \rho_{a\theta} \sigma_a \sigma_\theta & \sigma_\theta^2 \end{bmatrix}. \quad (19)$$

We denote the correlation coefficient between action a_i and a_j shorthand by ρ_a rather than ρ_{aa} . Above, we describe the equilibrium joint distribution in terms of (a_i, a_j, θ) , but alternatively we can describe it, after replacing the individual action a_j by the average action A , through the triple (a_i, A, θ) . After all, the covariance between the individual, but symmetrically distributed, actions a_i and a_j , given by $\rho_a \sigma_a^2$ has to be equal to the variance of the average action, or $\sigma_A^2 = \rho_a \sigma_a^2$. Similarly, the covariance between the individual action and the average action has to be equal to the covariance of any two, symmetric, individual action profiles, or $\rho_{aA} \sigma_a \sigma_A = \rho_a \sigma_a^2$. Likewise, the covariance between the individual (but symmetric) action a_i and the state θ has to equal to the covariance between the average action and the state θ , or $\rho_{a\theta} \sigma_a \sigma_\theta = \rho_{A\theta} \sigma_A \sigma_\theta$. The variance-covariance matrix of the equilibrium joint distribution over individual actions a_i , A and state θ is given by:

$$\Sigma_{a_i, A, \theta} = \begin{bmatrix} \sigma_a^2 & \rho_a \sigma_a^2 & \rho_{a\theta} \sigma_a \sigma_\theta \\ \rho_a \sigma_a^2 & \rho_a \sigma_a^2 & \rho_{a\theta} \sigma_a \sigma_\theta \\ \rho_{a\theta} \sigma_a \sigma_\theta & \rho_{a\theta} \sigma_a \sigma_\theta & \sigma_\theta^2 \end{bmatrix}.$$

Now, given the characterization of the unique Bayes Nash equilibrium in Proposition 2 above, we can express either of the variance-covariance matrices in terms of the equilibrium coefficients (α_x, α_y) and the variances of the underlying random variables $(\theta, \varepsilon_i, \varepsilon)$, or

$$\Sigma_{a_i, A, \theta} = \begin{bmatrix} \alpha_x^2 \sigma_x^2 + \alpha_y^2 \sigma_y^2 + \sigma_\theta^2 (\alpha_x + \alpha_y)^2 & \alpha_y^2 \sigma_y^2 + \sigma_\theta^2 (\alpha_x + \alpha_y)^2 & \sigma_\theta^2 (\alpha_x + \alpha_y) \\ \alpha_y^2 \sigma_y^2 + \sigma_\theta^2 (\alpha_x + \alpha_y)^2 & \alpha_y^2 \sigma_y^2 + \sigma_\theta^2 (\alpha_x + \alpha_y)^2 & \sigma_\theta^2 (\alpha_x + \alpha_y) \\ \sigma_\theta^2 (\alpha_x + \alpha_y) & \sigma_\theta^2 (\alpha_x + \alpha_y) & \sigma_\theta^2 \end{bmatrix}.$$

Given the structure of the variance-covariance matrix, we can express the equilibrium coefficients α_x^* and α_y^* directly in terms of the variance and covariance terms that they generate:

$$\alpha_x^* = \frac{\sigma_a}{\sigma_\theta} \rho_{a\theta} - \frac{\sigma_a}{\sigma_y} \sqrt{\rho_a - \rho_{a\theta}^2}, \quad \alpha_y^* = \frac{\sigma_a}{\sigma_y} \sqrt{\rho_a - \rho_{a\theta}^2}. \quad (20)$$

In other words, we attribute to the private signal x , through the weight α_x^* , the residual correlation between a and θ , where the residual is obtained by removing the correlation between a and θ which is due to the public signal. In turn, the weight attributed to the public signal is proportional to the difference between the correlation across actions and across action and signal. We recall that the actions of any two agents are correlated as they respond to the same underlying fundamental state θ . Thus, even if their private signals are independent conditional on the true state of the world θ , their actions are correlated due to the correlation with the hidden random variable θ . Now, if these conditionally independent signals were the only sources of information, and the correlation between action and the hidden state θ were $\rho_{a\theta}$, then all the correlation of the agents' action would have to come through the correlation with the hidden state, and in consequence the correlation across actions arises indirectly, in a two way passage through the hidden state, or $\rho_a = \rho_{a\theta} \cdot \rho_{a\theta}$. In consequence, any correlation ρ_a beyond this indirect path, or $\rho_a - \rho_{a\theta}^2$ is generated by means of a common signal, the public signal y .

Since the correlation coefficient of the actions has to be nonnegative, the above representation suggest that as long as the correlation coefficient $(\rho_a, \rho_{a\theta})$ satisfy:

$$0 \leq \rho_a \leq 1, \text{ and } \rho_a - \rho_{a\theta}^2 \geq 0, \quad (21)$$




Figure 1: Bayes Nash equilibrium of beauty contest, $r = 1/4$, with varying degree of precision τ_x of private signal.

we can find information structures τ such the coefficients resulting from (20) indeed are equilibrium coefficients of the associated Bayes Nash equilibrium strategy.

Proposition 4 (Correlation and Information)

For every $(\rho_a, \rho_{a\theta})$ such that $0 \leq \rho_a \leq 1$, and $\rho_a - \rho_{a\theta}^2 \geq 0$, there exists a unique information structure τ such that the associated Bayes Nash equilibrium displays the correlation coefficients $(\rho_a, \rho_{a\theta})$.

In the two-dimensional space of the correlation coefficients $(\rho_a, \rho_{a\theta}^2)$, the set of possible Bayes Nash equilibria is described by the area below the 45° degree line. We illustrate how a particular Bayes Nash equilibrium with its correlation structure $(\rho_a, \rho_{a\theta})$ is generated by a particular information structure τ . In Figure 1, each level curve describes the correlation structure of the Bayes Nash equilibrium for a particular precision τ_x of the private signal. A higher precision τ_x generates a higher level curve. The upward sloping movement represents an increase in informativeness of the public signal, i.e. an increase in the precision τ_y . An increase in the precision of the public signal therefore leads to an increase in the correlation of action across agents as well as in the correlation between individual action and state of the world. For low levels of precision in the private and the public signal, an increase in the precision of the public signal first leads to an increase in the correlation of actions, and then only later into an increased correlation with the state of the world.

In Figure 2, we remain in the unit square of the correlation coefficients $(\rho_a, \rho_{a\theta}^2)$. But this time, each level curve is identified by the precision τ_y of the public signal. As the precision of the private signal increases, the level curve bends upward and first backward, and eventually forward. At low levels of the precision of the private signal, an increase in the precision of the private signal increases the dispersion across agents and hence decreases the correlation across agents. But as it gives each individual more information about the true state of the world, an increase in precision always leads to an increase in the correlation with the true state of the world, this is the upward movement. As the precision improves, eventually the noise becomes sufficiently small so that the underlying common value generated by θ dominates the noise, and then serves to both increase the correlation with the state and across actions. But in contrast to the private information, where the equilibrium sets moves mostly northwards, i.e. where the improvement




Figure 2: Bayes Nash equilibrium of beauty contest, $r = 1/4$, with varying degree of precision τ_x of public signal.

occurs mostly in the direction of an increase in the correlation between the state and the individual agent, the public information leads the equilibrium sets to move mostly eastwards, i.e. most of the change leads to an increase in the correlation across actions. In fact for a given correlation between the individual actions, represented by ρ_a , an increase in the precision of the public signal leads to the elimination of correlated equilibria with very low *and* with very high correlation between the state of the world and the individual action.

5 Bayes Correlated Equilibrium

We now characterize the set of Bayes correlated equilibria. We shall momentarily restrict our attention to symmetric and normally distributed correlated equilibria. We shall later discuss to what extent this restriction is without loss of generality.

With the focus on symmetric equilibria, we can characterize the correlated equilibria in two distinct, yet related, ways. We can either characterize the equilibria in terms of an arbitrary pair of individual actions, a_i and a_j and the state of the world θ . Or, we can characterize the equilibria in terms of the average action A and the deviation of the individual action a_i from the average action, $a_i - A$, and the state of the world θ . The first representation is closer to the description of the Bayes Nash equilibrium in terms of the individual action whereas the second representation puts more emphasis on the distributional properties of the correlated equilibrium.

5.1 Individual and Aggregate Variance

We shall begin the analysis with the first version. We consider the class of symmetric and normally distributed Bayes correlated equilibria. Any correlated equilibrium is then given by a multivariate normal distribution of the vector

(a_i, a_j, θ) of random variables:

$$\begin{pmatrix} a_i \\ a_j \\ \theta \end{pmatrix} \sim N \left(\begin{pmatrix} \mu_a \\ \mu_a \\ \mu_\theta \end{pmatrix}, \begin{pmatrix} \sigma_a^2 & \rho_{a\theta}\sigma_a\sigma_\theta & \rho_{a\theta}\sigma_a\sigma_\theta \\ \rho_a\sigma_a^2 & \sigma_a^2 & \rho_{a\theta}\sigma_a\sigma_\theta \\ \rho_{a\theta}\sigma_a\sigma_\theta & \rho_{a\theta}\sigma_a\sigma_\theta & \sigma_\theta^2 \end{pmatrix} \right). \quad (22)$$

Given the distribution (22) we can compute the expected values of θ and a_j given a_i , which we can think of as a recommendation for agent i . The best response of agent i given any recommendation a_i has to satisfy

$$a_i = -\frac{1}{\gamma_a} (\lambda_a + \gamma_{a\theta}\mathbb{E}[\theta|a_i] + \gamma_{aA}\mathbb{E}[A|a_i]). \quad (23)$$

With the hypothesis of a symmetric equilibrium in a continuum economy, we naturally have that

$$\mathbb{E}[A|a_i] = \mathbb{E}[a_j|a_i], \text{ for all } a_i \in \mathbb{R}.$$

Given that the best response property (23) has to hold for all a_i in the support of the correlated equilibrium, it follows that the above condition has to hold in expectation over all a_i , or by the law of total expectation:

$$\mathbb{E}[a_i] = -\frac{1}{\gamma_a} (\lambda_a + \gamma_{a\theta}\mathbb{E}[\mathbb{E}[\theta|a_i]] + \gamma_{aA}\mathbb{E}[\mathbb{E}[a_j|a_i]]). \quad (24)$$

But by symmetry, it follows that the expected action of agent i is equal to expected action of agent j , and hence we can use (24) to solve for the mean of the individual action and the average action:

$$\mathbb{E}[a_i] = \mathbb{E}[A] = -\frac{\lambda_a + \gamma_{a\theta}\mathbb{E}[\theta]}{\gamma_a + \gamma_{aA}} = -\frac{\lambda_a + \gamma_{a\theta}\mu_\theta}{\gamma_a + \gamma_{aA}}. \quad (25)$$

It thus follows that the mean of the individual action and the mean of the average action is uniquely determined by the mean value μ_θ of the state of the world and the interaction matrix Γ across *all* correlated equilibria.

The complete description of the set of correlated equilibria then rests on the description of the second moments of the multivariate distribution. The characterization of the second moments of the equilibrium distribution again uses the best response property of the individual action, see (23). But, now we use the property of the conditional expectation, rather than the iterated expectation to derive restriction on the covariates. The recommended action a_i has to constitute a best response in the entire support of the equilibrium distribution. Hence the best response has to hold for all $a_i \in \mathbb{R}$, and thus the conditional expectation of the state $\mathbb{E}[\theta|a_i]$ and the representative action a_j of the other agents, $\mathbb{E}[a_j|a_i]$, have to change with a_i at exactly the rate required to maintain the best response property:

$$1 = -\frac{1}{\gamma_a} \left(\gamma_{a\theta} \frac{\partial \mathbb{E}[\theta|a_i]}{\partial a_i} + \gamma_{aA} \frac{\partial \mathbb{E}[a_j|a_i]}{\partial a_i} \right), \text{ for all } a_i \in \mathbb{R}. \quad (26)$$

Given the multivariate normal distribution (2), the conditional expectations $\mathbb{E}[\theta|a_i]$ and $\mathbb{E}[a_j|a_i]$ are linear in a_i and given by

$$\mathbb{E}[\theta|a_i] = \left(1 + \frac{\rho_{a\theta}\sigma_\theta}{\sigma_a} \frac{\gamma_{a\theta}}{\gamma_a + \gamma_{aA}} \right) \mu_\theta + \frac{\rho_{a\theta}\sigma_\theta}{\sigma_a} a_i, \quad (27)$$

and

$$\mathbb{E}[a_j|a_i] = -\frac{\gamma_{a\theta}}{\gamma_a + \gamma_{aA}} \mu_\theta (1 - \rho_a) + \rho_a a_i. \quad (28)$$

The optimality of the best response property can then be expressed, using (27) and (28) as

$$1 = -\frac{1}{\gamma_a} \left(\gamma_{a\theta} \frac{\rho_{a\theta}\sigma_\theta}{\sigma_a} + \gamma_{aA}\rho_a \right). \quad (29)$$

It follows that we can express either one of the three elements in the description of the second moments, $(\sigma_a, \rho_a, \rho_{a\theta})$ in terms of the other two and the primitives of the game as described by the interaction matrix Γ . In fact, it is convenient to solve for the standard deviation of the individual actions σ_a , or

$$\sigma_a = -\frac{\sigma_\theta \gamma_{a\theta} \rho_{a\theta}}{\rho_a \gamma_{aA} + \gamma_a}. \quad (30)$$

The remaining restrictions on the correlation coefficients are imposed by the requirement that the variance-covariance matrix is nonnegative definite.

The condition on the variance of the individual action, given by (30), actually follows the same logic as the condition on the mean of the individual action, given by (25). To wit, for the mean, we used the law of total expectation to arrive at the equality restriction. Similarly, we could obtain the above restriction (30) by using the law of total variance and covariance. More precisely, we could require, using the equality (??), that the variance of the individual action matches the sum of the variances of the conditional expectation. Then, by using the law of total variance and covariance, we could represent the variance of the conditional expectation in terms of the variance of the original random variables, and obtain the exact same condition (30). Here we chose to directly use the linear form of the conditional expectation given by the multivariate normal distribution. We explain towards the end of the section that the later method, which restricts the moments via conditioning, remains valid beyond the multivariate normal distributions.

Proposition 5 (First and Second Moments of BCE)

A multivariate normal distribution of (a_i, a_j, θ) is a symmetric Bayes correlated equilibrium iff

1. *the mean of the individual action is:*

$$\mathbb{E}[a_i] = -\frac{\lambda_a + \gamma_{a\theta}\mu_\theta}{\gamma_a + \gamma_{aA}}; \quad (31)$$

2. *the standard deviation of the individual action is:*

$$\sigma_a = -\frac{\gamma_{a\theta}\rho_{a\theta}}{\rho_a\gamma_{aA} + \gamma_a}\sigma_\theta; \quad \text{and} \quad (32)$$

3. *the correlation coefficients satisfy the inequality:*

$$0 \leq 1 - \rho_a \leq 1, \quad \rho_a - \rho_{a\theta}^2 \geq 0. \quad (33)$$

The characterization of the first and second moments suggests that the mean μ_θ and the variance σ_θ^2 of the fundamental variable θ are the driving force of the moments of the equilibrium actions. The linear form of the best response function translates into a linear relationship in the first and second moment of the state of the world and the equilibrium action. In case of the standard deviation the linear relationship is affected by the correlation coefficients ρ_a and $\rho_{a\theta}$ which assign weights to the interaction parameter γ_{aA} and $\gamma_{a\theta}$, respectively. The set of all correlated equilibria is graphically represented in Figure 3 below:




Figure 3: Set of Bayes correlated equilibrium in terms of correlation coefficients ρ_a and $\rho_{a\theta}$

We should emphasize that at $\rho_{a\theta} = 0$, the only correlated equilibrium is given by $\rho_a = 1$, in other words, there is a discontinuity in the equilibrium set at $\rho_{a\theta} = 0$. In the symmetric equilibrium, if $\rho_{a\theta} = 0$, then this means that the action of each agent is completely insensitive to the realization of the true state θ . But this means, that the agents do not respond to any information about the state of the world θ beyond the expected value of the state, $\mathbb{E}[\theta]$. Thus, each agent acts as if he were in a complete information world where the true state of the world is the expected value of the state. But, we know from the earlier discussion, that in this environment, there is a unique correlated equilibrium where the agents all choose the same action and hence $\rho_a = 1$. The standard deviation σ_a can be expressed in terms of the elements of the interaction matrix and the correlation coefficients:

$$\sigma_a = -\frac{\sigma_\theta \gamma_{a\theta} \rho_{a\theta}}{\rho_a \gamma_{aA} + \gamma_a}, \quad (34)$$

and since by the assumption of moderate interaction, $\gamma_{aA} + \gamma_a < 0$, the denominator is negative for all $\rho_a \in [0, 1]$, or

$$\rho_a \gamma_{aA} + \gamma_a < 0, \quad (35)$$

it follows that to guarantee that $\sigma_a \geq 0$:

$$\gamma_{a\theta} \rho_{a\theta} \geq 0. \quad (36)$$

Thus the sign of the correlation coefficient $\rho_{a\theta}$ has to equal the sign of the interaction term $\gamma_{a\theta}$.

We saw in Proposition 5 that the relationship between the correlation coefficients ρ_a and $\rho_{a\theta}$ depended only on the sign of the information externality $\gamma_{a\theta}$, but not on the strength of the parameters γ_a, γ_{aA} and $\gamma_{a\theta}$. We can therefore focus our attention on the variance of the individual action and how it varies with the strength of the interaction as measured by the correlation coefficients $(\rho_a, \rho_{a\theta})$.

Proposition 6 (Variance of Individual Action)

1. If the game displays strategic complements, $\gamma_{aA} > 0$, then:

- (a) σ_a is increasing in ρ_a and $|\rho_{a\theta}|$;

(b) the maximal variance σ_a is obtained at $\rho_a = |\rho_{a\theta}| = 1$.

2. If the game displays strategic substitutes, $\gamma_{aA} < 0$, then

(a) σ_a is decreasing in ρ_a and increasing in $|\rho_{a\theta}|$;

(b) the maximal variance σ_a is obtained at

$$\rho_a = |\rho_{a\theta}|^2 = \min \left\{ \frac{\gamma_a}{\gamma_{aA}}, 1 \right\}. \quad (37)$$

In particular, we find that as the correlation in the actions across individuals increases, the variance in the action is amplified in the case of strategic complements, but attenuated in the case of strategic complements. An interesting implication of the attenuation of the individual variance is that the maximal variance of the individual action may not be attained under minimal or maximal correlation of the individual actions but rather at an intermediate level of interaction. In particular, if the interaction effect γ_{aA} is larger than the own effect γ_a , i.e.

$$|\gamma_{aA}| > |\gamma_a|,$$

then the maximal variance σ_a is obtained with an interior solution. Of course, in the case of strategic complements, the positive feed-back effect implies that the maximal variance is obtained when the actions are maximally correlated.

5.2 Volatility and Dispersion

So far we have described the correlation equilibrium in terms of the triple (a_i, a_j, θ) . Yet, a distinct but equivalent representation can be given in term of the average action A and the idiosyncratic difference, $a - A$ and the state θ . This alternative representation in terms of $(A, a - A, \theta)$ is frequently adopted in games with a continuum of agents. In this context we can interpret the conditional distribution of the agents' action a around the mean A as exact distribution of the actions in the population. The idiosyncratic difference $a - A$ describes the dispersion around the average action, and the variance of the average action A can be interpreted as the volatility of the game. The dispersion, $a - A$, measures how much the individual action can deviate from the average action, yet be justified consistently with the conditional expectation of each agent in equilibrium. The language for volatility and dispersion in the context of this environment was earlier suggested by Angeletos and Pavan (2007).

We therefore obtain an alternative representation of the set of correlated equilibria in terms of $(A, a - A, \theta)$ with:

$$\begin{pmatrix} A \\ a - A \\ \theta \end{pmatrix} \sim N \left(\begin{pmatrix} -\frac{\lambda_a + \gamma_{a\theta}\mu_\theta}{\gamma_a + \gamma_{aA}} \\ 0 \\ \mu_\theta \end{pmatrix}, \begin{pmatrix} \sigma_A^2 & 0 & \rho_{A\theta}\sigma_A\sigma_\theta \\ 0 & \sigma_{a-A}^2 & 0 \\ \rho_{A\theta}\sigma_A\sigma_\theta & 0 & \sigma_\theta^2 \end{pmatrix} \right). \quad (38)$$

By construction, the idiosyncratic difference $a - A$ describes the dispersion, which is statistically independent of the average action A and the state θ . By definition, the dispersion has zero mean and we denote the standard deviation of the dispersion by σ_{a-A} . Given the entries in the variance-covariance matrix (38), we can relate the distributions of (a_i, a_j, θ) and $(A, a - A, \theta)$ as follows.

Proposition 7

The second moments of the individual action a and the average action A display the following statistical relationships:

1. $\rho_{aA} = \sqrt{\rho_a}$;
2. $\rho_{A\theta} = \rho_{a\theta}/\sqrt{\rho_a}$;
3. $\sigma_A = \sqrt{\rho_a}\sigma_a$.

These statistical relations between the individual action and the average action in the continuum economy allows to describe the variance-covariance matrix of any correlated equilibrium in terms of the individual action and the state of the world. We can write the joint distribution of (a, A, θ) as follows

$$\begin{pmatrix} a \\ A \\ \theta \end{pmatrix} \sim N \left(\begin{pmatrix} -\frac{\lambda_a + \gamma_{a\theta}\mu_\theta}{\gamma_a + \gamma_{aA}} \\ -\frac{\lambda_a + \gamma_{a\theta}\mu_\theta}{\gamma_a + \gamma_{aA}} \\ \mu_\theta \end{pmatrix}, \begin{pmatrix} \sigma_a^2 & \rho_a\sigma_a^2 & \rho_{a\theta}\sigma_a\sigma_\theta \\ \rho_a\sigma_a^2 & \rho_a\sigma_a^2 & \rho_{a\theta}\sigma_a\sigma_\theta \\ \rho_{a\theta}\sigma_a\sigma_\theta & \rho_{a\theta}\sigma_a\sigma_\theta & \sigma_\theta^2 \end{pmatrix} \right) \quad (39)$$

and thus eliminate σ_A^2, ρ_{aA} and $\rho_{A\theta}$ after introducing the correlation coefficient ρ_a across individual actions using the results of Proposition 7. The dispersion is described by the variance of $a - A$, which is given by $(1 - \rho_a)\sigma_a^2$ whereas the aggregate volatility is given by $\rho_a\sigma_a^2$.

Proposition 8 (Volatility and Dispersion)

1. The volatility is increasing in ρ_a and $|\rho_{a\theta}|$
2. The dispersion is increasing in $|\rho_{a\theta}|$ and reaches an interior maximum at:

$$\rho_a = \frac{\gamma_a}{\gamma_{aA} + 2\gamma_a}, \quad \rho_{a\theta}^2 = \rho_a.$$

The dispersion, $a - A$, measures how much the individual action can deviate from the average action. The maximal level of dispersion occurs when the correlation with respect to the state θ is largest. But it reaches its maximum at an interior level of the correlation across the individual actions as we might expect. We note that relative to the variance of the individual action, see Proposition 6, the volatility, is increasing in the correlation coefficient ρ_a irrespective of the nature of the strategic interaction.

5.3 Matching Bayes Correlated and Nash Equilibria

The description of the Bayes correlated equilibria lead to a complete characterization of the equilibrium behavior of the agents. Yet, the construction of the equilibrium set did not give us any direct information as to how rich and complicated an information structure would have to be to support the behavior in terms of a related Bayes Nash equilibrium. We know from the epistemic result of Proposition 1 that such information structures exists, but we do not yet know which form they may take. We now describe the relationship between Bayes correlated and Bayes Nash

equilibria by constructing the information structure implicitly associated with every Bayes correlated equilibrium. We are going to describe a class of bivariate information structures, such that the union of the Bayes Nash equilibria generated by these information structures spans the entire set of Bayes correlated equilibria.

We observe that the Bayes Nash and correlated equilibria share the same mean. We can therefore match the respective equilibria if we can match the second moments of the equilibria. After inserting the coefficients of the linear strategies of the Bayes Nash equilibrium, we can match the moments of the two equilibrium notions. In the process, we get two equations relating the Bayes correlated and Nash equilibrium. The Bayes Nash equilibria are defined by the variance of the private and the public signal. The correlated equilibria are defined by the correlation coefficients of individual actions across agents, and individual actions and state θ .

Proposition 9 (Matching BCE and BNE)

For every interaction structure Γ , there is a bijection between Bayes correlated and Bayes Nash equilibrium.

Finally we observe that for a given finite precision of the information structure, i.e. $0 < \tau < \infty$, the associated Bayes Nash equilibrium is an interior point relative to the set of correlated equilibria. As the set of correlated equilibria is described by $\rho_a - \rho_{a\theta}^2 \geq 0$, and since we know that $\rho_a = (\rho_{aA})^2$ we have $\rho_{aA} > \rho_{a\theta}$. It follows that the Bayes Nash equilibrium is an interior equilibrium relative to the correlated equilibria in terms of the correlation coefficients, and certainly in terms of the variance of individual and average action. To put it differently, the equality $\rho_a = \rho_{a\theta}^2$ is obtained in the Bayes Nash equilibrium if and only if $\tau_y = 0$.

The above description of the bijection between Bayes correlated and Bayes Nash equilibrium was stated for the class of normally distributed Bayes Nash equilibria. An interesting aspect of the constructive approach was that a bivariate information structure was sufficient to generate the entire set of Bayes correlated equilibrium. We conjecture that the sufficiency of a bivariate information structures is likely to remain valid even with general distribution of fundamental uncertainty. After all, the correlation coefficients arise from idiosyncratic dispersion and aggregate volatility. The private signal supports the idiosyncratic dispersion and the public signal is sufficient to support the aggregate volatility.

5.4 Beyond Normal Distributions and Symmetry

Beyond Normal Distributions The above characterization of the mean and variance of the equilibrium distribution was obtained under the assumption that the distributions of the fundamental variable θ and resulting joint distribution was a multivariate normal distribution. Now, even if the distribution of the state of the world θ is a normally distributed, the joint equilibrium distribution does not necessarily have to be a normal distribution itself. If the equilibrium distribution is not a multivariate normal distribution anymore, then the first and second moments alone do not completely characterize the equilibrium distribution anymore. In other words, the first and second moment only impose restrictions on the higher moments, but do not completely identify the higher moments anymore. We observe however that the restrictions regarding the first and second moment remain to hold. In particular, the result regarding the mean of the action is independent of the distribution of the equilibrium or even the normality of the fundamental variable θ . With respect to the restrictions on the second moments, the restrictions still hold, but

outside of the class of multivariate normal distribution, the inequalities may not necessarily be achieved as equalities for some equilibrium distributions.

In this context, it is worthwhile to note that the equilibrium characterization of the first and second moments could alternatively be obtained by using the law of total expectation, and its second moment equivalents, the law of total variance and covariance. These “laws”, insofar as they relate marginal probabilities to conditional probabilities, naturally appeared in the equilibrium characterization of the best response function which introduce the conditional expectation over the state and the average action, and hence the conditional probabilities. For higher-order moments, a natural a simple and elegant generalization of this relationship exists, see Brillinger (1969), sometimes referred to as law of total cumulance, and as such would deliver further restrictions on higher-order moments if we were to consider equilibrium distributions beyond the normal distribution.

Beyond Symmetry The above characterization of the mean and variance of the equilibrium distribution pertained to the symmetric equilibrium distribution. But actually, the characterization remains entirely valid for *all* equilibrium distributions if we focus on the average action rather than the individual action. In addition, the result about the mean of the individual action remains true for all equilibrium distribution, and not only the symmetric equilibrium distribution. This later result suggests that the asymmetric equilibria only offer a richer set of possible second moments distributions across agents. Interestingly, in the finite agent environment, the asymmetry in the second moments does not lead to joint distributions over aggregates outcomes and state which cannot be obtain already with symmetric equilibrium distributions.

6 Prior Information

The description of the Bayes correlated equilibria displayed a rich set of possible equilibrium outcomes. In particular, the variance of the individual and the average action had a wide range across equilibria. The analysis of the Bayes Nash equilibrium shed light on the source of the variation. If the noisy signals of each agent contained little information about the state of the world, then the action of each agent did not vary much in the realization of the signal. On the other hand, with precise information about the true state of the world, the best response of each agent would vary substantially with the realized signal and hence would display a larger variance in equilibrium. In the spirit of the robust analysis, we began without any assumptions on the nature of the private information that the agents may have when they make their decisions. But in many circumstances, there may be prior knowledge about the nature of the private information of the agents. In particular, we may able to impose a lower bound on the private information that the agents may have. We can then ask how the prediction of the equilibrium behavior can be refined in the presence of prior restrictions on the private information of the agents.

Given the sufficiency of a bivariate information structure to support the entire equilibrium set, we present the lower bounds on the private information here in terms of a private and a public information source, each one given in terms of a normally distributed noisy signal. We maintain the notation from Section 4 and denote the private signal that each agent i observes by $x_i = \theta + \varepsilon_i$, and the public signal that all agents observe by $y = \theta + \varepsilon$, as defined earlier in (6) and (7), respectively.

The exogenous data on the payoff and belief environment of the game is now given by the multivariate normal distribution of the triple (θ, x_i, y) . The information contained in the private signal x_i and the public signal y represent the lower bound on the private information of the agents. Naturally, this analysis includes the situation of a lower bound on the private or the public information only. In this case, the precision of the complementary signal is simply assumed to be zero. Correspondingly, we can define a Bayes correlated equilibrium with given private information as a joint distribution over the exogenous data (θ, x, y) and the endogenous data (a, A) . We use the symmetry and the relationship between the individual action and the average action to obtain a compact representation of the variance-covariance matrix $\Sigma_{\theta, x, y, a, A}$:

$$\begin{pmatrix} \sigma_\theta^2 & \sigma_\theta^2 & \sigma_\theta^2 & \rho_{a\theta}\sigma_a\sigma_\theta & \rho_{a\theta}\sigma_a\sigma_\theta \\ \sigma_\theta^2 & \sigma_\theta^2 + \sigma_x^2 & \sigma_\theta^2 & \sigma_a\sigma_x\rho_{ax} + \sigma_a\sigma_\theta\rho_{a\theta} & \sigma_a\sigma_\theta\rho_{a\theta} \\ \sigma_\theta^2 & \sigma_\theta^2 & \sigma_\theta^2 + \sigma_y^2 & \sigma_a\sigma_y\rho_{ay} + \sigma_a\sigma_\theta\rho_{a\theta} & \sigma_a\sigma_y\rho_{ay} + \sigma_a\sigma_\theta\rho_{a\theta} \\ \rho_{a\theta}\sigma_a\sigma_\theta & \sigma_a\sigma_x\rho_{ax} + \sigma_a\sigma_\theta\rho_{a\theta} & \sigma_a\sigma_y\rho_{ay} + \sigma_a\sigma_\theta\rho_{a\theta} & \sigma_a^2 & \rho_a\sigma_a^2 \\ \rho_{a\theta}\sigma_a\sigma_\theta & \sigma_a\sigma_\theta\rho_{a\theta} & \sigma_a\sigma_y\rho_{ay} + \sigma_a\sigma_\theta\rho_{a\theta} & \rho_a\sigma_a^2 & \rho_a\sigma_a^2 \end{pmatrix} \quad (40)$$

The newly appearing correlation coefficients ρ_{ax} and ρ_{ay} represent the correlation between the individual action and the random terms, ε_i and ε , in the private and public signals, x_i and y , respectively. We can now analyze the correlated equilibrium conditions as before. The best response function must satisfy:

$$a = -\frac{1}{\gamma_a} (\gamma_{a\theta}\mathbb{E}[\theta | a, x, y] + \gamma_{aA}\mathbb{E}[A | a, x, y]), \quad \forall a, x, y. \quad (41)$$

In contrast to the analysis of the Bayes correlated equilibrium without prior information, the recommended action now has to form a best response conditional on the recommendation a and the realization of the private and public signals, x_i and y , respectively. In particular, the conditional expectation induced jointly by (a, x, y) has to vary at a specific rate with the realization of a, x, y so as to maintain the best response property (41) for all realizations of a, x, y :

$$\begin{aligned} 1 &= \frac{\partial}{\partial a} \left[-\frac{1}{\gamma_a} (\gamma_{a\theta}\mathbb{E}[\theta | a, x, y] + \gamma_{aA}\mathbb{E}[A | a, x, y]) \right], \quad \forall a, x, y; \\ 0 &= \frac{\partial}{\partial x} \left[-\frac{1}{\gamma_a} (\gamma_{a\theta}\mathbb{E}[\theta | a, x, y] + \gamma_{aA}\mathbb{E}[A | a, x, y]) \right], \quad \forall a, x, y; \\ 0 &= \frac{\partial}{\partial y} \left[-\frac{1}{\gamma_a} (\gamma_{a\theta}\mathbb{E}[\theta | a, x, y] + \gamma_{aA}\mathbb{E}[A | a, x, y]) \right], \quad \forall a, x, y. \end{aligned} \quad (42)$$

The complete characterization of the set of Bayes correlated equilibria with prior information requires the determination of a larger set of second moments, namely $(\sigma_a, \rho_a, \rho_{ax}, \rho_{ay}, \rho_{a\theta})$ than in the earlier analysis. As we gather the equilibrium restrictions in (42), we find that we also have a corresponding increase in the number of equality constraints on the equilibrium conditions. Indeed, from the conditions (42) we can determine $(\rho_{ay}, \rho_{ax}, \sigma_a)$ uniquely:

$$\sigma_a = -\frac{\sigma_\theta\gamma_{a\theta}\rho_\theta}{\gamma_a + \rho_a\gamma_{aA}}, \quad (43)$$

$$\rho_{ax} = \frac{\sigma_\theta}{\sigma_x\rho_{a\theta}} \left(\frac{\gamma_a + \rho_a\gamma_{aA}}{\gamma_a} - \rho_{a\theta}^2 \frac{\gamma_a + \gamma_{aA}}{\gamma_a} \right) \quad (44)$$

$$\rho_{ay} = \frac{\sigma_\theta}{\sigma_y\rho_{a\theta}} \left(\frac{\gamma_a + \rho_a\gamma_{aA}}{\gamma_a} - \rho_{a\theta}^2 \frac{\gamma_a + \gamma_{aA}}{\gamma_a} \right). \quad (45)$$

Notably, the characterization of the standard deviation of the individual action has not changed relative to the initial analysis. The novel restrictions on the correlation coefficients ρ_{ax} and ρ_{ay} only involve γ_a and γ_{aA} through the ratio γ_{aA}/γ_a , but the informational externality $\gamma_{a\theta}$ does not appear. The ratio γ_{aA}/γ_a already appeared earlier as we imposed individual and aggregate concavity on the payoff matrix. It is now convenient to define the ratio γ_{aA}/γ_a as the strategic interaction term r_A :

$$r_A \triangleq -\frac{\gamma_{aA}}{\gamma_a}.$$

Given the earlier assumptions, see (4), the ratio r_A can vary from $-\infty$ to 1. A negative value represents a game with strategic substitutes, whereas a positive value of r_A represents a game with strategic complements. Consequently, the relation between the correlation coefficients ρ_{ax} and ρ_{ay} can be written, using the conditions (44) and (45) as:

$$\rho_{ax}\sigma_x = \rho_{ay}\sigma_y(1 - r_A),$$

where the factor $1 - r_A$ corrects for the fact that the public signal receives a different weight than the private signal due to the interaction structure.

The additional inequality restrictions arise as the variance-covariance matrix of the multivariate normal distribution has to form a positive semidefinite matrix², or: \times

$$\sigma_a^4 \sigma_y^2 \sigma_x^2 \sigma_\theta^2 (1 - \rho_a - \rho_{ax}^2) (\rho_a - \rho_{a\theta}^2 - \rho_{ay}^2) \geq 0.$$

Thus the additional inequalities which completely describe the set of correlated equilibria are given by:

$$1 - \rho_a - \rho_{ax}^2 \geq 0, \tag{46}$$

$$\rho_a - \rho_{a\theta}^2 - \rho_{ay}^2 \geq 0. \tag{47}$$

We encountered the above inequalities before, see Proposition 5.3, but without the additional entries of ρ_{ax} and ρ_{ay} . The first inequality reflects the equilibrium restriction between ρ_a and ρ_{ax} . As ρ_{ax} represents the correlation between the individual action a and the idiosyncratic signal x , it imposes an upper bound on the correlation coefficient ρ_a among individual actions. If each of the individual actions are highly correlated with their private signal, then the correlation of the individual actions cannot be too high in equilibrium. Conversely, the second inequality states that either the correlation between individual action and public signal, or individual action and state of the world naturally force an increase in the correlation across individual actions. The correlation coefficients $\rho_{a\theta}$ and ρ_{ay} therefore impose a lower bound on the correlation coefficient ρ_a .

The equilibrium restrictions imposed by the private and public signal are separable. We can hence combine (44) and (46), or (45) and (47), respectively, to analyze how the private or the public signal restrict the set of Bayes correlated equilibria. Given that the mean action is constant across the Bayes correlated equilibria and that the variance σ_a^2 of the action is determined by the correlation coefficients $(\rho_a, \rho_{a\theta})$, see (43), we can describe the set of Bayes correlated equilibria exclusively in terms of correlation coefficients $(\rho_a, \rho_{a\theta})$. It is more natural to state the subsequent results in terms of the precision, rather than the variance of the private and public signal, and we denote:

$$\tau_x \triangleq \sigma_x^{-2}, \quad \tau_y \triangleq \sigma_y^{-2}. \tag{48}$$

²The variance-covariance matrix Σ is required to be positive semidefinite (sometimes called nonnegative definite) rather than positive definite as we allow for linear relationships between the components of the random vector (see DeGroot (1970)).




Figure 4: Set of BCE with given public and private information

We define the set of all Bayes correlated equilibria which are consistent with prior private information τ_x as the *private equilibrium set* $C_x(\tau_x, r_A)$:

$$C_x(\tau_x, r_A) \triangleq \{(\rho_a, \rho_{a\theta}) \in [0, 1] \times [-1, 1] \mid (\rho_a, \rho_{a\theta}, \rho_{ax}) \text{ satisfy (44), (46)}\}.$$

Similarly, we define the set of all Bayes correlated equilibria which are consistent with prior public information τ_y as the *public equilibrium set* $C_y(\tau_y, r_A)$:

$$C_y(\tau_y, r_A) \triangleq \{(\rho_a, \rho_{a\theta}) \in [0, 1] \times [-1, 1] \mid (\rho_a, \rho_{a\theta}, \rho_{ay}) \text{ satisfy (45), (47)}\}.$$

The intersection of the private and the public equilibrium sets defines the Bayes correlated equilibria consistent with the prior information $\tau = (\tau_x, \tau_y)$:

$$C(\tau, r_A) \triangleq C_x(\tau_x, r_A) \cap C_y(\tau_y, r_A) \subset [0, 1] \times [-1, 1].$$

The shape of the Bayes correlated equilibrium set is illustrated in Figure 4. Each forward bending curve describes the set of correlation coefficients $(\rho_a, \rho_{a\theta})$ which solve (45) and (46) as an equality, given a lower bound on the precision τ_x of the private information. Similarly, each backward bending curve traces out the set of correlation coefficients $(\rho_a, \rho_{a\theta})$ which solve (??) and (47) as an equality, given a lower bound on the precision τ_y of the public information. A lens formed by the intersection of a forward and a backward bending curve represents the Bayes correlated equilibria consistent with a lower bound on the precision of the private and the public signal.

As suggested by the behavior of the equilibrium set, any correlation in the actions cannot undo the given private and public information, but rather provides additional correlation opportunities over and above those contained in τ .

Proposition 10 (Prior Information) For all $r_A \in (-1, \infty)$:

1. The equilibrium set $C(\tau, r_A)$ is decreasing in τ ;
2. The lowest correlation coefficient $\rho_a \in C(\tau, r_A)$, $\min_{\rho_a \in C(\tau, r_A)} \rho_a$, is increasing in τ ;

3. The lowest correlation coefficient $\rho_{a\theta} \in C(\tau, r_A)$, $\min_{\rho_{a\theta} \in C(\tau, r_A)} \rho_{a\theta}$, is increasing in τ .

Thus, as the precision of the prior information increases, the set of Bayes correlated equilibria shrink. As the precision of the signal increases, the equilibrium set, as represented by the correlation coefficients becomes smaller. In particular, the lowest possible correlation coefficients of ρ_a and $\rho_{a\theta}$ that may emerge in any Bayes correlated equilibrium increase as the given private information increases.

As the preceding discussion suggests, we can relate the set of Bayes Correlated equilibria under the prior information with a corresponding set of Bayes Nash equilibria. If the correlated equilibrium contains no additional information in the conditioning through the recommended action a over and above the private and public signal, x and y , then the correlated equilibrium is simply equal to the Bayes Nash equilibrium with the specific information structure τ . This suggests that we identify the unique Bayes Nash equilibrium with information structure τ and interaction term r_A in terms of the correlation coefficients $(\rho_a, \rho_{a\theta})$ as $B(\tau, r_A) \in [-1, 1]^2$.

Corollary 1 (BCE and BNE with Prior Information)

For all $\tau = (\tau_x, \tau_y)$, we have:

$$C(\tau, r_A) = \bigcup_{\tau'_x \geq \tau_x, \tau'_y \geq \tau_y} B(\tau', r_A).$$


Proposition 10 described the dependence of the equilibrium set on the prior information. Alternatively, we can describe the dependence of the equilibrium set on the structural parameters of the game. In other words, we may ask how the set of equilibrium correlation coefficients varies with the interaction parameter r_A of the game.

Proposition 11 (Prior Information and Interaction)

For all $\tau \in \mathbb{R}_+^2$:

1. $C_x(\tau, r_A)$ is decreasing in r_A ;
2. $C_y(\tau, r_A)$ is increasing in r_A .

The comparative static results in the strategic interaction r_A are straightforward. The information in the private signal x leads each agent to choose an action which is less correlated with the average action than the same information contained in the public signal y . Now, as the interaction in the game tends towards strategic substitutability, each agent tends to rely more heavily on the private signal relative to the public signal. Thus, for every level of correlation with the state θ , expressed in terms of $\rho_{a\theta}$, there will be less correlation across actions, expressed in terms of ρ_a . Conversely, the restrictions imposed by the public information, represented by the set $C_y(\tau_y, r_A)$ become weaker as the game is moving away from strategic complements to strategic substitutes. After all, the public information correlates the agent's action because they rely on the same information. If we decrease the propensity to coordinate, and hence correlate, then all equilibria will display less correlation across actions, for a given correlation with respect to the state θ . We thus find that the comparative static results with respect to the strategic interaction are pointing in the opposite direction for the equilibrium sets $C_x(\tau_x, r_A)$ and $C_y(\tau_y, r_A)$, respectively. In consequence, the equilibrium set $C(\tau_y, r_A)$ formed by the intersection of the private and public equilibrium sets, $C_x(\tau_x, r_A) \cap C_y(\tau_y, r_A)$, does not display a monotone behavior in r_A in terms of set inclusion.

Figure 5: Bayes correlated equilibrium set with precision τ_x of prior private information

The behavior of the equilibrium sets $C_x(\tau, r_A)$ and $C_y(\tau, r_A)$ with respect to the strategic interaction r_A are illustrated in Figure 5 and 6. The set $C_x(\tau, r_A)$ is decreasing in r_A . As long as the strategic complementarity is not too large, then the boundary of the equilibrium set is described by a function $\rho_a(\rho_{a\theta})$ (rather than a correspondence), which identifies the largest correlation coefficient ρ_a consistent with the precision τ_x of the private signal and a given correlation coefficient $\rho_{a\theta}$. However, if the degree of complementarity becomes too large, or $r_A > 1/2$, then the boundary of the equilibrium set can no longer be described by a function $\rho_a(\rho_{a\theta})$. To wit, as we increase the precision τ_y of the public signal, the correlation coefficient ρ_a initially drops, before it recovers again. The failure of monotonicity of $\rho_{a\theta}$ in the precision τ_y of the public signal arises with a strong level of complementarity. Here an increase in the public information leads the agent to rely so heavily on the public signal that an improvement in the precision of the public signal, at a constant precision of the private signal, leads to a decrease in the correlation coefficient $\rho_{a\theta}$. In other words, the quality of the equilibrium decision, as measured by the correlation with the state θ , $\rho_{a\theta}$, decreases, even though the precision of the public signal increases, while the precision of the private signal stays constant. This strong responsiveness of the equilibrium action to the public information in games with strategic complements already appeared prominently in Morris and Shin (2002), where they established that the social value of public information can be negative in the presence of strategic complements. In the context of this result, Proposition 10 may be interpreted as the robust counterpart to the negative value of public information. Proposition 10.3 established that the lowest correlation coefficient $\rho_{a\theta}$ consistent with some information structure is increasing in the precision τ_y of the public signal. Thus, while the social value of public information may be negative for some information structure, it always increase the lowest value $\rho_{a\theta}$ across all possible information structures.

The behavior of the equilibrium set $C_y(\tau, r_A)$ generated by prior public information displays a similarly surprising behavior, but now for strategic substitutes rather than complements. As long as the degree of substitutability is not too large, or $r_A \geq -1$, then the equilibrium restriction is described by a function of $\rho_a(\rho_{a\theta})$. In particular, an increase in the private information always leads to an increase in the correlation coefficient $\rho_{a\theta}$. However, if $r_A < -1$, then an increase in the private information leads initially to a decrease in the correlation coefficient $\rho_{a\theta}$. Thus, again, despite the improvement in the available information, this time in the private information, while the public




Figure 6: Bayes correlated equilibrium set with precision τ_y of prior public information

information is held constant, the correlation between the action and state decreases. Now, the reason for the strong responsiveness to the private information lies in strategic substitutability. The private information allows each agent to choose an action away from the average action, and hence the private information receives considerably more weight than the public information relative to the informativeness of the signals. In consequence, an increase in the private information leads to a decrease in the correlation coefficient ρ_a as well as $\rho_{a\theta}$. This is illustrated in Figure 6.

7 Robust Identification

So far, our analysis has been concerned with the predictive implications of Bayes correlated and Bayes Nash equilibrium. In particular, we have been asking what are the restrictions imposed by the structural model on the observed endogenous statistics about the actions of the agents. In this section we pursue the converse question, namely the issue of identification. We ask what restrictions can be imposed on the parameters of interest, namely the structural parameters of the game by the observed variables? We are particularly interested in how the identification of the structural parameters is influenced by the solution concept, and hence the specification of the private information of the agents as known to the analyst.

Now, identification may depend critically on what types of data are available. We therefore consider first the possibility of identification with individual data, and then with aggregate data. The identification with aggregate data is centered around the canonical problem of demand and supply identification. In contrast to the received work on identification in the demand and supply model we allow for incomplete information by the market participants. The problem of identification in the demand and supply model relies only on aggregate data about the action of the agents, the aggregate quantity and the price, which constitutes a *linear combination* of the aggregate action and the realized state of the world.

7.1 Robust Identification and Individual Decision Making

First, we pose the question of robust identification in the context of individual decision making, without any element of strategic interaction. Within the context of individual decision-making, the question is how much we can learn from the observed data when the analyst cannot be certain about the information that the agent has when he chooses his action. In the single agent context, the noisy signal x that the agent receives about the state of world θ leads to noise in the predictor variable. The noise in the predictor variable introduces a bias, which is commonly referred to as “attenuation bias”. But importantly, in the single agent model, the sign of the parameter of the interest, the informational externality $\gamma_{a\theta}$ remains correctly identified, even though the information externality is set-identified rather than point-identified. Now, as we extend the analysis to strategic interaction, the “attenuation bias” critically affects the ability to identify the nature of the strategic interaction. In particular, the set-identified information externality “covers” the size of strategic externality to the extent that we may not even identify the sign of the strategic interaction, i.e. whether the agents are playing a game of strategic substitutes or complements. We show that the inability to identify the sign of the strategic interaction is due to the quest for robust identification. Namely, under the Bayes Nash equilibrium concept and a specific information structure, the sign, in fact the size of the strategic interaction is point identified. But, the identification of the sign of the strategic interaction fails under the Bayes correlated equilibrium.

For a short moment, we consider the decision-making of an individual agent without any element of strategic interaction. We therefore set $\gamma_{aA} = 0$ and in the absence of any strategic interaction it is sufficient to consider a one-dimensional private signal $x_i = \theta + \varepsilon_i$ about the state of the world θ , as defined earlier in (6). The best response of the agent, as a function of the signal x_i is given by:

$$a_i = -\frac{\lambda_a + \gamma_{a\theta}\mathbb{E}[\theta | x_i]}{\gamma_a}. \quad (49)$$

The best response of the agent, given the precision of the random variables θ and x_i is therefore given by:

$$a_i(x_i) = -\frac{\lambda_a}{\gamma_a} - \frac{\gamma_{a\theta}}{\gamma_a} \frac{x_i \tau_x + \mu_\theta \tau_\theta}{\tau_x + \tau_\theta}. \quad (50)$$

In contrast to the complete information environment, agent i responds to the signal x_i rather than to the state of the world θ . As the signal x_i represents noisy information about θ , the weight that the agent attaches to the signal is only:

$$\left| -\frac{\gamma_{a\theta}}{\gamma_a} \frac{\tau_x}{\tau_x + \tau_\theta} \right| < \left| -\frac{\gamma_{a\theta}}{\gamma_a} \right|,$$

and hence the response of the agent is attenuated under uncertainty about θ . As the agent’s best response to θ is characterized by the slope $-\frac{\gamma_{a\theta}}{\gamma_a}$, it follows that we can learn at most the ratio $-\frac{\gamma_{a\theta}}{\gamma_a}$, but not the individual terms γ_a and $\gamma_{a\theta}$.

Throughout this subsection we assume that the choice of the agent is observed with noise and that the realized state of the world θ is observed without noise. The error in the observation of the action a is denoted u , it has zero mean and is uncorrelated with θ and ε_i . The noisy observation is then given by a' :

$$a' = a + u. \quad (51)$$

The covariance between the observed random variable a' and θ estimates the agent's responsiveness to the underlying state of the world: $\widehat{\frac{\gamma_{a\theta}}{\gamma_a}}$. The estimate of the responsiveness of the agent to the state of the world confounds his true responsiveness, represented by the ratio $\frac{\gamma_{a\theta}}{\gamma_a}$, and the quality of his private information, represented by the ratio $\frac{\tau_x}{\tau_x + \tau_\theta}$ of the precisions:

$$\widehat{\frac{\gamma_{a\theta}}{\gamma_a}} = -\frac{\gamma_{a\theta}}{\gamma_a} \frac{\tau_x}{\tau_x + \tau_\theta}. \quad (52)$$

Thus, the noise in the predictor variable induces a bias, commonly referred to as “attenuation bias”. If the information structure of the agent were known, and hence the ratio $\frac{\tau_x}{\tau_x + \tau_\theta}$ were known to the econometrician, then he could recover from the estimate $\widehat{\frac{\gamma_{a\theta}}{\gamma_a}}$ the true responsiveness of the agent's preferences. But in the absence of precise knowledge of the information structure, we cannot identify the responsiveness $-\frac{\gamma_{a\theta}}{\gamma_a}$, but rather a set of possible ratios:

$$\frac{\gamma_{a\theta}}{\gamma_a} \in [\widehat{\frac{\gamma_{a\theta}}{\gamma_a}}, \infty). \quad (53)$$

In other words, without a specific hypothesis about the information structure of the agent, namely τ_x , we can only identify a set of possible parameters, given by (53). The set of parameters that we can identify in the absence of precise knowledge of the information structures is what we refer to as robust identification. Now, while the set is large, it allows us to identify the sign of the informational externality, $\gamma_{a\theta}$. However, as we move to the strategic environment with $\gamma_{aA} \neq 0$, we shall see that the set identification of $\gamma_{a\theta}$ is going to prevent the sign identification of the strategic interaction γ_{aA} .

7.2 Robust Identification with Individual Data

We now consider the strategic environment and assume for the remainder of this subsection that the econometrician observes the individual actions a_i and the state θ . In other words, the econometrician learns the first and second moment of the joint equilibrium distribution over actions and state: $(\mu_a, \sigma_a, \rho_a, \sigma_\theta, \rho_{a\theta})$. We begin the identification analysis under the hypothesis of Bayes Nash equilibrium and a given information structure τ of the agents. We find that the slope of the equilibrium response and the sign of the interaction parameters are identified for every information structure of the game.

Proposition 12 (Identification in BNE)

The Bayes Nash equilibrium with information structure τ ,

1. *identifies the sign of the informational externality $\gamma_{a\theta}$;*
2. *identifies the sign of the strategic interaction γ_{aA} if $0 < \tau_x, \tau_y < \infty$; and*
3. *identifies the ratios, $\gamma_{a\theta}/(\gamma_a + \gamma_{aA})$ and $\lambda_a/(\gamma_a + \gamma_{aA})$.*

The previously noted lack of point identification of the individual terms $\gamma_a, \gamma_{a\theta}$ and γ_{aA} remains in the Bayes Nash equilibrium. In particular, the marginal cost γ_a of the private action scales all the relevant equilibrium conditions, and hence prevents point identification of the interaction parameters $\gamma_{a\theta}$ and γ_{aA} .

We contrast the point identification for any specific information structure with the set identification in the Bayes correlated equilibrium. We do not make a specific hypothesis regarding the information structure of the agents, and

ask what we learn from the data in the absence of specific knowledge of the information structure. Now, from the observation of the covariance $\rho_{a\theta}\sigma_a\sigma_\theta$ and the observation of the aggregate variance $\rho_a\sigma_a^2$, we can identify the values of $\rho_{a\theta}$ and ρ_a . The equilibrium conditions which tie the data to the structural parameters are given by the following conditions on mean and variance:

$$\mu_a = -\frac{\lambda_a + \mu_\theta\gamma_{a\theta}}{\gamma_a + \gamma_{aA}}, \quad (54)$$

and

$$\sigma_a = -\frac{\sigma_\theta\gamma_{a\theta}\rho_{a\theta}}{\rho_a\gamma_{aA} + \gamma_a}. \quad (55)$$

We thus have two restrictions to identify the four unknown structural parameters $(\lambda_a, \gamma_a, \gamma_{a\theta}, \gamma_{aA})$. We can solve for two of the unknowns in terms of the remaining unknowns. In particular, when we solve for $(\lambda_a, \gamma_{a\theta})$ in terms of the remaining unknowns, γ_a and γ_{aA} , we obtain expressions for the equilibrium intercept and the equilibrium slope in terms of the moments and the remaining unknown structural parameters:

$$-\frac{\lambda_a}{\gamma_a + \gamma_{aA}} = \mu_a - \frac{\sigma_a\mu_\theta(\gamma_a + \rho_a\gamma_{aA})}{\sigma_\theta(\gamma_a + \gamma_{aA})\rho_{a\theta}}, \quad (56)$$

and

$$-\frac{\gamma_{a\theta}}{\gamma_a + \gamma_{aA}} = \frac{\sigma_a}{\rho_{a\theta}\sigma_\theta} \frac{\gamma_a + \rho_a\gamma_{aA}}{\gamma_a + \gamma_{aA}}. \quad (57)$$

Now, except for the case of $\rho_a = 1$, in which the agents are all perfectly correlated, we find that the ratio on the lhs is not uniquely determined. As the strategic interaction parameter γ_{aA} can vary, or $\gamma_{aA} \in (-\infty, -\gamma_a)$, it follows that we can only partially identify the above ratios, namely,

$$-\frac{\lambda_a}{\gamma_a + \gamma_{aA}} \in \left(-\infty, \mu_a - \frac{\rho_a\sigma_a}{\rho_{a\theta}\sigma_\theta}\right),$$

and

$$-\frac{\gamma_{a\theta}}{\gamma_a + \gamma_{aA}} \in \left(\frac{\rho_a\sigma_a}{\rho_{a\theta}\sigma_\theta}, \infty\right),$$

which describes the respective sets into which each ratio can be identified.

Proposition 13 (Partial Identification in BCE)

The Bayes correlated equilibria:

1. *identify the sign of the informational externality $\gamma_{a\theta}$;*
2. *do not identify the sign of the strategic interaction γ_{aA} ;*
3. *identify a set of response ratios, if $\rho_a < 1$:*

$$-\frac{\gamma_{a\theta}}{\gamma_a + \gamma_{aA}} \in \begin{cases} \left(\frac{\rho_a\sigma_a}{\rho_{a\theta}\sigma_\theta}, \infty\right) & \text{if } \frac{\sigma_a}{\rho_{a\theta}\sigma_\theta} > 0; \\ \left(-\infty, \frac{\rho_a\sigma_a}{\rho_{a\theta}\sigma_\theta}\right) & \text{if } \frac{\sigma_a}{\rho_{a\theta}\sigma_\theta} < 0. \end{cases} \quad (58)$$

Thus, in comparison to the Bayes Nash equilibrium, the Bayes correlated equilibrium, weakens the possibility of identification in two respects. First, we fail to identify the sign of the strategic interaction γ_{aA} ; second, we can identify only a set of possible interaction ratios. It is then natural to ask whether prior information can improve the

identification of the structural parameters, just as prior information could improve the equilibrium prediction. This is the content of the next propositions.

The econometrician observes the following data (or moments), denoted by $m = (\sigma_a, \sigma_\theta, \rho_a, \rho_\theta, \mu_a, \mu_\theta)$. In addition, the econometrician is assumed to have prior information about the information structure $\tau = (\tau_x, \tau_y)$. We investigate how the identification of the sign of γ_{aA} and the point identification of the interaction ratio $\frac{\gamma_{a\theta}}{\gamma_a + \gamma_{aA}}$ depends on the prior information. We denote the lower and upper bound of the identified set for the equilibrium slope $\frac{\gamma_{a\theta}}{\gamma_a + \gamma_{aA}}$ by $\underline{\gamma}(\tau, m)$ and $\bar{\gamma}(\tau, m)$, respectively. The bounds depend naturally on the prior information and the observed data. Similarly, we denote the lower and upper bound of the identified set for the strategic interaction γ_{aA} by $\underline{\gamma}_{aA}(\tau, m)$ and $\bar{\gamma}_{aA}(\tau, m)$, respectively. Now, the rhs in the identification of the equilibrium slope in (57) indicates that the range of the identified set depends on the observed data and the ratio

$$\frac{\gamma_a + \rho_a \gamma_{aA}}{\gamma_a + \gamma_{aA}} = \frac{1 + \rho_a r_A}{1 + r_A}.$$

The observed data, in particular $\rho_{a\theta}$ and ρ_a , determine a point in the unit square which describes the location of the correlation coefficients $(\rho_a, \rho_{a\theta})$. In Section 6, we analyzed the impact of the prior information on the possible equilibrium coefficients. As $\rho_a < 1$, the value of this ratio is decreasing in r_A , and hence the smallest possible value of r_A , consistent with the data, provides the upper bound for the above ratio. Now, the analysis in Proposition 11 shows that this lower bound on r_A is obtained through the public equilibrium set, and conversely, the upper bound on r_A is obtained by the private equilibrium set. We therefore obtain the following result.


Proposition 14 (Equilibrium Slope and Prior Information)

1. The lower bound $\underline{\gamma}(\tau, m)$ is increasing in τ_x and the upper bound $\bar{\gamma}(\tau, m)$ is decreasing in τ_y ;
2. The lower bound and the upper bound converge as the prior information becomes precise:

$$\lim_{\tau_x \uparrow \infty} \underline{\gamma}((\tau_x, \tau_y), m) = \lim_{\tau_y \uparrow \infty} \bar{\gamma}((\tau_x, \tau_y), m) = \frac{\rho_{a\theta} \sigma_a}{\sigma_\theta}.$$

The above statement shows that the identification improves monotonically with the prior information. Figure 7 illustrates how prior information improves the set identification. The x -axis represents the possible values of the slope of the equilibrium response (multiplied by the mean μ_θ of the state θ), whereas the y -axis represents the intercept of the equilibrium response. The observed mean of the equilibrium action restricts the relationship between the slope and the intercept parameter to a one-dimensional line with slope -1 . The shaded blue lines indicate the possible pair of intercept and slope as a function of the observed data. As we increase the precision of the public information, the length of the blue line shrinks (from below), which indicates that the identified set shrinks with an improvement in the lower bound of the information.

Conversely, if we were to increase the lower bound on the private information, then we would impose additional restrictions on the linear relationship from above. We can establish a similar improvement with respect to the sign of the strategic interaction γ_{aA} , namely the set $[\underline{\gamma}_{aA}(\tau, m), \bar{\gamma}_{aA}(\tau, m)]$.

Figure 7: Set Identification and Prior Public Information: $\tau_y > \tau'_y > \tau''_y$ **Proposition 15 (Interaction Sign and Prior Information)**

If $\rho_{a\theta} < 1$ and $\gamma_{aA} \neq 0$, then either there exists $\bar{\tau}_x$ such that for all $0 < \bar{\tau}_x < \tau_x$:

$$0 < \underline{\gamma}_{aA}((\tau_x, \tau_y), m) < \bar{\gamma}_{aA}((\tau_x, \tau_y), m),$$

or there exists $\bar{\tau}_y$ such that for all $0 < \bar{\tau}_y < \tau_y$:

$$\underline{\gamma}_{aA}((\tau_x, \tau_y), m) < \bar{\gamma}_{aA}((\tau_x, \tau_y), m) < 0.$$

The identification results here, in particular the contrast between Bayes Nash equilibrium and Bayes correlated equilibrium, are related to, but distinct from the results presented in Aradillas-Lopez and Tamer (2008). They compare identification results under different solution concepts, most notably level k -rationalizability and Nash equilibrium. In their analysis of an entry game with incomplete information, they document the loss in identification power that arises with a more permissive solution concept, i.e. level k -rationalizability. As we compare Bayes Nash and Bayes correlated equilibrium, we show that the lack of identification is not necessarily due to lack a common prior, as associated with rationalizability, but rather the richness of possible private information structures (but all with a common prior). Interestingly, the entry game in Aradillas-Lopez and Tamer (2008) displayed multiple equilibria even under complete information. In contrast, the class of games we analyze all have a unique equilibrium under either complete or incomplete information, still a more permissive equilibrium concepts leads to a loss in identification.

We should emphasize that the current payoff environment describes a common value environment, i.e. the state of the world is the same for all the agents. In contrast, much of the small, but growing literature on identification in games with incomplete information is concerned with a private value environment, in which the private information of agent i only affects the utility of agent i , as for example in Sweeting (2009), Bajari, Hong, Krainer, and Nekipelov (2010) or Paula and Tang (2010). A second important distinction is that in the above mentioned papers, the identification is about some partial aspect of the utility functions *and* the distribution of the (idiosyncratic) states of the world, whereas the present identification seeks to identify the entire utility function but assumes that the states

of the world are observed by the econometrician.

An alternative exercise would limit the identification to a certain subset of parameters, say the interaction term γ_{aA} , but then identify the distribution of the states of the world rather than observe it. For example, Bajari, Hong, Krainer, and Nekipelov (2010) estimate the peer effect in the recommendation of stocks among stock market analysts in a private value environment. There, the observables are the recommendations of the stock analysts and analyst specific information about the relationship of the analyst to the recommended firm. The present analysis suggest that a similar exercises could be pursued in a common value environment, much like a beauty contest. A natural extension here would be use of the actual performance of the recommended stocks to in fact identify the information structure of the stock analysts.

Finally, in many of the recent contributions the assumption of conditional independence of the private information, relative to public observables, is maintained. For example, in Paula and Tang (2010), the conditional independence assumption is used to characterize the joint action equilibrium distribution in terms of the marginal probabilities of every action. Paula and Tang (2010) uses the idea that if private signals are i.i.d. across individuals, then the players actions must be independent in a single equilibrium, “but correlated when there are multiple equilibria” to provide a test for multiple equilibria. In contrast, in our model, we have uniqueness of the Bayes Nash equilibrium, but the unobserved information structure of the agents could lead to correlation, which would then be interpreted in the above test as evidence of multiple equilibria, but could simply be due to the unobserved correlation rather than multiplicity of equilibria.

In a different context, the issue of identification also arises in Saint-Paul (2011). He considers a macroeconomic model in which the government (or the public) only observes a subset of the relevant random variables. Due to the limited information of the government, a difference between the true model and the perceived model may arise. A self-interested adviser (with knowledge of the true model) may therefore suggest a socially optimal policy for the perceived model, that while consistent with the perceived model actually maximizing the biased objective function of the adviser within the true model.

7.3 Robust Identification with Aggregate Data: Demand and Supply

Next, we consider the possibility of robust identification in the context of demand and supply. We consider the linear model of demand and supply. We assume that the (inverse) demand function of the representative consumer is given by:

$$P_d = \lambda_d + \gamma_{ad}A + \gamma_d\theta_d, \quad (59)$$

and that the inverse supply function of the industry (and the representative firm) is given by:

$$P_s = \lambda_s + \gamma_{as}A + \gamma_s\theta_s. \quad (60)$$

The demand and supply functions (59) and (60) can, as usual, be generated by quadratic utility and cost functions respectively. The supply of the individual firm is given by a_i and the aggregate supply is given by A . The canonical analysis of the demand and supply identification uses demand as well as cost shocks. In the interest of space, here we

shall consider only a common supply shock.³ An important aspect of the demand and supply identification is that it canonically uses only aggregate data, namely the aggregate quantity, rather than the individual supply data, but then uses a second source of data, namely the aggregate price, to facilitate the identification. The present identification analysis adheres to this convention. We can rewrite the demand and supply function in terms of the general notation adopted in Section 2 by setting:

$$\gamma_d \triangleq 0, \gamma_s \triangleq \gamma_{a\theta}, \gamma_{as} \triangleq \gamma_a, \gamma_{ad} \triangleq \gamma_{aA}. \quad (61)$$

In contrast to the classic identification literature, which considers complete information environments, we are interested in an environment with incomplete information. The state of the world now describes the supply shock, $\theta = \theta_s$, say the level of an input price, and remains normally distributed $\theta \sim N(\mu_\theta, \sigma_\theta^2)$. The incomplete information arises as each firm has to make its supply decision, a_i , on the basis of private, but possibly noisy information, about the common cost shock. The market is cleared by the price which results from the supplied aggregate quantity A and the level of demand. We determine whether the demand and supply function can be identified under incomplete information, and more precisely whether it can be robustly identified, independent of the knowledge of the structure of the private information of the agents.

We begin the analysis with the Bayes Nash equilibrium as a benchmark. We gave a complete analysis of the Bayes Nash equilibrium in Section 4 and use it for the present identification result. We maintain the assumption that the information structure, τ , but not the signal realizations (x_i, y) are known to the econometrician, and hence can be used to identify the structure of demand and supply. Now, we observed earlier that the variance of the aggregate action, here the aggregate supply, and the covariance with the observed cost shock is sufficient to identify the linear coefficients of the Bayes Nash equilibrium strategy. With the separate observation of the equilibrium price P , we can point identify the demand parameter γ_{aA} . The variance and covariance then allow us to point identify the supply parameters, $\gamma_{a\theta}$ and γ_a .

Proposition 16 (Point Identification in BNE)

The demand function, γ_{aA} and the supply function $(\gamma_a, \gamma_{a\theta})$, is point identified for every information structure $0 < \tau_x, \tau_y < \infty$.

We note that the above identification result uses the positive noise in the private and public signal to distinguish the demand parameter γ_{aA} from the supply parameter γ_a . The supply of the individual firm is more responsive to the private signal than to the public signal in this game with substitutes. After all, the private signals are less correlated across agents than the public signal. In an environment with complete information, where all the agents receive the same noiseless signal about the cost shock, the parameters of the supply function, γ_a and $\gamma_{a\theta}$, could not be identified separately. Of course, the slope of the equilibrium response would still be point-identified.

Next, we consider the possibility of robust identification. As in the Bayes Nash equilibrium, the presence of the equilibrium price data allows us to point identify the demand parameter γ_{aA} . The identification of the supply function is a different matter. In the presence of the aggregate data, we can only learn from the aggregate variance σ_A^2

³The present analysis can be extended to the canonical model with demand *and* cost shocks. Here, we do not pursue the analysis of the associated model with a two-dimensional state $\theta = (\theta_d, \theta_s)$, as this extension, while entirely straightforward, is substantially more elaborate as we would have to keep track of the additional interaction and covariance terms.

and the covariance with the cost shock, $\rho_{A\theta}\sigma_A\sigma_\theta$, which allows us to identify the correlation coefficient $\rho_{A\theta}$. But in contrast to the Bayes Nash equilibrium, where we could infer from the aggregate data the behavior of the individual choices, in particular the correlation coefficients ρ_a and $\rho_{a\theta}$, we now have to contend with additional limitations. In Proposition 7 we established a relationship between the individual choices and the aggregate choices, and the aggregate covariance $\rho_{A\theta}$ restricts the idiosyncratic correlation coefficients:

$$\rho_{A\theta} = \frac{\rho_{a\theta}}{\sqrt{\rho_a}}. \quad (62)$$

In other words, the data, in particular $\rho_{A\theta}$, allows us to identify a curve in the unit square of the correlation coefficients $(\rho_a, \rho_{a\theta})$, but not more. In turn, the observation of the mean and variance of the aggregate supply imposes restriction on the parameters of the supply function, γ_a and $\gamma_{a\theta}$:

$$\mu_A = -\frac{\lambda_a + \mu_\theta\gamma_{a\theta}}{\gamma_a + \gamma_{aA}}, \quad (63)$$

and

$$\sigma_A = -\frac{\sigma_\theta\gamma_{a\theta}\rho_{A\theta}\sqrt{\rho_a}}{\rho_a\gamma_{aA} + \gamma_a}, \quad (64)$$

yet these two restrictions are not sufficient to point identify either γ_a and $\gamma_{a\theta}$, or even the slope of the equilibrium supply, $\gamma_{a\theta}/(\gamma_a + \gamma_{aA})$. In fact by inspecting (63) and (64), we find that we face three unknown structural parameters $(\lambda_a, \gamma_a, \gamma_{a\theta})$ and an unknown endogenous parameter, ρ_a .

Proposition 17 (Partial Identification in BCE)

The demand function, γ_{aA} , remains point identified, but the supply function, $(\gamma_a, \gamma_{a\theta})$, is not point identified.

The partial identification is notably asymmetric. The demand structure is completely identified, but the supply structure is only partially identified. This asymmetry is due entirely to the asymmetry in which the private information enters the equilibrium variables. We assumed that the firms make their supply decision on the basis of their private information about the cost shocks, but that the equilibrium price, and hence the consumer's decision was made after the shock had been realized. Thus, the variation in the equilibrium price, conditional on the supplied quantity, was identified without the interference of the private information. The identification problem would become more symmetric if instead we were to assume that the consumer were to make their decision under private information as well, and hence before the realization of the demand shock. We did not pursue this more elaborate information environment here. Yet, if we were to pursue it, the basic insight, namely that the private information renders the identification incomplete would now apply a fortiori, and in particular it would render the demand structure partially identified as well.

Given the observed variance of the aggregate action, the analyst cannot know whether the variance in the supply should be attributed to the marginal cost, γ_a or to the sensitivity with respect to cost shock, $\gamma_{a\theta}$. In particular, if the supply is very sensitive to the cost shock, then a given variance of the aggregate quantity, must be explained by a sufficiently large marginal cost to dampen the response of the firm. Using the above restriction on mean and variance, we know that the equilibrium response is given by:

$$\frac{\gamma_{a\theta}}{\gamma_a + \gamma_{aA}} = -\frac{\sigma_A}{\sigma_\theta\rho_{A\theta}} \frac{1}{\sqrt{\rho_a}} \frac{\gamma_a + \gamma_{aA}\rho_a}{\gamma_a + \gamma_{aA}}. \quad (65)$$




Figure 8: Identified set of strategic interaction term r_A as a function of prior information with precision $\tau_x = \tau_y = \tau$

The linear condition (65) is similar to the identification issue we faced with individual data in (57). But there, we could at least infer the correlation coefficient ρ_a from the individual data. By contrast, in the context of demand and supply, we are restricted to use aggregate observables, and hence do not know ρ_a . The equilibrium slope is therefore only set-identified as given the limited data, we can only assert that:


$$\frac{\gamma_{a\theta}}{\gamma_a + \gamma_{aA}} \in -\frac{\sigma_A}{\sigma_\theta \rho_{\theta A}} \cdot \left(\sqrt{\rho_a}, \frac{1}{\sqrt{\rho_a}} \right) \Leftrightarrow \frac{\gamma_{a\theta}}{\gamma_a + \gamma_{aA}} \in -\frac{\sigma_A}{\sigma_\theta \rho_{\theta A}} \cdot (0, \infty), \quad (66)$$

where the second inclusion follows from the fact that we cannot infer the correlation coefficient ρ_a from the aggregate data.

Now, the interaction between the correlation coefficient ρ_a and the parameter γ_a of the supply function in (65) suggests that prior information may substantially improve the identification result, just as it was shown previously with the individual data. Earlier we noticed that the lower bounds on the private information imposed lower bounds on the correlation coefficients ρ_a and $\rho_{a\theta}$, see Proposition 10. Now, the aggregate data only leads to information about the ratio of these two coefficients, but not their size. Thus, with the lower bounds we narrow the possible range of the individual behavior, and hence the possible range of the correlation coefficients. At the same time, we saw that the private and public equilibrium sets respond differently to a change in the parameter of the supply function. In Proposition 11 we showed that as we increase γ_a , the private equilibrium set shrinks, whereas the public equilibrium set grows.

Now, for a given pair of correlation coefficients $(\rho_a, \rho_{a\theta})$, consistent with (62), we ask how an increase in the prior information improves the identification. Figure 8 illustrates this with a symmetric increase in the precision of the public and private signal, or $\tau_x = \tau_y$. With zero prior information, we see that we cannot restrict the value of the interaction ratio. But as we increase the precision of the prior information, the identified set shrinks rapidly until it is point identified. Now for a given pair of correlation coefficients in the interior of the unit square, or $0 < \rho_a, \rho_{a\theta} < 1$, eventually the increase in the precision would require too large a correlation to be consistent with the observed data, which explains the collapse of the identified set beyond a specific value of the precision τ .

Now, with aggregate data, an entire range of correlation coefficients $(\rho_a, \rho_{a\theta})$ is consistent with the aggregate correlation coefficient $\rho_{A\theta}$. This is illustrated in Figure 9 where we fix the aggregate correlation coefficient $\rho_{A\theta}$ and

Figure 9: Identified set of γ_a with aggregate data $\rho_{A\theta}$

a particular information structure τ and ask what values of the interaction ratio r_A are consistent with a particular $\rho_{a\theta}$ and ρ_a . For a given $\rho_{A\theta}$, we plot on the x -axis the possible values of $\rho_{a\theta}$ (which then implicitly defines the corresponding values of ρ_a). Then given the information structure τ , every vertical segment between the lower boundary defined through the public equilibrium set and the upper boundary defined through the private equilibrium set identifies the set of interaction ratios consistent with the data. The entire range of interaction ratios consistent with the observed data is therefore given by the lower bound and the upper bound on the interaction ratio across all $\rho_{a\theta}$ with a nonempty intersection. While in general the interaction term r_A can be negative, given that the demand is point identified with a negative slope, $\gamma_{aA} < 0$, it follows that the values of the interaction term $r_A = \frac{\gamma_{aA}}{\gamma_a}$ have to be positive everywhere. The shaded area in between the lower and upper boundary, subject to the nonnegativity condition then presents the identified set.

We denote the lower and upper bound of the identified set for the strategic interaction γ_a by $\underline{\gamma}_a(\tau, m)$ and $\bar{\gamma}_a(\tau, m)$, respectively. The bounds depend naturally on the prior information and the observed data. Earlier we introduce the lower and upper bounds of the identified set for the equilibrium response $\frac{\gamma_{a\theta}}{\gamma_a + \gamma_{aA}}$ as $\underline{\gamma}(\tau, m)$ and $\bar{\gamma}(\tau, m)$, respectively.

Proposition 18 (Prior Information and Identification)

1. The lower bound $\underline{\gamma}_a(\tau, m)$ increases in the precision τ_y and the upper bound $\bar{\gamma}_a(\tau, m)$ decreases in the precision τ_x .
2. The lower bound $\underline{\gamma}(\tau, m)$ increases in the precision τ_x and the upper bound $\bar{\gamma}(\tau, m)$ decreases in the precision of τ_y .
3. The lower bounds and the upper bounds converge as the prior information becomes precise:

$$\lim_{\tau_x \uparrow \infty} \underline{\gamma}_a(\tau, m) = \lim_{\tau_y \uparrow \infty} \partial \bar{\gamma}_a(\tau, m) = \gamma_a, \quad \lim_{\tau_x \uparrow \infty} \underline{\gamma}(\tau, m) = \lim_{\tau_y \uparrow \infty} \partial \bar{\gamma}(\tau, m) = \frac{\gamma_{a\theta}}{\gamma_a + \gamma_{aA}}.$$

We analyzed the identification problem in the demand and supply model under the restriction that we only observed the aggregate supply data rather than the individual supply data. In the case of the Bayes Nash equilibrium,

we showed that the restriction to aggregate data is actually without loss of identification power. However, in the case of the Bayes Correlated equilibrium, additional information in terms of the individual supply data would ease the identification problem. The additional information would clearly help, as the description of the identified set for the equilibrium slope, see (66), indicated. With the individual supply data, we could infer ρ_a and $\rho_{a\theta}$, rather than just the ratio of the two, which is contained in the aggregate correlation coefficient $\rho_{A\theta}$, see (62).

8 Conclusion

It was the objective of this paper to derive robust equilibrium predictions for a large class of games. We began with an epistemic result that related the class of Bayes Nash equilibria with the class of Bayes correlated equilibria. The equivalence results allowed us to focus on the characterization of the Bayes correlated equilibria which proceeded without reference to a specific information structure held by the agents. Within a class of quadratic payoff environments, we gave a full characterization of the equilibria in terms of moment restrictions on the equilibrium distributions. The robust analysis allowed us to make equilibrium predictions independent of the information structure, the nature of the private information that the agents might have access to.

We then reversed the point of view and considered the problem of identification rather than the problem of prediction. We asked what are the implication of a robust point of view for identification, namely the ability to infer the unobservable structural parameters of the game from the observable data. Here we showed that in the presence of robustness concerns, the ability to identify the underlying parameters of the game is weakened in important ways, yet does not completely eliminate the possibility of identification. The current perspective, namely to analyze the set of correlated equilibria rather than the Bayes Nash equilibria under a specific information structure, is potentially useful in the emerging econometric analysis of games of incomplete information. There the identification question is typically pursued for a given information structure, say independently distributed payoff types, and it is of interest to know how sensitive the identification results are to the structure of the private information. In this context, the robust identification might be particularly important as we rarely observe data about the nature of the information structure directly.

We considered a continuum of action spaces, but binary or discrete action games might be of interest as well. A natural example would be voting games where the unknown state of the world is the quality of the candidate, and we could then examine what the probability of a qualified candidate being elected is and how it depends on the information structure or the election regime being adopted. We would then have robust policy implications for the voting rules.

In the present analysis, we use the structure of the quadratic payoffs, in particular the linear best response property to derive the first and second moments of the correlated equilibrium set. A natural next step would be to bring the present analysis to Bayesian games with nonlinear strategies. For example, it would be of considerable interest to ask how the allocations and the revenues in the first price auction differ across belief environments.

We would also like to use the equilibrium predictions to offer robust versions of policy and welfare analysis. In many incomplete information environments, a second best or otherwise welfare improving policy typically relies on and is sensitive to the specification of the belief environment. With the current analysis, we might be able to

recommend robust taxation or information disclosure policies which are welfare improving across a wide range of belief environments. In particular, we might ask how the nature of the policy depends on the prior information of the policy maker about the belief environment of the agents.

9 Appendix

Proof of Proposition 2. The characterization of the linear Bayes Nash equilibrium strategy in the same environment of quadratic payoffs is given in Proposition 1 of Angeletos and Pavan (2007), the only difference being a different labeling of the linear return and interaction parameters, λ and Γ . ■

Proof of Proposition 7. The covariance $cov(a_i, a_j)$ is given by, using the expansion suggested by (38):

$$cov(a_i, a_j) = cov(a_i - A + A, a_j - A + A) = cov(A, A) = \rho_a \sigma_a^2.$$

Similarly, the covariance $cov(a_i, \theta)$ is given by, using the expansion suggested by (38):

$$cov(a_i, \theta) = cov(a_i - A + A, \theta) = cov(\theta, A) = \rho_{a\theta} \sigma_a \sigma_\theta,$$

and all of the results follow directly from rearranging the relevant expressions. ■

Proof of Proposition 5. Given the distribution (2) we can compute the expected values of θ and A given a . The standard formulae for the conditional distributions lead to:

$$\mathbb{E}[\theta|a] = \left(1 + \frac{\rho_{a\theta}\sigma_\theta}{\sigma_a} \frac{\gamma_{a\theta}}{\gamma_a + \gamma_{aA}}\right) \mu_\theta + \frac{\rho_{a\theta}\sigma_\theta}{\sigma_a} a, \quad (67)$$

and

$$\mathbb{E}[A|a] = -\frac{\gamma_{a\theta}}{\gamma_a + \gamma_{aA}} \mu_\theta (1 - \rho_a) + \rho_a a. \quad (68)$$

The optimality property of the best response function then requires:

$$1 = -\left(\frac{\rho_{a\theta}\sigma_\theta}{\sigma_a} \frac{\gamma_{a\theta}}{\gamma_a} + \rho_a \frac{\gamma_{aA}}{\gamma_a}\right) \Leftrightarrow \sigma_a = -\frac{\sigma_\theta \gamma_{a\theta} \rho_{a\theta}}{\rho_a \gamma_{aA} + \gamma_a}. \quad (69)$$

We observe that the conditional expectation is only valid if $\sigma_a > 0$. In contrast for $\sigma_a = 0$, we have a unique Bayes correlated equilibrium which is equal to the complete information correlated equilibrium with $\theta = \mu_\theta$.

We can then insert the condition (69) into the normal distribution (22) and as the variance-covariance matrix has to be a nonnegative definite matrix, a necessary and sufficient condition is that the determinant of the variance-covariance matrix is nonnegative, or,

$$-\sigma_\theta^6 \rho_{a\theta}^4 \gamma_{a\theta}^4 (\rho_a - 1) \frac{\rho_a - \rho_{a\theta}^2}{(\gamma_a + \rho_a \gamma_{aA})^4} \geq 0 \Rightarrow \rho_a - \rho_{a\theta}^2 \geq 0,$$

which concludes the proof. ■

Proof of Proposition 6. The variance σ_a^2 is given by (3), and inserting $\rho_a = \rho_{a\theta}^2$ we obtain:

$$-\frac{\sigma_\theta \gamma_{a\theta} \rho_{a\theta}}{\rho_{a\theta}^2 \gamma_{aA} + \gamma_a}$$

which is maximized at:

$$|\rho_{a\theta}| = \sqrt{\frac{\gamma_a}{\gamma_{aA}}},$$

or $\rho_a = \gamma_a / \gamma_{aA}$. ■

Proof of Proposition 8. (1.) The volatility σ_A^2 , which is given by:

$$\rho_a \sigma_a^2 = \rho_a \left(-\frac{\sigma_\theta \gamma_{a\theta} \rho_{a\theta}}{\rho_a \gamma_{aA} + \gamma_a} \right)^2,$$

is increasing in the correlation coefficients ρ_a and $|\rho_{a\theta}|$. The partial derivatives with respect to ρ_a and $|\rho_{a\theta}|$ are:

$$\frac{\sigma_\theta^2 \rho_a^2 \gamma_{a\theta}^2}{(\gamma_a + \rho_a \gamma_{aA})^3} (\gamma_a - \rho_a \gamma_{aA}) > 0,$$

and

$$\frac{2r |\rho_{a\theta}| \sigma_\theta^2 \gamma_{a\theta}^2}{(\gamma_a + \rho_a \gamma_{aA})^2} > 0,$$

respectively.

(2.) The dispersion σ_{a-A}^2 , using (??), is given by:

$$(1 - \rho_a) \sigma_a^2 = (1 - \rho_a) \left(-\frac{\sigma_\theta \rho_{a\theta} \gamma_{a\theta}}{\rho_a \gamma_{aA} + \gamma_a} \right)^2, \quad (70)$$

and it follows that the dispersion is increasing in $|\rho_{a\theta}|$. The dispersion σ_{a-A}^2 is monotone decreasing in ρ_a if it is game of strategic substitutes, and not necessarily monotone if it is a game of strategic complements. The partial derivative with respect to ρ_a is given by

$$-\frac{\sigma_\theta^2 \rho_a^2 \gamma_{a\theta}^2 (\gamma_a + \gamma_{aA} + (1 - \rho_a) \gamma_{aA})}{(\gamma_a + \rho_a \gamma_{aA})^3}.$$

However by Proposition 5, it follows that $\rho_{a\theta}^2 \leq \rho_a$, and we therefore obtain the maximal dispersion at $\rho_{a\theta}^2 = \rho_a$. Consequently, we have

$$(1 - \rho_a) \sigma_a^2 = (1 - \rho_a) \rho_a \left(-\frac{\sigma_\theta \gamma_{a\theta}}{\rho_a \gamma_{aA} + \gamma_a} \right)^2,$$

and the dispersion reaches an interior maximum at

$$\rho_a = \frac{\gamma_a}{2\gamma_a + \gamma_{aA}} \in (0, 1),$$

irrespective of the nature of the game. ■

Proof of Proposition 9. Every Bayes Nash equilibrium with an information structure τ is clearly a Bayes correlated equilibrium. It remains to establish that every Bayes correlated equilibrium can be replicated by a Bayes Nash equilibrium for some information structure τ . By Proposition 5, the set of Bayes correlated equilibria is completely characterized by the pair of correlation coefficients $(\rho_a, \rho_{a\theta})$ with

$$0 \leq \rho_a - \rho_{a\theta}^2 \leq 1, \quad (71)$$

and

$$0 \leq \rho_a \leq 1. \quad (72)$$

The correlation coefficients ρ_a and $\rho_{a\theta}$ of the Bayes Nash equilibrium can be expressed in terms of the equilibrium coefficients α_x and α_y and variances σ_θ^2 , σ_x^2 and σ_y^2 as:

$$\rho_{a\theta} = \frac{\sigma_\theta (\alpha_x + \alpha_y)}{\sqrt{\alpha_x^2 \sigma_x^2 + \alpha_y^2 \sigma_y^2 + \sigma_\theta^2 (\alpha_x + \alpha_y)^2}}, \quad (73)$$

and

$$\rho_a = \frac{\alpha_y^2 \sigma_y^2 + \sigma_\theta^2 (\alpha_x + \alpha_y)^2}{\alpha_x^2 \sigma_x^2 + \alpha_y^2 \sigma_y^2 + \sigma_\theta^2 (\alpha_x + \alpha_y)^2}. \quad (74)$$

It now follows immediately from (73) and (74) that the inequalities of (33) are satisfied for all σ_x^2 and σ_y^2 . It remains to show that all pairs $(\rho_a, \rho_{a\theta})$ satisfying (71) and (72) can be generated by some (τ_x, τ_y) . We observe that relative to $\rho_{a\theta}$, the correlation coefficient ρ_a has the additional term $\alpha_y^2 \sigma_y^2$. It then follows by the equality of (19):

$$\alpha_y^* = \alpha_x^* \frac{\sigma_x^2}{\sigma_y^2} \left(\frac{\gamma_a}{\gamma_a + \gamma_{aA}} \right),$$

or

$$\alpha_y^2 \sigma_y^2 = \alpha_x^2 \sigma_x^4 \sigma_y^2 \left(\frac{\gamma_a}{\gamma_a + \gamma_{aA}} \right)^2,$$

that we can choose (τ_x, τ_y) so that $\rho_{a\theta}$ stays constant but that ρ_a varies freely between $\rho_{a\theta}^2$ and 1. Conversely, we can choose ρ_a fixed and vary $\rho_{a\theta}$ between 0 and $\sqrt{\rho_a}$. ■

Proof of Proposition 10. We form the conditional expectation using (40) and the equilibrium conditions for the Bayes correlated equilibrium are then given by (42) and the solution to theses equations is given by (43)-(45).

(1.) The equilibrium set is described as the set which satisfies the inequalities (46) and (47), where the correlation coefficients ρ_{ax}^2 and ρ_{ay}^2 appear separately. By determination of (44) and (45), the square of the correlation coefficient is strictly decreasing in σ_x and σ_y , which directly implies that the respective inequalities become less restrictive, and hence the equilibrium set increases as either σ_x or σ_y increase.

(2.) The lowest value of the correlation coefficient ρ_a is achieved when the inequality (46) is met as an equality. It follows that the minimum is reached at the exterior of the equilibrium set. The equilibrium set is increasing in σ by the previous argument in (1), and hence the resulting strict inequality.

(3.) The lowest value of the correlation coefficient $\rho_{a\theta}$ is achieved when the inequalities (46) and (47) are met as equalities. It follows that the minimum is reached at the exterior of the equilibrium set. The equilibrium set is increasing in σ by the previous argument in (1), and hence the resulting strict inequality. ■

Proof of Proposition 11. The comparative static results follow directly from the description of the correlation coefficients ρ_{ax} and ρ_{ay} given by (44) and (45). These correlation coefficients are a function of γ_a and γ_{aA} . We insert their solution into the inequalities (46) and (47) and solve for the relation between ρ_a and $\rho_{a\theta}$ as we restrict the the inequalities (46) and (47) to be equalities. ■

Proof of Proposition 12. (1.) Given the knowledge of $\sigma_\theta^2, \sigma_x^2$ and σ_y^2 and the information about the covariates, we can recover the value of the linear coefficients α_x^2 and α_y^2 from variance-covariance matrix (19), say:

$$\alpha_x^2 = \frac{\sigma_a^2 - \sigma_A^2}{\sigma_x^2}, \quad \alpha_y^2 = \frac{\sigma_A^2 (1 - \rho_{A\theta}^2)}{\sigma_y^2}. \quad (75)$$

The value of covariate $\rho_{A\theta} \sigma_A \sigma_\theta$, given by $\sigma_\theta^2 (\alpha_x + \alpha_y)$ directly identifies the sign of the externality $\gamma_{a\theta}$, given the composition of the equilibrium coefficients α_x^* and α_y^* of the Bayes Nash equilibrium, see (16) and (17).

(2.) We have from the description of the Bayes Nash equilibrium in Proposition 2 that in every Bayes-Nash equilibrium, α_x^* and α_y^* satisfy the linear relationship:

$$\alpha_y^* = \alpha_x^* \frac{\sigma_x^2}{\sigma_y^2} \frac{\gamma_a}{\gamma_a + \gamma_{aA}}.$$

Now, if $0 < \sigma_x^2, \sigma_y^2 < \infty$, then we can identify the sign of γ_{aA} .

(3.) Given the identification of α_x^* and α_y^* , we can identify the ratios $\lambda_a \setminus (\gamma_a + \gamma_{aA})$ and $\gamma_{a\theta} \setminus (\gamma_a + \gamma_{aA})$. We recover the mean action μ_a and the coefficients of the linear strategy, i.e. α_x^* and α_y^* , from the equilibrium data. We therefore have three equalities, but four unknown structural parameters $(\lambda_a, \gamma_a, \gamma_{aA}, \gamma_{a\theta})$. From the equilibrium conditions, namely (15), (16), and (17), we have the values of μ_a, α_x and α_y . This allows us to solve for λ_a, γ_{aA} and $\gamma_{a\theta}$ as a function of $\mu_a, \alpha_x, \alpha_y$:

$$\begin{aligned} \lambda_a &= \frac{1}{\alpha_y \sigma^2 \sigma_y^2} (\alpha_x^2 \sigma^2 \gamma_a \sigma_x^2 \mu_\theta + \alpha_x \mu_a \sigma^2 \gamma_a \sigma_x^2 - \alpha_x \alpha_y \sigma^2 \gamma_a \sigma_y^2 \mu_\theta + \alpha_x \alpha_y \gamma_a \sigma_x^2 \sigma_y^2 \mu_\theta) \\ \gamma_{aA} &= \frac{1}{\alpha_y \sigma_y^2} (\alpha_x \gamma_a \sigma_x^2 - \alpha_y \gamma_a \sigma_y^2) \\ \gamma_{a\theta} &= -\frac{1}{\alpha_y \sigma^2 \sigma_y^2} (\alpha_x^2 \sigma^2 \gamma_a \sigma_x^2 - \alpha_x \alpha_y \sigma^2 \gamma_a \sigma_y^2 + \alpha_x \alpha_y \gamma_a \sigma_x^2 \sigma_y^2), \end{aligned} \quad (76)$$

where all the expressions depend on γ_a , and hence are not identified, as γ_a is not identified. But if we form the ratios $\lambda_a \setminus (\gamma_a + \gamma_{aA})$ and $\gamma_{a\theta} \setminus (\gamma_a + \gamma_{aA})$ with the expressions on the rhs of (76), then we obtain expressions which do only depend on the observable data, and are hence point identified, and in particular

$$\frac{\lambda_a}{\gamma_a + \gamma_{aA}} = \frac{(\sigma^2 \mu_a \sigma_x^2 + \sigma^2 \alpha_x \sigma_x^2 \mu_\theta - \sigma^2 \alpha_y \sigma_y^2 \mu_\theta + \sigma_x^2 \alpha_y \sigma_y^2 \mu_\theta)}{\sigma^2 \sigma_x^2}, \quad (77)$$

and

$$\frac{\gamma_{a\theta}}{\gamma_a + \gamma_{aA}} = -\frac{\alpha_x \sigma^2 \sigma_x^2 - \alpha_y \sigma^2 \sigma_y^2 + \alpha_y \sigma_x^2 \sigma_y^2}{\sigma^2 \sigma_x^2}. \quad (78)$$

which completes the proof of identification. We observe that, using (75), we could express the ratios (77) and (78) entirely in terms of the second moments of observed data. ■

Proof of Proposition 13. (1.) From the observation of the covariance $\rho_{a\theta} \sigma_a \sigma_\theta$ we can infer the sign and the size of $\rho_{a\theta}$, see (??). Given the information about lhs of (55) and the information of $\rho_{a\theta}$, we can infer the sign of $\gamma_{a\theta}$.

(2.) Even though the sign of $\gamma_{a\theta}$ can be established, we cannot extract the unknown variables on the lhs of (54) in the presence of the linear return term λ_a , and hence it follows that we cannot sign γ_{aA} .

(3.) From the observation of the covariance $\rho_{a\theta} \sigma_a \sigma_\theta$ and the observation of the aggregate variance $\rho_a \sigma_a^2$, we can infer the value of $\rho_{a\theta}$ and ρ_a . The equilibrium conditions then impose the conditions (54) and (55) on mean and variance. We thus have two equations to identify the four unknown structural parameters $(\lambda_a, \gamma_a, \gamma_{a\theta}, \gamma_{aA})$. We can solve for $(\lambda_a, \gamma_{a\theta})$ in terms of the remaining unknowns γ_a and γ_{aA} to obtain:

$$\lambda_a = \frac{\sigma_a \gamma_a \mu_\theta + \sigma_a \rho_a \mu_\theta \gamma_{aA} - \mu_a \sigma_\theta \gamma_{aA} \rho_{a\theta} - \gamma_a \mu_a \sigma_\theta \rho_{a\theta}}{\sigma_\theta \rho_{a\theta}},$$

and

$$\gamma_{a\theta} = -\frac{\sigma_a \gamma_a + \sigma_a \rho_a \gamma_{aA}}{\sigma_\theta \rho_{a\theta}}.$$

In particular, we would like to know whether this allows us to identify the ratios:

$$-\frac{\lambda_a}{\gamma_a + \gamma_{aA}} = \mu_a - \frac{\sigma_a \mu_\theta (\gamma_a + \rho_a \gamma_{Aa})}{\sigma_\theta (\gamma_a + \gamma_{Aa}) \rho_{a\theta}} \quad (79)$$

or

$$-\frac{\gamma_{a\theta}}{\gamma_a + \gamma_{aA}} = \frac{\frac{1}{\sigma_\theta} \frac{\sigma_a \gamma_a + \sigma_a \rho_a \gamma_{Aa}}{\rho_{a\theta}}}{\gamma_a + \gamma_{aA}} = \frac{\sigma_a}{\rho_{a\theta} \sigma_\theta} \frac{\gamma_a + \rho_a \gamma_{Aa}}{\gamma_a + \gamma_{Aa}} \quad (80)$$

in terms of the observables. But, except for the case of $\rho_a = 1$, we see that this is not the case. As $\gamma_{aA} \in (-\infty, -\gamma_a)$, it follows that we can only partially identify the above ratios, namely,

$$-\frac{\lambda_a}{\gamma_a + \gamma_{aA}} \in \left(-\infty, \mu_a - \frac{\rho_a \sigma_a}{\rho_{a\theta} \sigma_\theta} \right), \quad (81)$$

and

$$-\frac{\gamma_{a\theta}}{\gamma_a + \gamma_{aA}} \in \left(\frac{\rho_a \sigma_a}{\rho_{a\theta} \sigma_\theta}, \infty \right), \quad (82)$$

which describes the respective set into which each ratio can be identified. ■

Proof of Proposition 14. We know from Proposition 13 that the interaction ratio is a function of the observed data and the unobserved interaction parameters:

$$-\frac{\gamma_{a\theta}}{\gamma_a + \gamma_{aA}} = \frac{\sigma_a}{\rho_{a\theta} \sigma_\theta} \frac{\gamma_a + \rho_a \gamma_{Aa}}{\gamma_a + \gamma_{Aa}}.$$

The prior information on the private and public information restricts the possible values of γ_a and γ_{aA} , and hence the values that the above interaction ratio can attain.

We begin the argument with the public equilibrium set which will provide an upper bound on the ratio

$$\frac{\gamma_a + \rho_a \gamma_{Aa}}{\gamma_a + \gamma_{Aa}}. \quad (83)$$

The above ratio appears in the correlation coefficient ρ_{ay} as described in (45). The value of the ratio is maximized when the inequality constraint (47) of the public equilibrium set holds as an equality, and thus

$$\frac{\gamma_a + \rho_a \gamma_{Aa}}{\gamma_a + \gamma_{Aa}} = \rho_{a\theta}^2 + \frac{\sigma_y}{\sigma_\theta} \rho_{a\theta} \sqrt{\rho_a - \rho_{a\theta}^2}, \quad (84)$$

and hence

$$\bar{\gamma}(\tau, m) = \frac{\sigma_a}{\rho_{a\theta} \sigma_\theta} \left(\rho_{a\theta}^2 + \frac{\sigma_y}{\sigma_\theta} \rho_{a\theta} \sqrt{\rho_a - \rho_{a\theta}^2} \right).$$

It follows that if the variance σ_y decreases, then the largest value of the above ratio decreases and as σ_y^2 decreases to zero:

$$\frac{\gamma_a + \rho_a \gamma_{Aa}}{\gamma_a + \gamma_{Aa}} = \rho_{a\theta}^2,$$

which implies that

$$\lim_{\sigma_y \downarrow 0} \bar{\gamma}(\tau, m) = \frac{\sigma_a \rho_{a\theta}}{\sigma_\theta}.$$

Now, consider the private equilibrium set. The ratio (83) does not appear directly in the correlation coefficient ρ_{ax} , rather the structural parameters γ_a and γ_{aA} appear in the form of a ratio. We therefore ask what feasible pair in the private equilibrium set minimizes the ratio (83). The minimal ratio is achieved by a pair (γ_a, γ_{aA}) which solves

the inequality (46) as an equality. We can therefore solve for γ_{aA} in terms of γ_a . The resulting function $\gamma_{aA}(\gamma_a)$ is linear in γ_a :

$$\gamma_{aA} = \gamma_a \frac{\sigma_x \rho_{a\theta} \sqrt{1 - \rho_a} - (1 - \rho_{a\theta}^2) \sigma_\theta}{(\rho_a - \rho_{a\theta}^2) \sigma_\theta}.$$

Hence the minimal value of the ratio (83) does not depend on either γ_a nor γ_{aA} , and is given by:

$$\frac{\gamma_a + \rho_a \gamma_{aA}}{\gamma_a + \gamma_{aA}} = \frac{\gamma_a + \rho_a \left(\gamma_a \frac{\sigma_x \rho_{a\theta} \sqrt{1 - \rho_a} - (1 - \rho_{a\theta}^2) \sigma_\theta}{(\rho_a - \rho_{a\theta}^2) \sigma_\theta} \right)}{\gamma_a + \left(\gamma_a \frac{\sigma_x \rho_{a\theta} \sqrt{1 - \rho_a} - (1 - \rho_{a\theta}^2) \sigma_\theta}{(\rho_a - \rho_{a\theta}^2) \sigma_\theta} \right)} = \frac{(\rho_a - \rho_{a\theta}^2) \sigma_\theta + \rho_a (\sigma_x \rho_{a\theta} \sqrt{1 - \rho_a} - (1 - \rho_{a\theta}^2) \sigma_\theta)}{(\rho_a - \rho_{a\theta}^2) \sigma_\theta + (\sigma_x \rho_{a\theta} \sqrt{1 - \rho_a} - (1 - \rho_{a\theta}^2) \sigma_\theta)}, \quad (85)$$

and hence

$$\underline{\gamma}(\tau, m) = \frac{\sigma_a}{\rho_{a\theta} \sigma_\theta} \frac{(\rho_a - \rho_{a\theta}^2) \sigma_\theta + \rho_a (\sigma_x \rho_{a\theta} \sqrt{1 - \rho_a} - (1 - \rho_{a\theta}^2) \sigma_\theta)}{(\rho_a - \rho_{a\theta}^2) \sigma_\theta + (\sigma_x \rho_{a\theta} \sqrt{1 - \rho_a} - (1 - \rho_{a\theta}^2) \sigma_\theta)}.$$

It is immediate to verify that $\underline{\gamma}(\tau, m) \leq \bar{\gamma}(\tau, m)$ for all τ and m . It follows from the determination of $\underline{\gamma}(\tau, m)$ that as a function of σ_x , $\underline{\gamma}(\tau, m)$ is decreasing in the standard deviation σ_x , or in other words, it is increasing in the precision of the private signal. Moreover as σ_x^2 decreases to zero:

$$\frac{\gamma_a + \rho_a \gamma_{aA}}{\gamma_a + \gamma_{aA}} = \rho_{a\theta}^2,$$

and hence

$$\lim_{\sigma_x \downarrow 0} \underline{\gamma}(\tau, m) = \frac{\sigma_a \rho_{a\theta}}{\sigma_\theta},$$

which concludes the proof. ■

Proof of Proposition 15. For any given ρ_a , with $0 \leq \rho_a < 1$, the ratio

$$\frac{\gamma_a + \rho_a \gamma_{aA}}{\gamma_a + \gamma_{aA}}$$

is larger than 1 if and only $\gamma_{aA} > 0$. Thus we can identify the sign of γ_{aA} if we can establish that the ratio on the rhs of (85), which determined the lower bound on the equilibrium slope, is larger than 1. Now, clearly if $\rho_a < 1$, and if for some $\bar{\sigma}_x$:

$$\bar{\sigma}_x \rho_{a\theta} \sqrt{1 - \rho_a} - (1 - \rho_{a\theta}^2) \sigma_\theta < 0,$$

we have

$$\frac{(\rho_a - \rho_{a\theta}^2) \sigma_\theta + \rho_a (\bar{\sigma}_x \rho_{a\theta} \sqrt{1 - \rho_a} - (1 - \rho_{a\theta}^2) \sigma_\theta)}{(\rho_a - \rho_{a\theta}^2) \sigma_\theta + (\bar{\sigma}_x \rho_{a\theta} \sqrt{1 - \rho_a} - (1 - \rho_{a\theta}^2) \sigma_\theta)} > 1,$$

then the above ratio will remain above 1 for all $\sigma_x < \bar{\sigma}_x$.

Similarly, if for given data the expression on the rhs of (84) is smaller than 1 for some $\bar{\sigma}_y$:

$$\rho_{a\theta}^2 + \frac{\bar{\sigma}_y}{\sigma_\theta} \rho_{a\theta} \sqrt{\rho_a - \rho_{a\theta}^2} < 1,$$

then it will remain smaller than 1 for all $0 \leq \sigma_y < \bar{\sigma}_y$. ■

Proof of Proposition 17. From the observation of the demand, we get γ_{aA} point identified, but then still have to identify γ_a and $\gamma_{a\theta}$. We observe the covariance $\rho_{A\theta} \sigma_A \sigma_\theta$ and the aggregate variance σ_A^2 . This allows us to identify the correlation coefficient $\rho_{A\theta}$, and using Proposition 7:

$$\mu_A = -\frac{\lambda_a + \mu_\theta \gamma_{a\theta}}{\gamma_a + \gamma_{aA}},$$

and

$$\sigma_A = -\frac{\sigma_\theta \gamma_{a\theta} \rho_{A\theta} \sqrt{\rho_a}}{\rho_a \gamma_{aA} + \gamma_a}.$$

We thus have two equations to identify the four unknown structural parameters $(\lambda_a, \gamma_a, \gamma_{a\theta}, \rho_a)$. We can solve for $(\lambda_a, \gamma_{a\theta})$ in terms of the remaining unknowns γ_a and ρ_a to obtain:

$$\lambda_a = \frac{\sigma_A \gamma_a \mu_\theta + \sigma_A \gamma_{aA} \rho_a \mu_\theta - \gamma_{aA} \mu_A \sqrt{\rho_a} \sigma_\theta \rho_{A\theta} - \mu_A \gamma_a \sqrt{\rho_a} \sigma_\theta \rho_{A\theta}}{\sqrt{\rho_a} \sigma_\theta \rho_\theta},$$

and

$$\gamma_{a\theta} = -\frac{\sigma_A \gamma_a + \sigma_A \gamma_{aA} \rho_a}{\sqrt{\rho_a} \sigma_\theta \rho_{A\theta}}.$$

In particular, we would like to know whether this allows us to identify the ratio:

$$\frac{\gamma_{a\theta}}{\gamma_a + \gamma_{aA}} = -\frac{\sigma_A}{\sigma_\theta \rho_{\theta A}} \frac{1}{\sqrt{\rho_a}} \frac{\gamma_a + \gamma_{aA} \rho_a}{\gamma_a + \gamma_{aA}}, \quad (86)$$

and even the value of $\gamma_{a\theta}$:

$$\gamma_{a\theta} = -\frac{\sigma_A}{\sigma_\theta \rho_{\theta A}} \frac{\gamma_a + \gamma_{aA} \rho_a}{\sqrt{\rho_a}}. \quad (87)$$

We see from (86) and (87) that this is impossible (without bounds on the information.) ■

Proof of Proposition 16. The parameter of the demand function, γ_{aA} , is identified as we observe the aggregate quantity and the aggregate price. The covariance between P and A then identifies γ_{aA} . By Proposition 2, if we know the information structure (τ_x, τ_y) , we can infer from the equilibrium variance σ_A^2 and covariance $\rho_{A\theta} \sigma_A \sigma_\theta$ of the aggregate action A , the value of the linear coefficients α_x^* and α_y^* . Given the equilibrium relationship between α_x^* and α_y^* as documented by (18), this allows to identify the value of γ_a given the known value of γ_{aA} . But, now the composition of the equilibrium coefficient α_x^* and α_y^* allow to point identify $\gamma_{a\theta}$ as well. ■

Proof of Proposition 18. It remains to set identify the value of γ_a and $\gamma_{a\theta}$. From the aggregate data we can infer $\rho_{A\theta}$, but not ρ_a or $\rho_{a\theta}$. We know from Proposition 7 that the aggregate covariance is related to ρ_a and $\rho_{a\theta}$ through

$$\rho_{A\theta} = \frac{\rho_{a\theta}}{\sqrt{\rho_a}}. \quad (88)$$

Thus given $\rho_{A\theta}$, there is a curve, given by (88) in the unit square of $(\rho_a, \rho_{a\theta})$ which is compatible with the given data m , and in particular $\rho_{A\theta}$. Now consider a prior information structure (τ_x, τ_y) . We then ask which pairs $(\rho_a, \rho_{a\theta})$ are compatible with the given information structure. To this end, we need to know whether for a given pair $(\rho_a, \rho_{a\theta})$ there exists γ_a for the identified value of γ_{aA} such that the pair $(\rho_a, \rho_{a\theta})$ is in the intersection of the equilibrium set $C_x(\tau, r_A) \cap C_y(\tau, r_A)$.

By Proposition 11, the public equilibrium set $C_y(\tau, r_A)$ increases in γ_a and the private equilibrium set $C_x(\tau, r_A)$ decreases with γ_a . It follows that for every $(\rho_a, \rho_{a\theta})$ there is a lower bound and an upper bound, $\underline{\gamma}_a(\tau, m)$ and $\bar{\gamma}_a(\tau, m)$, respectively, for which if γ_a satisfies $\underline{\gamma}_a(\tau, m) \leq \gamma_a \leq \bar{\gamma}_a(\tau, m)$, then $C_x(\tau, r_A) \cap C_y(\tau, r_A) \cap (\rho_a, \rho_{a\theta}) \neq \emptyset$. From Proposition 10, we know that an increase in the precision of the prior information (τ_x, τ_y) lead $C_x(\tau, r_A)$ and $C_y(\tau, r_A)$ to shrink. Eventually as (τ_x, τ_y) decrease sufficiently, $\underline{\gamma}_a(\tau, m) = \bar{\gamma}_a(\tau, m)$. At this point, we have point-identified γ_a at a singleton $(\rho_a, \rho_{a\theta})$, and hence it follows from (87) that we also have point identified $\gamma_{a\theta}$. Since

the interval is $\left[\underline{\gamma}_a(\tau, m), \bar{\gamma}_a(\tau, m)\right]$ is shrinking with a decrease in the variance, it follows from the monotonicity of the rhs of (86) that the identified set is also shrinking with a increase in the precision of the public and private information. ■

References

- ANGELETOS, G., AND A. PAVAN (2007): “Efficient Use of Information and Social Value of Information,” *Econometrica*, 75, 1103–1142.
- ARADILLAS-LOPEZ, A., AND E. TAMER (2008): “The Identification Power of Equilibrium in Simple Games,” *Journal of Business and Economic Statistics*, 26, 261–283.
- AUMANN, R. (1987): “Correlated Equilibrium as an Expression of Bayesian Rationality,” *Econometrica*, 55, 1–18.
- BAJARI, P., H. HONG, J. KRAINER, AND D. NEKIPELOV (2010): “Estimating Static Models of Strategic Interaction,” *Journal of Business and Economic Statistics*, 28, 469–482.
- BERGEMANN, D., AND S. MORRIS (2005): “Robust Mechanism Design,” *Econometrica*, 73, 1771–1813.
- (2011): “Correlated Equilibrium in Games of Incomplete Information,” Discussion paper, Cowles Foundation for Research in Economics, Yale University and Princeton University.
- BERGEMANN, D., S. MORRIS, AND S. TAKAHASHI (2010): “Interdependent Preferences and Strategic Distinguishability,” Discussion Paper 1772, Cowles Foundation for Research in Economics, Yale University.
- BERGEMANN, D., AND M. PESENDORFER (2007): “Information Structures in Optimal Auctions,” *Journal of Economic Theory*, 137, 580–609.
- BERGEMANN, D., AND J. VÄLIMÄKI (2002): “Information Acquisition and Efficient Mechanism Design,” *Econometrica*, 70, 1007–1033.
- BLACKWELL, D. (1953): “Equivalent Comparison of Experiments,” *Annals of Mathematics and Statistics*, 24, 265–272.
- BRILLINGER, D. (1969): “The Calculation of Cumulants Via Conditioning,” *Annals of the Institute of Statistical Mathematics*, 21, 215–218.
- CHWE, M. (2006): “Incentive Compatibility Implies Signed Covariance,” Discussion paper, UCLA.
- DASGUPTA, P., AND E. MASKIN (2000): “Efficient Auctions,” *Quarterly Journal of Economics*, 115, 341–388.
- DEGROOT, M. (1970): *Optimal Statistical Decisions*. McGraw Hill, New York.
- FORGES, F. (1993): “Five Legitimate Definitions of Correlated Equilibrium in Games with Incomplete Information,” *Theory and Decision*, 35, 277–310.
- (2006): “Correlated Equilibrium in Games with Incomplete Information Revisited,” *Theory and Decision*, 61, 329–344.
- GOSSNER, O. (2000): “Comparison of Information Structures,” *Games and Economic Behavior*, 30, 44–63.

- GUESNERIE, R. (1992): “An Exploration of the Eductive Justifications of the Rational-Expectations Hypothesis,” *American Economic Review*, 82, 1254–1278.
- JEHIEL, P., AND B. MOLDOVANU (2001): “Efficient Design with Interdependent Valuations,” *Econometrica*, 69, 1237–1259.
- JEHIEL, P., B. MOLDOVANU, M. MEYER-TER-VEHN, AND B. ZAME (2006): “The Limits of Ex Post Implementation,” *Econometrica*, 74, 585–610.
- KAJII, A., AND S. MORRIS (1997): “The Robustness of Equilibria to Incomplete Information,” *Econometrica*, 65, 1283–1309.
- KAMENICA, E., AND M. GENTZKOW (2010): “Bayesian Persuasion,” *American Economic Review*, p. forthcoming.
- LEHRER, E., D. ROSENBERG, AND E. SHMAYA (2006): “Signaling and Mediation in Bayesian Games,” Discussion paper, Tel Aviv University and University of Paris.
- (2010): “Signaling and Mediation in Games with Common Interest,” *Games and Economic Behavior*, 68, 670–682.
- MORRIS, S., AND H. SHIN (2002): “Social Value of Public Information,” *American Economic Review*, 92, 1521–1534.
- NEYMAN, A. (1997): “Correlated Equilibrium and Potential Games,” *International Journal of Game Theory*, 26, 223–227.
- PAULA, A. D., AND X. TANG (2010): “Inference of Signs of Interaction Effects in Simultaneous Games with Incomplete Information,” Discussion Paper PIER Working Paper 10-010, University of Pennsylvania.
- SAINT-PAUL, G. (2011): “Ideological Bias and the Production of Macroeconomic Theories,” Discussion paper, University of Toulouse.
- SWEETING, A. (2009): “The Strategic Timing of Radio Commercials: An Empirical Analysis Using Multiple Equilibria,” *RAND Journal of Economics*, 40, 710–742.