

**ROBUST PREDICTIONS
IN GAMES WITH INCOMPLETE INFORMATION**

By

Dirk Bergemann and Stephen Morris

**September 2011
Revised December 2011**

COWLES FOUNDATION DISCUSSION PAPER NO. 1821R



**COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
YALE UNIVERSITY
Box 208281
New Haven, Connecticut 06520-8281**

<http://cowles.econ.yale.edu/>

Robust Predictions

in

Games with Incomplete Information*

Dirk Bergemann[†]

Stephen Morris[‡]

December 10, 2011

Abstract

We analyze games of incomplete information and offer equilibrium predictions which are valid for, and in this sense robust to, all possible private information structures that the agents may have. We completely characterize the set of Bayes correlated equilibria in a class of games with quadratic payoffs and normally distributed uncertainty in terms of restrictions on the first and second moments of the equilibrium action-state distribution. We derive exact bounds on how prior knowledge about the private information refines the set of equilibrium predictions.

We consider information sharing among firms under demand uncertainty and find newly optimal information policies via the Bayes correlated equilibria. Finally, we reverse the perspective and investigate the identification problem under concerns for robustness to private information. The presence of private information leads to set rather than point 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, Matthew Gentzkow, Phil Haile, Emir Kamenica, Marc Henry, Arthur Lewbel, Larry Samuelson, and Elie Tamer, and research assistance from Brian Baisa and Aron Tobias. 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 at the 2011 North American Winter Meetings of the Econometric Society.

[†]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 a posterior belief about the payoff states, and a posterior belief about the beliefs of the other agents. The posterior belief about the payoff state represents knowledge about the *payoff environment*, whereas the posterior belief about the beliefs of the other agents represents knowledge about the *belief environment*. In turn, the private information of the agent, the type in the language of Bayesian games, influences the optimal strategy of the agent, and ultimately the equilibrium distribution over actions and states. 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* the common prior over the payoff states. The fundamental uncertainty about the set of feasible payoffs is thus completely described by the common prior over the payoff states, which we also refer to as the fundamental states. 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 states. The belief environment then describes a potentially rich type space which is subject only to the constraint that the marginal distribution over the fundamental states coincides with the common prior. A pair of payoff environment and belief environment form a standard Bayesian game. Yet importantly, for a given payoff environment, there are many belief environments, and each distinct belief environment may lead to a 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 the Bayes Nash equilibria of all possible belief environments 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 distributions for all possible belief environments and the class of Bayes correlated equilibrium distributions. 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 set of 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.

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 state and the equilibrium distribution itself are normally distributed then the characterization of the equilibrium is completely given by the first and second moments. 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 it is not a complete characterization in the sense that the determination of the higher moments is incomplete.

The relationship between the Bayes Nash equilibrium and the Bayes correlated equilibrium is shown to lead to new insights into the relationship between information structure and the nature of the Bayes Nash equilibrium. The compact representation of the Bayes correlated equilibria allows us to assess the private and social welfare across the entire set of possible information structures and associated Bayes Nash equilibria. We illustrate this in the context of information sharing among firms. A striking result by Clarke (1983) was the finding that firms, when facing uncertainty about a common parameter of demand, will never find it optimal to share information. The present analysis of the Bayes correlated equilibrium allows us to modify this insight - implicitly by allowing for richer information structures than previously considered - and we find that the Bayes correlated equilibrium that maximizes the private welfare of the firms is not necessarily obtained with either zero or full information disclosure.

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 no additional information beyond the common prior about the payoff state. 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 state. 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. Finally, we therefore reverse the perspective of our analysis and consider the issue of identification rather than prediction. Namely, we are asking whether the observable data, actions and payoff states, can identify the structural parameters of the payoff functions, and thus of the game, without overly narrow 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 which restrictions are common to all possible belief environments given a specific payoff environment. We find that we can robustly identify the sign of some interaction parameters, but have to leave the sign and size of other parameters, in particular whether the agents are playing a game of strategic substitutes or complements, unidentified. 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 Bayes 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), Bergemann and Morris (2005), and Perry and Reny (2002) 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

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.

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. 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 an analyst to infer the nature of the payoff interaction between the agent's choice and the random variable. 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 (2011) derive the sender-optimal information structure.

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 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 state. Section 5 begins with the analysis of the Bayes correlated equilibrium. We give a complete description of the equilibrium set in terms of moment restrictions on the joint equilibrium distribution. In Section 6 we analyze how prior information about the belief environment can further restrict the equilibrium predictions. In Section 7 we consider the optimal sharing of information among firms. By rephrasing the choice of information policy as a choice over information structures, we derive newly optimal information policies through the lens of Bayes correlated equilibria. In Section 8, we turn from prediction to the issue of identification. Section 9 discusses some possible extensions and offers concluding remarks. The Appendix collects the proofs from the main body of the text.

2 Set-Up

We first define the solution concept of Bayes correlated equilibrium. We then relate the notion of Bayes correlated equilibrium to robust equilibrium predictions in a class of continuum player games with symmetric payoff. In the companion paper, Bergemann and Morris (2011a), we develop this solution concept and its relationship to robust predictions in canonical finite player and finite action games. In the companion paper, we also show how the results there can be adapted and refined first to symmetric payoffs and then to the continuum of agents and continuum of actions analyzed here.

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 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 state $\theta \in \Theta$ whereas it does not specify the private information the agents may have access to.

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

$$\mathbb{E}_{\hat{\mu}(\cdot|a)} u(a, h, \theta) \geq \mathbb{E}_{\hat{\mu}(\cdot|a)} u(a', h, \theta) \quad (1)$$

for each $a \in \mathbb{R}$ and $a' \in \mathbb{R}$; and

$$\text{marg}_{\Theta} \mu = \psi. \quad (2)$$

This 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 will show that Bayes correlated equilibrium captures all behavior that could arise if players observed additional private information (in a symmetric way) and played according to a symmetric Bayes Nash equilibrium. To formalize this, we first introduce the relevant (symmetric) information structures for this continuum agent economy. 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 structure”, is parameterized by (T, π) .

The payoff environment (u, ψ) and the belief environment (T, π) together define a game of incomplete information $((u, \psi), (T, \pi))$. A symmetric 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 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 $\mu = \psi \circ \pi \circ \sigma$.

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 structure (T, π) .

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

In our companion paper, Bergemann and Morris (2011a), the notion of Bayes correlated equilibrium is defined in a canonical game theory environment - with a finite number of actions, agents and states - where the players have additional information from an information structure (T, π) , and thus a Bayes correlated equilibrium is a joint distribution over action, states and types, i.e. a distribution $\nu \in \Delta(A \times \Theta \times T)$. In the language of the more general notion offered there, the Bayes correlated equilibrium defined here is the Bayes correlated information with the “null information structure”, i.e. the case in which the agents are not assumed a priori to have access to a specific information structure. Here, we choose this minimal notion of a Bayes correlated equilibrium to obtain robust predictions for an observer who only knows the payoff environment but has “null” information about the belief environment of the game. But, just as in the companion paper, Bergemann and Morris (2011a), we can analyze the impact of private information on the size of the Bayes correlated equilibrium set. In fact in Section 6, we analyze how prior knowledge of the belief environment can refine the set of equilibrium predictions. We maintain our restriction to normally distributed uncertainty, now normally distributed types, to obtain explicit descriptions of the resulting restriction on the equilibrium set. By contrast, in Bergemann and Morris (2011a), we allow for general information structures and derive a many player generalization of the ordering of Blackwell (1953) as a necessary and sufficient condition to order the set of Bayes correlated equilibrium. However, within this general environment, we do not obtain an explicit and compact description of the equilibrium set in terms of the first and second moments of the equilibrium distributions, as we do in the present analysis.

The general notion of Bayes correlated information also facilitates the discussion of the relationships between the notion of Bayes correlated equilibrium, and related, but distinct notions of correlated equilibrium in games of incomplete information, most notably in the work of Forges (1993), which is titled and identifies “five legitimate definitions of correlated equilibrium in games with incomplete information”. We refer to the reader to the companion paper, Bergemann and Morris (2011a) for a detailed discussion and comparison.

3 Linear Best Response and Normal Uncertainty

For the remainder of this paper, we shall consider a payoff environment with linear best responses and normally distributed uncertainty. Thus we assume that player i sets his action equal to a linear function of his expectations of the average action of others A and a payoff relevant state θ . Thus we have

$$a_i = r\mathbb{E}_i(A) + s\mathbb{E}_i(\theta) + u, \tag{3}$$

where $r, s, u \in \mathbb{R}$ are the parameters of the best response function and are assumed to be identical across players. The average action of the all players but i is represented by A . In the case of a finite number I of

players, it is therefore the sum:

$$A \triangleq \frac{1}{I-1} \sum_{j \neq i} a_j, \quad (4)$$

and in the case of continuum of agents, $j \in [0, 1]$, it is an integral:

$$A \triangleq \int_0^1 a_j dj. \quad (5)$$

With the linear best response, the equilibrium behavior with a finite, but large number of players converges to the equilibrium behavior with a continuum of players. The model with a continuum of players has the advantage that we do not need to keep track of the relative weight of the individual player i , namely $1/I$, and the weight of all the other players, namely $(I-1)/I$. In consequence, the expression of the equilibrium strategies are frequently more compact with a continuum of players. In the subsequent analysis, we will focus on the game with a continuum of players, but report on the necessary adjustments with a finite player environment.

The parameter r represents the strategic interaction among the players, and we therefore refer to it as the “interaction parameter”. If $r < 0$, then we have a game of strategic substitutes, if $r > 0$, then we have a game of strategic complementarities. The case of $r = 0$ represents the case of single person decision problem where each player i simply responds the state of the world θ , but is not concerned about his interaction with the other players.

The parameter s represents the informational response of player i , and it can be either negative or positive. We shall assume that the state of the world θ matters for the decision of agent i , and hence $s \neq 0$.

We shall assume that the interaction parameter r is bounded above, or

$$r \in (-\infty, 1), \quad (6)$$

which is a necessary and sufficient condition for the complete information game to have an interior Nash equilibrium. In fact, with the restriction (6), the Nash equilibrium in the game with complete information is unique and given by:

$$a_i(\theta) = \frac{u}{1-r} + \frac{s}{1-r}\theta, \text{ for all } i \text{ and } \theta. \quad (7)$$

Moreover, under complete information about the state of the world θ , even the correlated equilibrium is unique; Neyman (1997) gives an elegant argument.

The payoff state, or the state of the world, θ is assumed to be distributed normally with

$$\theta \sim N(\mu_\theta, \sigma_\theta^2). \quad (8)$$

The present environment of linear best response and normally distributed uncertainty encompasses a wide class of interesting economic environments. The following three 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 A taken in the population. Morris and Shin (2002) 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 notation, the beauty contest model set $s = 1 - r$ and $u = 0$ with $0 \leq r < 1$.

Example 2 (Competitive and Strategic Markets) Guesnerie (1992) presents 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. If there is uncertainty about the demand intercept, we can write the demand curve as $p(A) = s\theta + rA + u$ with $r < 0$ while the cost of firm i is $c(a_i) = \frac{1}{2}a_i^2$. Individual firm profits are now given by

$$a_i p(A) - c(a_i) = (rA + s\theta + u)a_i - \frac{1}{2}a_i^2.$$

In an alternative interpretation, we can have a common cost shock, so the demand curve is $p(A) = rA + u$ with $r < 0$ while the cost of firm i is $c(a_i) = -s\theta a_i + \frac{1}{2}a_i^2$. Such an economy can be derived as the limit of large, but finite, Cournot markets, as shown by Vives (1988), (2011).

Example 3 (Quadratic Economies and the Social Value of Information) Angeletos and Pavan (2009) consider a general class of quadratic economies (games) with a continuum of agents and private information about a common state $\theta \in \mathbb{R}$. There the payoff of agent i is given by a symmetric quadratic utility function $u_i(a_i, A, \theta)$, which depends on the individual action a_i , the average action A and the payoff state $\theta \in \mathbb{R}$:

$$u_i(a_i, A, \theta) \triangleq \frac{1}{2} \begin{pmatrix} a_i \\ A \\ \theta \end{pmatrix}' \begin{pmatrix} U_{aa} & U_{aA} & U_{a\theta} \\ U_{aA} & U_{AA} & U_{A\theta} \\ U_{a\theta} & U_{A\theta} & U_{\theta\theta} \end{pmatrix} \begin{pmatrix} a_i \\ A \\ \theta \end{pmatrix}, \quad (9)$$

where the matrix $U = \{U_{kl}\}$ represents the payoff structure of the game. In the earlier working paper version, Bergemann and Morris (2011b), we also represented the payoff structure of the game by the matrix U . Angeletos and Pavan (2009) assume that the payoffs are concave in the own action: $U_{aa} < 0$, and that the interaction of the individual action and the average action (the “indirect effect”) is bounded by the own action (the “direct effect”):

$$-U_{aA}/U_{aa} < 1 \Leftrightarrow U_{aa} + U_{aA} < 0. \quad (10)$$

The best response in the quadratic economy (with complete information) is given by:

$$a_i = -\frac{U_{aA}A + U_{a\theta}\theta}{U_{aa}}.$$

The quadratic term of the own cost, U_{aa} simply normalizes the terms of the strategic and informational externality, U_{aA} and $U_{a\theta}$. In terms of the present notation we have

$$r = -\frac{U_{aA}}{U_{aa}}, s = -\frac{U_{a\theta}}{U_{aa}}.$$

Their restriction (10) is equivalent to the present restriction (6). The entries in the payoff matrix U which do not refer to the individual action a , i.e. the entries in the lower submatrix of U , namely

$$\begin{bmatrix} U_{AA} & U_{A\theta} \\ U_{A\theta} & U_{\theta\theta} \end{bmatrix}$$

are not relevant for the determination of either the Bayes Nash or the Bayes correlated equilibrium. These entries may be relevant for welfare analysis (as in Angeletos and Pavan (2009)), but for the welfare analysis in this paper they are not and can be uniformly set to zero.

4 Bayes Nash Equilibrium

We first report as a benchmark a standard approach to analyzing this class of games of incomplete information. Starting with the payoff environment described in the previous section, we add a description of the belief environment, i.e., what players know about the state and others' beliefs. Specifically, we assume that each player observes a two-dimensional signal. 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 given variances, and hence for a given type space, we then analyze the Bayes Nash equilibrium/a of the basic game. 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:

$$x_i = \theta + \varepsilon_i, \tag{11}$$

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

$$y = \theta + \varepsilon. \quad (12)$$

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:

$$\tau_x \triangleq \sigma_x^{-2}, \tau_y \triangleq \sigma_y^{-2}, \tau_\theta \triangleq \sigma_\theta^{-2} \quad \text{and} \quad \tau \triangleq \sigma_\theta^{-2} + \sigma_x^{-2} + \sigma_y^{-2};$$

we refer to the vector (τ_x, τ_y) 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. The best response:

$$a_i = rA + s\theta + u,$$

reflects the, possibly conflicting, objectives that agent i faces. Each agent has to solve a prediction-like problem in which he wishes to match his action, with the state θ and the average action A simultaneously. The interaction parameters, s and r , determine the weight that each component, θ and A , receives in the deliberation of the agent. If there is zero strategic interaction, or $r = 0$, then each agent faces a pure prediction problem. Now, we observed earlier, see (7), that the resulting Nash equilibrium strategy is given by:

$$a^*(\theta) \triangleq \frac{u}{1-r} + \frac{s}{1-r}\theta. \quad (13)$$

We refer to the terms in equilibrium strategy (13), $u/(1-r)$ and $s/(1-r)$, as the *equilibrium intercept* and the *equilibrium slope*, respectively.

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. In the game with incomplete information, agent i receives a pair of signals, x_i and y , generated by the information structure (11) and (12). 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' signals either, he can only form an expectation about their actions, 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 :

$$a_i = r\mathbb{E}[A|x_i, y] + s\mathbb{E}[\theta|x_i, y] + u.$$

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 : $a(x_i, y) = \alpha_0 + \alpha_x x_i + \alpha_y y$. The equilibrium is then identified by the linear coefficients $\alpha_0, \alpha_x, \alpha_y$, which we expect to depend on the interaction terms (r, s, u) and the information structure (τ_x, τ_y) .

Proposition 2 (Linear Bayes Nash Equilibrium)

The unique Bayes Nash equilibrium is a linear equilibrium: $a(x, y) = \alpha_0^ + \alpha_x^* x + \alpha_y^* y$, with the coefficients given by:*

$$\alpha_0^* = \frac{u}{1-r} + \frac{s}{1-r} \frac{\mu_\theta \tau_\theta}{\tau - r\tau_x}, \quad \alpha_x^* = s \frac{\tau_x}{\tau - r\tau_x}, \quad \alpha_y^* = \frac{s}{1-r} \frac{\tau_y}{\tau - r\tau_x}. \quad (14)$$

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). With the normalization of the average action given by (4) and (5), the above equilibrium strategy is independent of the number of players, and in particular independent of the finite or continuum version of the environment.

The Bayes Nash equilibrium shares the uniqueness property with the Nash equilibrium, its complete information counterpart. We observe that the linear coefficients α_x^* and α_y^* display the following relationship:

$$\frac{\alpha_y^*}{\alpha_x^*} = \frac{\tau_y}{\tau_x} \frac{1}{1-r}. \quad (15)$$

Thus, if there is no strategic interaction, or $r = 0$, 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, $r \neq 0$, then the relative weights also reflect the informativeness of the signal with respect to the average action. Thus if the game displays strategic complements, $r > 0$, 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 substitutability, $r < 0$, 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, (13) and (14), 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.

Proposition 3 (Attenuation)

The mean action in equilibrium is:

$$\mathbb{E}[a] = \alpha_0^* + \alpha_x^* \mu_\theta + \alpha_y^* \mu_\theta = \frac{u + s\mu_\theta}{1 - r},$$

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

$$|\alpha_x^* + \alpha_y^*| = \left| \frac{s}{1 - r} \right| \left(1 - \frac{\tau_\theta}{\tau - r\tau_x} \right) \leq \left| \frac{s}{1 - r} \right|.$$

Thus, the mean of the individual action, $\mathbb{E}[a]$, is independent of the information structure (τ_x, τ_y) . 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. 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 varies 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:

$$\Sigma_{a_i, a_j, \theta} = \begin{bmatrix} \sigma_a^2 & \rho_a \sigma_a^2 & \rho_{a\theta} \sigma_a \sigma_\theta \\ \rho_a \sigma_a^2 & \sigma_a^2 & \rho_{a\theta} \sigma_a \sigma_\theta \\ \rho_{a\theta} \sigma_a \sigma_\theta & \rho_{a\theta} \sigma_a \sigma_\theta & \sigma_\theta^2 \end{bmatrix}. \tag{16}$$

We denote the correlation coefficient between action a_i and a_j shorthand by ρ_a rather than ρ_{aa} .

With a continuum of agents, we can describe the equilibrium distribution, after replacing the individual action a_j by the average action A , through the triple (a_i, A, θ) . 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 $\sigma_A^2 = \rho_a \sigma_a^2$.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 or $\rho_{a\theta} \sigma_a \sigma_\theta = \rho_{A\theta} \sigma_A \sigma_\theta$.

With the characterization of the unique Bayes Nash equilibrium in Proposition 2, we can express the variance-covariance matrix of the equilibrium joint distribution over (a_i, A, θ) in terms of the equilibrium coefficients (α_x, α_y) and the variances of the underlying random variables $(\theta, \varepsilon_i, \varepsilon)$:

$$\Sigma_{a_i, A, \theta} = \begin{bmatrix} \alpha_x^2 \sigma_x^2 + \alpha_y^2 \sigma_y^2 + \sigma_\theta^2 (\alpha_x + \alpha_y)^2 & \alpha_y^2 \sigma_y^2 + \sigma_\theta^2 (\alpha_x + \alpha_y)^2 & \sigma_\theta^2 (\alpha_x + \alpha_y) \\ \alpha_y^2 \sigma_y^2 + \sigma_\theta^2 (\alpha_x + \alpha_y)^2 & \alpha_y^2 \sigma_y^2 + \sigma_\theta^2 (\alpha_x + \alpha_y)^2 & \sigma_\theta^2 (\alpha_x + \alpha_y) \\ \sigma_\theta^2 (\alpha_x + \alpha_y) & \sigma_\theta^2 (\alpha_x + \alpha_y) & \sigma_\theta^2 \end{bmatrix}. \quad (17)$$

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

$$\alpha_x^* = \frac{\sigma_a}{\sigma_\theta} \rho_{a\theta} - \alpha_y^*, \quad \alpha_y^* = \pm \frac{\sigma_a}{\sigma_y} \sqrt{\rho_a - \rho_{a\theta}^2}. \quad (18)$$

Thus, 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 where $\rho_{a\theta}$, 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:

$$0 \leq \rho_a \leq 1, \text{ and } \rho_a - \rho_{a\theta}^2 \geq 0, \quad (19)$$

²With a finite number of agents and the definition of the average action given by: $A = (1/(I-1)) \sum_{j \neq i} a_j$, the variance of A is given by $\sigma_A^2 = \left(\frac{1}{I-1} + \frac{I-2}{I-1} \rho_a \right) \sigma_a^2$ and hence the variance-covariance matrix in the continuum version is only an approximation, but not exact. We present the exact restrictions in Corollary 1 in the next section.

we can find information structures (τ_x, τ_y) such the coefficients resulting from (18) are indeed the equilibrium coefficients of the associated Bayes Nash equilibrium strategy.

Proposition 4 (Information and Correlation)

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_x, τ_y) such that the associated Bayes Nash equilibrium displays the correlation coefficients $(\rho_a, \rho_{a\theta})$:

$$\tau_x = \frac{(1 - \rho_a) \rho_{a\theta}^2}{\left((1 - \rho_a) + (1 - r) (\rho_a - \rho_{a\theta}^2) \right)^2 \sigma_\theta^2},$$

and

$$\tau_y = \frac{(\rho_a - \rho_{a\theta}^2) \rho_{a\theta}^2 (1 - r)^2}{\left((1 - \rho_a) + (1 - r) (\rho_a - \rho_{a\theta}^2) \right)^2 \sigma_\theta^2}.$$

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_x, τ_y) . 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 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




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

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 Bayes Nash 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 restrict attention to symmetric and normally distributed correlated equilibria and discuss the extent to which these are without loss of generality at the end of this Section. We begin the analysis with a continuum of agents and subsequently describe how the equilibrium restrictions are modified in a finite player environment.

We can characterize the Bayes correlated equilibria in two distinct, yet related, ways. With a continuum of agents, we can characterize the equilibria in terms of the realized average action A and the deviation of the individual action a_i from the average action, $a_i - A$. Under the continuum hypothesis, the distribution around the realized average action A represents the exact distribution of actions by the agents, conditional on the realized average action A . Alternatively we can characterize the equilibria in terms of an arbitrary pair of individual actions, a_i and a_j , and the state of the world θ . The first approach puts more emphasis on the distributional properties of the correlated equilibrium, and is convenient when we go beyond symmetric and normally distributed equilibria, whereas the second approach is closer to the description of the Bayes Nash equilibrium in terms of the individual action.




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

5.1 Equilibrium Moment Restrictions

We consider the class of symmetric and normally distributed Bayes correlated equilibria. With the hypothesis of a normally distributed Bayes correlated equilibrium, the aggregate distribution of the state of the world θ and the average action A is described by:

$$\begin{pmatrix} \theta \\ A \end{pmatrix} \sim N \left(\begin{pmatrix} \mu_\theta \\ \mu_A \end{pmatrix}, \begin{pmatrix} \sigma_\theta^2 & \rho_{A\theta}\sigma_A\sigma_\theta \\ \rho_{A\theta}\sigma_A\sigma_\theta & \sigma_A^2 \end{pmatrix} \right).$$

In the continuum economy, we can describe the individual action a as centered around the average action A with some dispersion σ_η^2 , so that $a = A + \eta$, for some $\eta \sim N(0, \sigma_\eta^2)$. In consequence, the joint equilibrium distribution of (θ, A, a) is given by:

$$\begin{pmatrix} \theta \\ A \\ a \end{pmatrix} \sim N \left(\begin{pmatrix} \mu_\theta \\ \mu_A \\ \mu_a \end{pmatrix}, \begin{pmatrix} \sigma_\theta^2 & \rho_{A\theta}\sigma_A\sigma_\theta & \rho_{A\theta}\sigma_A\sigma_\theta \\ \rho_{A\theta}\sigma_A\sigma_\theta & \sigma_A^2 & \sigma_A^2 \\ \rho_{A\theta}\sigma_A\sigma_\theta & \sigma_A^2 & \sigma_A^2 + \sigma_\eta^2 \end{pmatrix} \right). \quad (20)$$

The analysis of the Bayes correlated equilibrium proceeds by deriving restrictions on the joint equilibrium distribution (20). In other words, we seek to identify the restrictions on the moments of the equilibrium distribution. Given that we presently restrict attention to a multivariate normal distribution, it is sufficient to derive restrictions in terms of the first and second moments of the equilibrium distribution (20). The equilibrium restrictions arise from two sources: (i) the best response conditions of the individual agents:

$$a_i = r\mathbb{E}[A|a_i] + s\mathbb{E}[\theta|a_i] + u, \quad \text{for all } i \text{ and } a_i \in \mathbb{R}, \quad (21)$$

and (ii) the consistency condition, see Definition 1, where the latter condition, namely that the marginal distribution over θ is equal to the common prior over θ , is satisfied by construction of the joint equilibrium

distribution (20). The best response condition (21) of the Bayes correlated equilibrium allows the agent to form his expectation over the average action A and the state of the world θ by conditioning on the information that is contained in his “recommended” equilibrium action a_i .

As the best response condition (21) uses the expectation of the individual agent, it is convenient to introduce the following change of variable for the equilibrium random variable. By hypothesis of the symmetric equilibrium, we have:

$$\mu_a = \mu_A \text{ and } \sigma_a^2 = \sigma_A^2 + \sigma_\eta^2.$$

The covariance between the individual action and the average action is given by $\rho_{aA}\sigma_a\sigma_A = \sigma_A^2$, and is identical, by construction, to the covariance between the individual actions:

$$\rho_a\sigma_a^2 = \sigma_A^2. \quad (22)$$

We can therefore express the correlation coefficient between individual actions, ρ_a , as:

$$\rho_a = \frac{\sigma_A^2}{\sigma_A^2 + \sigma_\eta^2}, \quad (23)$$

and the correlation coefficient between individual action and the state θ as:

$$\rho_{a\theta} = \rho_{A\theta} \frac{\sigma_A}{\sigma_a}. \quad (24)$$

In consequence, we can rewrite the joint equilibrium distribution of (θ, A, a) in terms of the moments of the state of the world θ and the individual action a as:

$$\begin{pmatrix} \theta \\ A \\ a \end{pmatrix} \sim N \left(\begin{pmatrix} \mu_\theta \\ \mu_a \\ \mu_a \end{pmatrix}, \begin{pmatrix} \sigma_\theta^2 & \rho_{a\theta}\sigma_a\sigma_\theta & \rho_{a\theta}\sigma_a\sigma_\theta \\ \rho_{a\theta}\sigma_a\sigma_\theta & \rho_a\sigma_a^2 & \rho_a\sigma_a^2 \\ \rho_{a\theta}\sigma_a\sigma_\theta & \rho_a\sigma_a^2 & \sigma_a^2 \end{pmatrix} \right). \quad (25)$$

With the joint equilibrium distribution described by (25), we now use the best response property (21), to completely characterize the moments of the equilibrium distribution. Note that this corresponds to imposing the obedience condition (1) in the general setting of Section 2.

As the best response property (21) 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:

$$\mathbb{E}[a_i] = u + s\mathbb{E}[\mathbb{E}[\theta | a_i]] + r\mathbb{E}[\mathbb{E}[A | a_i]]. \quad (26)$$

But by symmetry, it follows that the expected action of each agent is equal to expected average action A , and hence we can use (26) to solve for the mean of the individual action and the average action:

$$\mathbb{E}[a_i] = \mathbb{E}[A] = \frac{u + s\mathbb{E}[\theta]}{1 - r} = \frac{u + s\mu_\theta}{1 - r}. \quad (27)$$

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 parameters (r, s, u) 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 (21). But, now we use the property of the conditional expectation, rather than the iterated expectation to derive restrictions 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 of the average action, $\mathbb{E}[A | a_i]$, have to change with a_i at exactly the rate required to maintain the best response property:

$$1 = \left(s \frac{d\mathbb{E}[\theta | a_i]}{da_i} + r \frac{d\mathbb{E}[A | a_i]}{da_i} \right), \text{ for all } a_i \in \mathbb{R}.$$

Given the multivariate normal distribution (25), the conditional expectations $\mathbb{E}[\theta | a_i]$ and $\mathbb{E}[A | a_i]$ are linear in a_i and given by

$$\mathbb{E}[\theta | a_i] = \left(1 - \frac{\rho_{a\theta}\sigma_\theta}{\sigma_a} \frac{s}{1-r} \right) \mu_\theta + \frac{\rho_{a\theta}\sigma_\theta}{\sigma_a} \left(a_i - \frac{u}{1-r} \right), \quad (28)$$

and

$$\mathbb{E}[A | a_i] = \frac{u + s\mu_\theta}{1-r} (1 - \rho_a) + \rho_a a_i. \quad (29)$$

The optimality of the best response property can then be expressed, using (28) and (29) as

$$1 = s \frac{\rho_{a\theta}\sigma_\theta}{\sigma_a} + r\rho_a.$$

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 (r, s) . In fact, it is convenient to solve for the standard deviation of the individual actions σ_a , or

$$\sigma_a = \frac{\sigma_\theta s \rho_{a\theta}}{1 - \rho_a r}. \quad (30)$$

The remaining restrictions on the correlation coefficients ρ_a and $\rho_{a\theta}$ are coming in the form of inequalities from the change of variables in (22)-(24), where

$$\rho_{a\theta}^2 = \rho_{A\theta}^2 \frac{\sigma_A^2}{\sigma_a^2} = \rho_{A\theta}^2 \rho_a \leq \rho_a. \quad (31)$$

Finally, the standard deviation has to be positive, or $\sigma_a \geq 0$. Now, it follows from the assumption of moderate interaction, $r < 1$, and the nonnegativity restriction of ρ_a implied by (31) that $1 - \rho_a r > 0$, and thus to guarantee that $\sigma_a \geq 0$, it has to be that $s\rho_{a\theta} \geq 0$. Thus the sign of the correlation coefficient $\rho_{a\theta}$ has to equal the sign of the interaction term s . We summarize these results.




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

Proposition 5 (First and Second Moments of BCE)

A multivariate normal distribution of (a_i, A, θ) is a symmetric Bayes correlated equilibrium if and only if

1. the mean of the individual action is:

$$\mathbb{E}[a_i] = \frac{u}{1-r} + \mu_\theta \frac{s}{1-r}; \tag{32}$$

2. the standard deviation of the individual action is:

$$\sigma_a = \frac{s\rho_{a\theta}}{1-\rho_a r} \sigma_\theta; \quad \text{and} \tag{33}$$

3. the correlation coefficients ρ_a and $\rho_{a\theta}$ satisfy the inequalities:

$$\rho_{a\theta}^2 \leq \rho_a \quad \text{and} \quad s \cdot \rho_{a\theta} \geq 0. \tag{34}$$

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 the 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 r and s , respectively. The set of all correlated equilibria is graphically represented in Figure 3.

The restriction on the correlation coefficients, namely $\rho_{a\theta}^2 \leq \rho_a$, emerged directly from the above change of variable, see (22)-(24). Alternatively, but equivalently, we could have disregarded the restrictions implied by the change of variables, and simply insisted that the matrix of second moments of (25) is indeed a

legitimate variance-covariance matrix, i.e. that is a nonnegative definite matrix. A necessary and sufficient condition for the nonnegativity of the matrix is that the determinant of the variance-covariance matrix is nonnegative, or,

$$\sigma_\theta^6 \rho_{a\theta}^4 s^4 (1 - \rho_a) \frac{\rho_a - \rho_{a\theta}^2}{(1 - \rho_a r)^4} \geq 0 \quad \Rightarrow \quad \rho_{a\theta}^2 \leq \rho_a.$$

Later, we extend the analysis from the pure common value environment analyzed here, to an interdependent value environment (in Section 5.5) and to prior private information (in Section 6). In these extensions, it will be convenient to extract the equilibrium restrictions in form of the correlation inequalities, directly from the restriction of the nonnegative definite matrix, rather than trace them through the relevant change of variable. In any case, these two procedures establish the same equilibrium restrictions.

We observe 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 $\rho_a = 1$.

At this point, it is appropriate to describe how the analysis of the Bayes correlated equilibrium would be modified by the presence of a finite number I of agents. We remarked in Section 3 that the best response function of the agent i is constant in the number of players. As the best response is independent of the number of players, it follows that the equilibrium equality restrictions, namely (32) and (33), are unaffected by the number, in particular the finiteness, of the players. The only modification arises with the change of variable, see (22)-(24), which relied on the continuum of agents. By contrast, the inequality restrictions with a finite number of players can be recovered directly from the fact that variance-covariance matrix $\Sigma_{a_1, \dots, a_I, \theta}$ of the equilibrium random variables $(a_1, \dots, a_I, \theta)$ has to be a nonnegative definite matrix.

Corollary 1 (First and Second Moments of BCE with Finitely Many Players)

A multivariate normal distribution of $(a_1, \dots, a_I, \theta)$ is a symmetric Bayes correlated equilibrium if and only if it satisfies (32), (33), and the correlation coefficients ρ_a and $\rho_{a\theta}$ satisfy the inequalities:

$$\rho_a \geq -\frac{1}{I-1}, \quad \rho_a - \rho_{a\theta}^2 \geq -\frac{1 - \rho_{a\theta}^2}{I-1}, \quad s \cdot \rho_{a\theta} \geq 0. \quad (35)$$

It is immediate to verify that the restrictions of the correlation structure in (35) converge towards the one in (34) as $I \rightarrow \infty$. We observe that the restrictions in (35) are more permissive with a smaller number

of agents, and in particular allow for moderate negative correlation across individual actions with a finite number of agents. By contrast, with infinitely many agents, it is a statistical impossibility that all actions are mutually negatively correlated.

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 (27). 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 (21), that the variance of the individual action matches the sum of the variances of the conditional expectations. 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.

5.2 Volatility and Dispersion

Proposition 5 documents that the relationship between the correlation coefficients ρ_a and $\rho_{a\theta}$ depends only on the sign of the information externality s , but not on the strength of the parameters r and s . 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, $r > 0$, then: (i) σ_a is increasing in ρ_a and $|\rho_{a\theta}|$; (ii) the maximal σ_a is obtained at $\rho_a = |\rho_{a\theta}| = 1$.*
2. *If the game displays strategic substitutes, $r < 0$, then: (i) σ_a is decreasing in ρ_a and increasing in $|\rho_{a\theta}|$; (ii) the maximal σ_a is obtained at*

$$\rho_a = |\rho_{a\theta}|^2 = \min \left\{ -\frac{1}{r}, 1 \right\}. \quad (36)$$

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 substitutes. 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 r is large,

namely $|r| > 1$, 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.

So far we have described the Bayes correlated equilibrium in terms of the triple (θ, A, a) . Yet, a distinct but equivalent representation can be given in terms of $(\theta, A, a - A)$: the state θ , the average action A , the idiosyncratic difference, $a - A$. In games with a continuum of agents, we can interpret the conditional distribution of the agents' action a around the mean A as the 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). 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 $\sigma_A^2 = \rho_a \sigma_a^2$.

Proposition 7 (Volatility and Dispersion)

1. *The volatility is increasing in $|\rho_{a\theta}|$, and increasing in ρ_a if and only if $r \geq -1/\rho_a$;*
2. *The dispersion is increasing in $|\rho_{a\theta}|$ and reaches an interior maximum at:*

$$\rho_a = \rho_{a\theta}^2 = \frac{1}{2 - r}.$$

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 for a larger range of strategic interaction parameters, including moderate strategic substitutes.

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 .

Corollary 2 (Matching BCE and BNE)

For every interaction structure (r, s, u) , 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. $0 < (\tau_x, \tau_y) < \infty$, 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 the precision of the public signal satisfies $\tau_y = 0$.

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 equilibria. 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 Interdependent Value Environment

So far, we have restricted our analysis to the common value environment in which the state of the world is the same for every agent. However, the analysis of the Bayes correlated equilibrium set easily extends to

a model with interdependent, but not necessarily common values. We describe a suitable generalization of the common value environment to an interdependent value environment: the payoff type of agent i is now given by $\theta_i = \theta + \nu_i$, where θ is the common value component and ν_i is the private value component. The distribution of the common component θ is given, as before, by $\theta \sim N(\mu_\theta, \sigma_\theta^2)$, and the distribution of the private component ν_i is given by $\nu_i \sim N(0, \sigma_\nu^2)$. It follows that by increasing σ_ν^2 at the expense of σ_θ^2 , we can move from a model of pure common values to a model of pure private values, and in between are in a canonical model of interdependent values.

The analysis of the Bayes correlated equilibrium can proceed as in Section 5.1. The earlier representation of the Bayes correlated equilibrium in terms of the variance-covariance matrix of the individual action a , the aggregate action A and the common value θ simply has to be augmented by distinguishing between the common value component θ and the private value component ν :

$$\Sigma_{a,A,\theta,\nu} = \begin{bmatrix} \sigma_a^2 & \rho_a \sigma_a^2 & \rho_{a\theta} \sigma_a \sigma_\theta & \rho_{a\nu} \sigma_a \sigma_\nu \\ \rho_a \sigma_a^2 & \rho_a \sigma_a^2 & \rho_{a\theta} \sigma_a \sigma_\theta & 0 \\ \rho_{a\theta} \sigma_a \sigma_\theta & \rho_{a\theta} \sigma_a \sigma_\theta & \sigma_\theta^2 & 0 \\ \rho_{a\nu} \sigma_a \sigma_\nu & 0 & 0 & \sigma_\nu^2 \end{bmatrix}.$$

The new correlation coefficient $\rho_{a\nu}$ represents the correlation between the individual action a and the individual value, the private component ν . The set of the Bayes correlated equilibria are affected by the introduction of the private component in a systematic manner. The equilibrium conditions, in terms of the best response, are given by:

$$a = r\mathbb{E}[A|a] + s\mathbb{E}[\theta + \nu|a] + u. \quad (37)$$

As the private component ν has zero mean, it is centered around the common value θ , the private component does not change the mean action in equilibrium. However, the addition of the private value component does affect the variance and covariance of the Bayes correlated equilibria. In fact, the best response condition (37), restricts the variance of the individual action to:

$$\sigma_a = \frac{s(\sigma_\theta \rho_{a\theta} + \sigma_\nu \rho_{a\nu})}{1 - \rho_a r},$$

so that the standard deviation σ_a of the individual action is now composed of the weighted sum of the common and private value sources of payoff uncertainty. Finally, the additional restrictions that arise from the requirement that the matrix $\Sigma_{a,A,\theta,\nu}$ is indeed a variance-covariance matrix, i.e. that it is a positive definite matrix, simply appear integrated in the original conditions:

$$\rho_a - \rho_{a\theta}^2 \geq 0, \quad 1 - \rho_{a\nu}^2 - \rho_a \geq 0. \quad (38)$$

In other words, to the extent that the individual action is correlated with the private component, it imposes a bound on how much the individual actions can be correlated, or $\rho_a \leq 1 - \rho_{a\nu}^2$. Thus to the extent that

the individual agent's action is correlated with the private component, it also limits the extent to which the individual action can be related with the public component, as by construction, the private and the public component are independently distributed. In Section 6, we consider the role of prior information on the structure of the equilibrium set, and a natural case of prior information is that each agent knows his own payoff type $\theta_i = \theta + \nu_i$, but does not necessarily know the composition of his own payoff state in terms of the private and public component.

5.5 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, an 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 distributions, 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 of 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 (11) and (12), 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. 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 $\Sigma_{\theta,x,y,a,A}$:

$$\begin{bmatrix} \sigma_\theta^2 & \sigma_\theta^2 & \sigma_\theta^2 & \rho_{a\theta}\sigma_a\sigma_\theta & \rho_{a\theta}\sigma_a\sigma_\theta \\ \sigma_\theta^2 & \sigma_\theta^2 + \sigma_x^2 & \sigma_\theta^2 & \sigma_a\sigma_x\rho_{ax} + \sigma_a\sigma_\theta\rho_{a\theta} & \sigma_a\sigma_\theta\rho_{a\theta} \\ \sigma_\theta^2 & \sigma_\theta^2 & \sigma_\theta^2 + \sigma_y^2 & \sigma_a\sigma_y\rho_{ay} + \sigma_a\sigma_\theta\rho_{a\theta} & \sigma_a\sigma_y\rho_{ay} + \sigma_a\sigma_\theta\rho_{a\theta} \\ \rho_{a\theta}\sigma_a\sigma_\theta & \sigma_a\sigma_x\rho_{ax} + \sigma_a\sigma_\theta\rho_{a\theta} & \sigma_a\sigma_y\rho_{ay} + \sigma_a\sigma_\theta\rho_{a\theta} & \sigma_a^2 & \rho_a\sigma_a^2 \\ \rho_{a\theta}\sigma_a\sigma_\theta & \sigma_a\sigma_\theta\rho_{a\theta} & \sigma_a\sigma_y\rho_{ay} + \sigma_a\sigma_\theta\rho_{a\theta} & \rho_a\sigma_a^2 & \rho_a\sigma_a^2 \end{bmatrix}. \quad (39)$$

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 analyze the correlated equilibrium conditions as before. The best response function must satisfy:

$$a = r\mathbb{E}[A|a, x, y] + s\mathbb{E}[\theta|a, x, y] + u, \quad \forall a, x, y. \quad (40)$$

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 (40) for all realizations of a, x, y . 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 from (40), we find that we also have a corresponding increase in the number of equality constraints on the equilibrium conditions, from one to three. Indeed, we can determine $(\rho_{ay}, \rho_{ax}, \sigma_a)$ uniquely:

$$\sigma_a = \frac{\sigma_\theta s \rho_{a\theta}}{1 - \rho_a r}, \quad \rho_{ax} = \frac{\sigma_\theta \left((1 - \rho_a r) - \rho_{a\theta}^2 (1 - r) \right)}{\sigma_x \rho_{a\theta}}, \quad \rho_{ay} = \frac{\sigma_\theta}{\sigma_y \rho_{a\theta}} \left(\frac{1 - \rho_a r}{1 - r} - \rho_{a\theta}^2 \right). \quad (41)$$

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 r , but the informational externality s does not appear.

Consequently, the relation between the correlation coefficients ρ_{ax} and ρ_{ay} can be written, using the conditions (41) as $\rho_{ax}\sigma_x = \rho_{ay}\sigma_y(1 - r)$, where the factor $1 - r$ 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 matrix, or:

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

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

$$1 - \rho_a - \rho_{ax}^2 \geq 0, \quad (42)$$

$$\rho_a - \rho_{a\theta}^2 - \rho_{ay}^2 \geq 0. \quad (43)$$

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 (41) with (42), or with (43), 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 (41), we can describe the set of Bayes correlated equilibria exclusively in terms of correlation coefficients $(\rho_a, \rho_{a\theta})$.

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

$$C_x(\tau_x, r) \triangleq \{(\rho_a, \rho_{a\theta}) \in [0, 1] \times [-1, 1] \mid (\rho_a, \rho_{a\theta}, \rho_{ax}) \text{ satisfy (19), (41), (42)}\}.$$

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

$$C_y(\tau_y, r) \triangleq \{(\rho_a, \rho_{a\theta}) \in [0, 1] \times [-1, 1] \mid (\rho_a, \rho_{a\theta}, \rho_{ay}) \text{ satisfy (19), (41), (43)}\}.$$

The intersection of the private and the public equilibrium sets defines the Bayes correlated equilibria consistent with the prior information (τ_x, τ_y) :

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

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 (41) and (42) as an equality, given a lower




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

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 (41) and (43) 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 additional correlation device cannot undo the given private and public information, but rather provides additional correlation opportunities over and above those contained in (τ_x, τ_y) .

Proposition 8 (Prior Information)

For all $r \in (-\infty, 1)$:

1. The equilibrium set $C((\tau_x, \tau_y), r)$ is decreasing in (τ_x, τ_y) ;
2. The lowest correlation coefficient $(\rho_{a\theta}, \cdot) \in C((\tau_x, \tau_y), r)$, is increasing in (τ_x, τ_y) ;
3. The lowest correlation coefficient $(\rho_a, \cdot) \in C((\tau_x, \tau_y), r)$, is increasing in (τ_x, τ_y) .

Thus, as the precision of the prior information increases, the set of Bayes correlated equilibria shrinks. 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 precision of 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_x, τ_y) . This suggests that we identify the unique Bayes Nash equilibrium with information structure (τ_x, τ_y) and interaction term r in terms of the correlation coefficients $(\rho_a, \rho_{a\theta})$ as $B((\tau_x, \tau_y), r) \subseteq [0, 1] \times [-1, 1]$.

Corollary 3 (BCE and BNE with Prior Information)

For all (τ_x, τ_y) , we have:

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

In Section 5.4, we extended the analysis of the Bayes correlated equilibrium from an environment with pure common values to an environment with interdependent values. Similarly, we could extend the analysis of the prior information, pursued here in some detail for the environment with pure common values to the one with interdependent values.

7 Information Sharing and Information Structure

We are often interested in analyzing what is the best information structure in a strategic setting, either for the players in the game or for an outside observer who cares about choices in the game. For example, recent work by Rayo and Segal (2010) and Kamenica and Gentzkow (2011) have considered this problem in the context of single person games, i.e., decision problems; Bergemann and Pesendorfer (2007) characterizes the revenue-maximizing information structure in an auction with many bidders; and a large literature reviewed below has examined the incentives of competing firms to share cost and demand information. Directly maximizing over all possible information structures, especially with many players, sounds intractable. Our compact representation of the Bayes correlated equilibria allows us to assess the private and/or social welfare across the entire set of possible information structures (and induced equilibrium distributions). In this section, we show how results developed in earlier sections allows us to easily do this and deliver novel economic insights. In particular, we identify settings where the information structure that turns out to be optimal was excluded from the parametric domain of information structures analyzed in earlier work. In the context of the application of information sharing among firms, we show that it is optimal to have firms transmit all information that they have, but have that information observed with noise by other firms.

The problem of information sharing among firms was pioneered in work by Novshek and Sonnenschein (1982), Clarke (1983) and Vives (1984), who examined to what extent competing firms have an incentive to share information in an uncertain environment. In this strand of literature, which is surveyed in Vives

(1990) and embedded in a very general framework by Raith (1996), each firm receives a private signal about a source of uncertainty, say a demand or cost shock. The central question then is under which conditions the firms have an incentive to commit ex-ante to an agreement to share information in some form. A striking result by Clarke (1983) was the finding that in a Cournot oligopoly with uncertainty about a common parameter of demand, the firms will never find it optimally to share information. The complete lack of information sharing, independent of the precision of the private signal and the number of competing firms, is surprising. After all, it would be socially optimal to reduce the uncertainty about demand and a reasonable conjecture would be that the firms could at least partially appropriate the social gains of information. The result of Clarke (1983) appeared in the context of a linear inverse demand with normally distributed uncertainty, and a constant marginal cost. In subsequent work, the strong result of zero information sharing was shown to rely on constant marginal cost, and with a quadratic cost of production, it was shown that either zero *or* complete information sharing can be optimal, where the information sharing result appears when the cost of production is sufficiently convex for each firm, and hence information becomes more valuable, see Kirby (1988) and Raith (1996).

In the above cited work, the individual firms receive a private, idiosyncratic and noisy signal x_i about the state of demand θ . Each firm can commit to transmit the information, noisy or noiseless, to an intermediary, such as a trade association, which aggregates the information. The intermediary then discloses the aggregate information to the firms. Importantly, while the literature did consider the possibility of noisy *or* noiseless transmission of the private information, it a priori restricted the disclosure policy to be noiseless, which implicitly restrict the information policy to disclose the same, common signal to all the firms. An information policy is then a pair a information transmission and information disclosure policies. The present analysis of the Bayes correlated equilibrium allows us to substantially modify the earlier insights. Interestingly, Proposition 9 establishes that it is with substantial loss in generality to restrict attention to a common disclosure policy.

We described the payoffs of the quantity setting firms with uncertainty about demand in Example 2, where $s > 0$ represents the positive informational effect of a higher state θ of demand and $r < 0$ represents the fact the firms are producing (homogeneous) substitutes. An increase in the absolute value of the (negative) interaction parameter r then represents an increase in the slope the demand curve. We first ask what information structure maximizes firms' profits, by finding the firm optimal Bayes correlated equilibrium. We will then consider how to attain that information structure through information sharing.

Correlation of output with demand ($\rho_{a\theta}$) increases profits but correlation between firms' output (ρ_a) decreases profit. Thus it is always optimal to set $\rho_{a\theta}$ as high as possible consistent with BCE, and thus $\rho_{a\theta} = \sqrt{\rho_a}$. If the demand curve is sufficiently steep, it is optimal to have complete information but

otherwise there is an interior solution.

Proposition 9 (Information Sharing and Profit)

1. If $r \geq -1$, then the firm optimal BCE is achieved at $\rho_a = \rho_{a\theta} = 1$.
2. If $r < -1$, then the firm optimal BCE occurs with less than perfect correlation across actions:

$$\rho_a^* = -\frac{1}{r} < 1 \text{ and } \rho_{a\theta}^* = \sqrt{\rho_a^*} < 1. \tag{44}$$

We can now translate the structure of the profit maximizing Bayes correlated equilibrium into the corresponding Bayes Nash equilibrium and its associated information policy and information structure. Suppose that each of the continuum of firms receives only a private signal x_i with variance σ_θ^2 . We characterized in Section 6 the set of Bayes correlated equilibria consistent with prior private information of a certain precision. If there is one corresponding to the firm optimal BCE, then we can identify an information sharing technology that will attain the first best. If not, we can identify the firm second best BCE and how that can be achieved.

If all information is publicly shared, then we reach the complete information equilibrium with $\rho_a = \rho_{a\theta} = 1$. The first part of Proposition implies that full public disclosure is the optimal information policy if the slope of the demand curve is sufficiently low. But the second part of the Proposition indicates that the optimal disclosure policy may require noisy and idiosyncratic disclosure of the transmitted information, rather than noiseless disclosure as previously analyzed in the literature. In fact, if slope of the demand curve is sufficiently large, then the profit maximizing Bayes correlated equilibrium arises under the correlation coefficient of the actions, $\rho_a = \rho_{a\theta}^2 = -\frac{1}{r} < 1$. As we learned from Proposition 6, these are the correlation coefficients which maximizes the variance of the individual action, i.e. the individual supply decisions. Now if private signals were sufficiently accurate, ρ_a would already be too high even without any information transmission. But if private signals were not too accurate, then it will be possible to attain the firm optimal BCE. From Proposition 5, we know that $\rho_a = \rho_{a\theta}^2$ is at the boundary of the set of Bayes correlated equilibrium and that the boundary can only be reached with idiosyncratic information, i.e. information which is conditionally on state θ independent across agents. Thus the optimal disclosure policy requires noisy and idiosyncratic disclosure of the transmitted information, rather than noiseless disclosure as previously assumed in the literature.

Proposition 10 (Noisy and Idiosyncratic Disclosure Policy)

1. If $r \geq -1$, then full disclosure is firm optimal.

2. If $r < -1$ and

(a) if $-\frac{1}{r} > \sigma_\theta^2 / (\sigma_x^2 + \sigma_\theta^2)$, then the firm optimal disclosure policy is to have each firm observe a noisy signal of the average of their private signals (which equals the true state).

(b) if $-\frac{1}{r} \leq \sigma_\theta^2 / (\sigma_x^2 + \sigma_\theta^2)$, then no disclosure is firm optimal.

The sharing of the private information impacts the profit of the firms through two channels. First, shared information about level of demand improves the supply decision of the firms, and unambiguously increases the profits. Second, shared information increases the correlation in the strategies of the actions. In an environment with strategic substitutes, this second aspect is undesirable from the point of view of each individual firm. Now, the literature only considered noiseless disclosure. In the context of our analysis, this represents a public signal; after all a noiseless disclosure means that all the firms receive the same information. Thus, the choice of the optimal disclosure regime can be interpreted as the choice of the precision τ_y of the public signal, and hence a point along a level curve for a given τ_x , see Figure 4. But now we realize that the disclosure in form of a public signal requires a particular trade-off between the correlation coefficient ρ_a across actions and the correlation $\rho_{a\theta}$ of action and state. In particular, an increase in the correlation coefficient $\rho_{a\theta}$ is achieved only at the cost of substantially increasing the undesirable correlation across actions. This trade-off, necessitated by the public information disclosure, meant that the optimal disclosure is either to not disclose *any* information or disclose *all* information. The present analysis suggests a more subtle result which is to disclose some information, so that the private information of all the firms is improved, but to do so in way that does not increase the correlation across actions more than necessary. This is achieved by an idiosyncratic, that is private and noisy disclosure policy, which necessarily does not reveal all the private information of the agents, as they would otherwise achieve complete correlation in their action.

The very last result of Proposition 10 reaffirms the earlier result of Clarke (1983), which presented conditions under which zero information transmission was optimal. The necessary and sufficient condition for zero information transmission:

$$\rho_a^* = -\frac{1}{r} < \sigma_\theta^2 / (\sigma_x^2 + \sigma_\theta^2), \quad (45)$$

is best understood in light of the role of the prior information. We established in Proposition 8 that the lowest correlation coefficient is increasing in the precision of the prior information τ_x and τ_y . Now, if the precision of the public signal is zero or $\tau_y = 0$, then the correlation coefficient induced by the private signal

with precision $\tau_x = \sigma_x^{-2}$ is given by the right hand side of (45). In other words, if the profit maximizing level of correlation ρ_a , is below the level already induced by the prior information τ_x , then, but only then, do the firms prefer zero information transmission and disclosure. We should mention that in contrast to the literature, we present and establish the above results, in line with rest of the present analysis, for the environment with a continuum of firms. However, the results carry over to the environment with a finite number of firms as the only relevant determinant is the structure of the best response as discussed in Section 3. The only modification that arises in the analysis with finite number of firms is the extent of the correlation $\rho_{a\theta}$ with respect to the state θ . If there are only a finite number of firms, and hence only a finite number of signals about the true state of the world, then even complete sharing of the available information will not allow the firms to achieve $\rho_{a\theta} = 1$, even though their actions will be completely correlated or $\rho_a = 1$. The finite information then acts as a constraint on the amount of information shared, but does not affect the preference for or against information sharing.

8 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, the structural parameters of the game (r, s, u) , by the observed variables? We are particularly interested in how the identification of the structural parameters (r, s, u) is influenced by the solution concept, and hence the specification of the private information of the agents as known to the analyst.

Now, identification depends critically on what types of data are available. Here, we consider the possibility of identification with individual data and assume that the econometrician observes the realized individual actions a_i and the realized state θ .³ 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_x, τ_y) of the agents.

³In Bergemann and Morris (2011b), we also analyze the robust identification with aggregate data. As a leading example we consider the canonical problem of demand and supply identification. The identification in the linear demand and supply model relies on the aggregate data, namely market quantity and market price. In contrast to the received work on identification in the demand and supply model we allow for incomplete information by the market participants about the cost and demand factors.

For a given information structure (τ_x, τ_y) and observed moments of the Bayes Nash equilibrium distribution, $(\mu_a, \sigma_a, \rho_a, \sigma_\theta, \rho_{a\theta})$, we can identify the weights on the private signal and the public signal, α_x^* and α_y^* , directly from the variance of the (aggregate) action and the covariance of the (aggregate) action with the state, see (18). Now, we can use the property of the equilibrium strategy, namely that the ratio of the weights is exactly equal the precision of the private and public signal, deflated by their (strategic) weight, see (15):

$$\frac{\alpha_x^*}{\alpha_y^*} = \frac{\tau_x}{\tau_y} (1 - r).$$

Thus given the knowledge of the information structure, we can infer the sign of the strategic interaction term r from the ratio of the linear weights, α_x^* and α_y^* . In particular, we can determine how much of the variance in the action, individual or aggregate, is attributable to the private and the public signal respectively. Given the *known* strength of the signals, the covariance of the action and the state then allow us to identify the slope of the equilibrium response. We thus find that the parameters of equilibrium response and the sign of the interaction parameters are identified for every *known* information structure of the game.

Proposition 11 (Point Identification in BNE)

The Bayes Nash equilibrium outcomes with information structure (τ_x, τ_y) ,

1. *identifies the informational externality s ;*
2. *identifies the strategic interaction r if $0 < \tau_x, \tau_y < \infty$; and*
3. *identifies the equilibrium slope and equilibrium intercept, the ratios $s/(1 - r)$ and $u/(1 - r)$.*

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:

$$\mu_a = \frac{u + \mu_\theta s}{1 - r}, \quad \sigma_a = \frac{\sigma_\theta s \rho_{a\theta}}{1 - \rho_a r}. \tag{46}$$

We thus have two restrictions to identify the three unknown structural parameters (r, s, u) . We can solve for two of the unknowns in terms of the remaining unknowns. In particular, when we solve for (s, u) in

terms of the remaining unknown r , we obtain expressions for the equilibrium intercept and the equilibrium slope in terms of the moments and the remaining unknown structural parameters:

$$\frac{u}{1-r} = \mu_a - \frac{\sigma_a \mu_\theta (1 - \rho_a r)}{\sigma_\theta (1-r) \rho_{a\theta}}, \quad \frac{s}{1-r} = \frac{\sigma_a}{\rho_{a\theta} \sigma_\theta} \frac{1 - \rho_a r}{1-r}. \quad (47)$$

Now, except for the case of $\rho_a = 1$, in which the actions of the agents are perfectly correlated, we find that the ratio on the left hand side is not uniquely determined. As the strategic interaction parameter r can vary, or $r \in (-\infty, 1)$, it follows that we can only partially identify the above ratios, namely,

$$\frac{u}{1-r} \in \begin{cases} \left(-\infty, \mu_a - \frac{\mu_\theta \rho_a \sigma_a}{\rho_{a\theta} \sigma_\theta}\right) & \text{if } \frac{\mu_\theta}{\rho_{a\theta}} > 0; \\ \left(\mu_a - \frac{\mu_\theta \rho_a \sigma_a}{\rho_{a\theta} \sigma_\theta}, \infty\right) & \text{if } \frac{\mu_\theta}{\rho_{a\theta}} < 0; \end{cases} \quad (48)$$

and the above ratio is point-identified if $\mu_\theta = 0$. Similarly,

$$\frac{s}{1-r} \in \begin{cases} \left(\frac{\rho_a \sigma_a}{\rho_{a\theta} \sigma_\theta}, \infty\right) & \text{if } \rho_{a\theta} > 0; \\ \left(-\infty, \frac{\rho_a \sigma_a}{\rho_{a\theta} \sigma_\theta}\right) & \text{if } \rho_{a\theta} < 0. \end{cases} \quad (49)$$

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

Proposition 12 (Set Identification in BCE)

The Bayes correlated equilibrium outcomes:

1. identify the sign of the informational externality s ;
2. do not identify the sign of the strategic interaction r ;
3. identify a set of equilibrium slopes, given by (49), if $\rho_a < 1$.

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 r ; second, we can identify only a set of possible interaction ratios. Given the sharp differences in the identification under Bayes Nash and Bayes correlated equilibrium, we now try to provide some intuition as to the source of the contrasting results. In the identification under the hypothesis of the Bayes correlated equilibrium, the econometrician observes and uses the same data as under the Bayes Nash equilibrium, but does not know anymore how precise or noisy the information of the agents is. Thus, the econometrician now face an attribution problem as the observed covariance between the action and the state could be large either because the individual preferences are very responsive to the state, i.e. s is large, or because the agents have very precise information about the state and hence respond strongly to the precise information, even though they are only moderately sensitive to the state, i.e. s is low.

This attribution problem, which is present when the agent’s information structure is not known, is often referred to as “attenuation bias” in the context of individual decision making. The basic 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, the “attenuation bias”. Yet in the single agent model, the sign of the parameter of interest, the informational externality s remains correctly identified, even though the information externality is set-identified rather than point-identified. Importantly, 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.

Given the lack of identification in the absence of knowledge regarding the information structure, it is natural to ask whether prior information can improve the identification of the structural parameters, just as prior information could improve the equilibrium prediction. In Section 6, we showed that the knowledge of the information structure (τ_x, τ_y) systematically restricts the equilibrium predictions of the coefficients $(\rho_a, \rho_{a\theta})$. Now, as we consider the identification of the structural parameters, we might use the knowledge of the information structure (τ_x, τ_y) together with the data to restrict the set of structural parameters consistent with the data and the prior information (τ_x, τ_y) . In the working paper, Bergemann and Morris (2011b), we formally analyze how the set of possible equilibrium coefficients $(\rho_a, \rho_{a\theta})$ depends on the prior information (τ_x, τ_y) and the interaction parameter r . This, then allows us ask which values of the interaction parameter r are consistent with the observed data, in particular the equilibrium correlation coefficients $(\rho_a, \rho_{a\theta})$. We show that the set-identification improves with an increase in the precision of the prior information and converges to point-identification as the precision of the prior information becomes arbitrarily large.

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 (2011). 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 interesting extension in the present setting would be to limit the identification to a certain subset of parameters, say the interaction term r , but then identify the distribution of the states of the world rather than assuming the observability of the states. 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 the public observables, is maintained. For example, in Paula and Tang (2011), 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 (2011) 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.

9 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.

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 discontinuous payoffs. For example, it would be of considerable interest to ask how the allocations and the revenues differ across belief environments and auction formats. In ongoing work Bergeman, Brooks, and Morris (2011) consider a private value environment in first price auction format. Abraham, Athey, Babaioff, and Grubb (2011) trace the implications of different information structures in a common value environment in a second price auction format.

Finally, we could 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.

10 Appendix

Proof of Proposition 1. 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

$$\mathbb{E}_{\widehat{\psi \circ \pi}(\cdot|t)} u(a, g \circ \sigma, \theta) \geq \mathbb{E}_{\widehat{\psi \circ \pi}(\cdot|t)} u(a', g \circ \sigma, \theta) \quad (50)$$

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 (50) establishes that $\psi \circ \pi \circ \sigma$ is a BCE. ■

Proof of Proposition 4. 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 $\sigma_\theta^2, \sigma_x^2$ and σ_y^2 as:

$$\rho_{a\theta} = \pm \frac{\sigma_\theta (\alpha_x + \alpha_y)}{\sqrt{\alpha_x^2 \sigma_x^2 + \alpha_y^2 \sigma_y^2 + \sigma_\theta^2 (\alpha_x + \alpha_y)^2}}, \quad (51)$$

and

$$\rho_a = \frac{\alpha_y^2 \sigma_y^2 + \sigma_\theta^2 (\alpha_x + \alpha_y)^2}{\alpha_x^2 \sigma_x^2 + \alpha_y^2 \sigma_y^2 + \sigma_\theta^2 (\alpha_x + \alpha_y)^2}. \quad (52)$$

It now follows immediately from (51) - (52), and the formulae of α_x^* and α_y^* , see (14), that we can recover the corresponding information structure (τ_x, τ_y) of the Bayes Nash equilibrium as

$$\sigma_x = \frac{((1 - \rho_a r) - \rho_{a\theta}^2 (1 - r)) \sigma_\theta}{\sqrt{1 - \rho_a |\rho_{a\theta}|}},$$

and

$$\sigma_y = \frac{((1 - \rho_a r) - \rho_{a\theta}^2 (1 - r)) \sigma_\theta}{\sqrt{\rho_a - \rho_{a\theta}^2 |\rho_{a\theta}| (1 - r)}},$$

which completes the proof. ■

Proof of Proposition 6. The variance σ_a^2 is given by (30), and inserting $\rho_a = \rho_{a\theta}^2$ we obtain $\sigma_\theta s \rho_{a\theta} / (1 - \rho_{a\theta}^2 r)$, which is maximized at $|\rho_{a\theta}| = \sqrt{-1/r}$, or $\rho_a = -1/r$. ■

Proof of Proposition 7. (1.) The volatility σ_A^2 , which is given by:

$$\rho_a \sigma_a^2 = \rho_a \left(\frac{\sigma_{\theta s} \rho_{a\theta}}{1 - \rho_a r} \right)^2,$$

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, respectively:

$$\frac{\sigma_{\theta}^2 \rho_{a\theta}^2 s^2}{(1 - \rho_a r)^3} (1 + \rho_a r),$$

where the later is positive if and only if

$$(1 + \rho_a r) \geq 0 \Leftrightarrow r \geq -\frac{1}{\rho_a},$$

and

$$\frac{2\rho_a |\rho_{a\theta}| \sigma_{\theta}^2 s^2}{(1 - \rho_a r)^2} > 0.$$

(2.) The dispersion, using (33), is given by:

$$(1 - \rho_a) \sigma_a^2 = (1 - \rho_a) \left(\frac{\sigma_{\theta} \rho_{a\theta} s}{1 - \rho_a r} \right)^2,$$

and it follows that the dispersion is increasing in $|\rho_{a\theta}|$. The dispersion 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

$$-\frac{\sigma_{\theta}^2 \rho_{a\theta}^2 s^2 (1 - r - (1 - \rho_a) r)}{(1 - \rho_a r)^3}.$$

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

$$(1 - \rho_a) \sigma_a^2 = (1 - \rho_a) \rho_a \left(\frac{\sigma_{\theta} s}{1 - \rho_a r} \right)^2,$$

and the dispersion reaches an interior maximum at $\rho_a = 1/(2 - r) \in (0, 1)$, irrespective of the nature of the game. ■

Proof of Proposition 8. We form the conditional expectation using (39) and the equilibrium conditions for the Bayes correlated equilibrium are then given by (40) and the solution to these equations is given by (41).

(1.) The equilibrium set is described as the set which satisfies the inequalities (42) and (43), where the correlation coefficients ρ_{ax}^2 and ρ_{ay}^2 appear separately. By determination of (41), 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 increases.

(2.) The lowest value of the correlation coefficient $\rho_{a\theta}$ is achieved when the inequalities (42) and (43) 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.

(3.) The lowest value of the correlation coefficient ρ_a is achieved when the inequality (43) 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. ■

Proof of Proposition 9. The ex post profit of the firm is given by:

$$(s\theta + rA)a + ua - \frac{1}{2}a^2,$$

and the interim expected profit is the above expectation and consists of terms that depend on the means μ_a and μ_θ plus

$$\sigma_a^2 \left(s\rho_{a\theta} \frac{\sigma_\theta}{\sigma_a} + r\rho_a - \frac{1}{2} \right).$$

Using the restriction on the variance of the individual action:

$$\sigma_a = \frac{\sigma_\theta s \rho_{a\theta}}{1 - r\rho_a},$$

we get

$$\frac{\sigma_\theta^2 s^2 \rho_{a\theta}^2}{2(1 - r\rho_a)^2} \tag{53}$$

The remaining restriction of the Bayes correlated equilibrium, see Proposition 5, is that $\rho_{a\theta}^2 \leq \rho_a$, and hence (53) can be rewritten as

$$\sigma_\theta^2 s^2 \frac{\rho_a}{2(1 - r\rho_a)^2},$$

i.e., it is always optimal to set the correlation coefficient $\rho_{a\theta}$ so that $\rho_{a\theta}^2 = \rho_a$. The relevant first order condition w.r.t. to ρ_a is given by:

$$\sigma_\theta^2 s^2 \frac{(1 + \rho_a r)}{(1 - \rho_a r)^3} = 0.$$

It follows that if $r > -1$, then there is no interior solution and the profit maximizing BCE is given by $\rho_a = \rho_{a\theta} = 1$. On the other hand, if $r < -1$, then the maximum is at an interior value of ρ_a :

$$\rho_a = -\frac{1}{r} < 1.$$

The validity of the second order conditions can be verified easily. ■

Proof of Proposition 10. By Proposition 9, if $r \geq -1$, then the profit maximizing equilibrium allocation requires $\rho_{a\theta} = \rho_a = 1$. Now, the Bayes Nash equilibrium associated with this correlation

structure requires that the agents have complete information about θ , but clearly with a large number of firms, here a continuum, this can be achieved by completely disclosing the private information of each individual firm (provided that $\sigma_x^2 < \infty$).

On the other hand, if $r < -1$, then the interior solution requires that $\rho_a < 1$ and $\rho_{a\theta}^2 = \rho_a$. By Proposition 2, we know that such a correlation structure can be achieved in the Bayes Nash equilibrium if and only if the agents make decisions on the basis of a private signal only, i.e. the variance of the public signal is required to be infinite. This in turn can be achieved if each agent receives information about the true state with an idiosyncratic noise, and hence with a private signal, which necessitates idiosyncratic and noisy information disclosure. Finally, given the initial private information of the agents, represented by σ_x^2 , we only need to complement the initial information if it does not already achieve or exceed $\rho_a = -1/r$. From (17), we find that the correlation coefficient in the Bayes Nash equilibrium without additional information beyond σ_x^2 is given by $\rho_a = \sigma_\theta^2 / (\sigma_\theta^2 + \sigma_x^2)$, which establishes the critical value for information sharing. ■

Proof of Proposition 11. (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 (16), say:

$$\alpha_x^2 = \frac{\sigma_a^2 - \sigma_A^2}{\sigma_x^2}, \quad \alpha_y^2 = \frac{\sigma_A^2 (1 - \rho_{A\theta}^2)}{\sigma_y^2}. \quad (54)$$

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 s , given the composition of the equilibrium coefficients α_x^* and α_y^* of the Bayes Nash equilibrium, see (14).

(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:

$$\alpha_y^* = \alpha_x^* \frac{\sigma_x^2}{\sigma_y^2} \frac{1}{1 - r}.$$

Now, if $0 < \sigma_x^2, \sigma_y^2 < \infty$, then we can identify r .

(3.) Given the identification of α_x^* and α_y^* , we can identify the ratios $u/(1 - r)$ and $s/(1 - r)$. 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. From the equilibrium conditions, see (14), we have the values of μ_a, α_x and α_y . This allows us to solve for r, s, u as a function of $\mu_a, \alpha_x, \alpha_y$:

$$\begin{aligned} u &= -\frac{\alpha_x^2 \sigma_x^2 \mu_\theta - \alpha_x \mu_a \sigma_x^2 - \alpha_x \alpha_y \sigma_y^2 \mu_\theta + \alpha_x \alpha_y \tau \sigma_x^2 \sigma_y^2 \mu_\theta}{\alpha_y \sigma_y^2}, \\ r &= \frac{\alpha_y \sigma_y^2 - \alpha_x \sigma_x^2}{\alpha_y \sigma_y^2}, \\ s &= \frac{\alpha_x^2 \sigma_x^2 - \alpha_x \alpha_y \sigma_y^2 + \alpha_x \alpha_y \tau \sigma_x^2 \sigma_y^2}{\alpha_y \sigma_y^2}. \end{aligned} \quad (55)$$

If we form the ratios $u/(1-r)$ and $s/(1-r)$ with the expressions on the rhs of (55), then we obtain expressions which do only depend on the observable data, and are hence point identified, and in particular

$$\frac{u}{1-r} = -\frac{-\mu_a\sigma_x^2 + \alpha_x\sigma_x^2\mu_\theta - \alpha_y\sigma_y^2\mu_\theta + \tau\sigma_x^2\alpha_y\sigma_y^2\mu_\theta}{\sigma_x^2}, \quad (56)$$

and

$$\frac{s}{1-r} = \frac{\alpha_x\sigma_x^2 - \alpha_y\sigma_y^2 + \tau\alpha_y\sigma_x^2\sigma_y^2}{\sigma_x^2}, \quad (57)$$

which completes the proof of identification. We observe that, using (54), we could express the ratios (56) and (57) entirely in terms of the first two moments of observed data. ■

Proof of Proposition 12. (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 (46). Given the information on left hand side and the information of $\rho_{a\theta}$, we can infer the sign of s .

(2.) Even though the sign of s can be established, we cannot extract the unknown variables on the rhs of (46) in the presence of the linear return term u , and hence it follows that we cannot sign r .

(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 on mean and variance, see (46). We thus have two equations to identify the three unknown structural parameters (r, s, u) . We can solve for (s, u) in terms of the remaining unknown r to obtain:

$$u = \frac{-\sigma_a\mu_\theta + \sigma_a\rho_a\mu_\theta r - \mu_a\sigma_\theta r\rho_{a\theta} + \mu_a\sigma_\theta\rho_{a\theta}}{\sigma_\theta\rho_{a\theta}}, \quad s = \frac{\sigma_a(1 - \rho_a r)}{\sigma_\theta\rho_{a\theta}}.$$

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

$$\frac{u}{1-r} = -\mu_a + \frac{\sigma_a\mu_\theta(1 - \rho_a r)}{\sigma_\theta(1-r)\rho_{a\theta}}, \quad \frac{s}{1-r} = -\frac{\sigma_a}{\rho_{a\theta}\sigma_\theta} \frac{(1 - \rho_a r)}{1-r},$$

in terms of the observables. But, except for the case of $\rho_a = 1$, we see that this is not the case. As $r \in (-\infty, 1)$, it follows that we can only partially identify the above ratios, namely (48) and (49) which describe the respective sets into which each ratio can be identified. ■

References

- ABRAHAM, I., S. ATHEY, M. BABAIOFF, AND M. GRUBB (2011): “Peaches, Lemons and Cookies: Designing Auction Markets with Dispersed Information,” Discussion paper, Microsoft and Harvard University and MIT.
- ANGELETOS, G., AND A. PAVAN (2007): “Efficient Use of Information and Social Value of Information,” *Econometrica*, 75, 1103–1142.
- (2009): “Policy with Dispersed Information,” *Journal of the European Economic Association*, 7, 11–60.
- 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.
- BERGEMAN, D., B. BROOKS, AND S. MORRIS (2011): “First Price Auction and Private Information,” Discussion paper, Yale University and Princeton University.
- BERGEMANN, D., AND S. MORRIS (2005): “Robust Mechanism Design,” *Econometrica*, 73, 1771–1813.
- (2011a): “Correlated Equilibrium in Games of Incomplete Information,” Discussion paper, Cowles Foundation for Research in Economics, Yale University and Princeton University.
- (2011b): “Robust Predictions in Games of Incomplete Information,” Discussion paper, 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.
- CLARKE, R. (1983): “Collusion and the Incentives for Information Sharing,” *Bell Journal of Economics*, 14, 383–394.
- DASGUPTA, P., AND E. MASKIN (2000): “Efficient Auctions,” *Quarterly Journal of Economics*, 115, 341–388.
- FORGES, F. (1993): “Five Legitimate Definitions of Correlated Equilibrium in Games with Incomplete Information,” *Theory and Decision*, 35, 277–310.
- 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 (2011): “Bayesian Persuasion,” *American Economic Review*, 101, 2590–2615.
- KIRBY, A. (1988): “Trade Associations as Information Exchange Mechanisms,” *RAND Journal of Economics*, 19, 138–146.
- 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.
- NOVSHEK, W., AND H. SONNENSCHN (1982): “Fulfilled Expectations Cournot Duopoly with Information Acquisition and Release,” *Bell Journal of Economics*, 13, 214–218.

- PAULA, A. D., AND X. TANG (2011): “Inference of Signs of Interaction Effects in Simultaneous Games with Incomplete Information,” *Econometrica*, forthcoming.
- PERRY, M., AND P. RENY (2002): “An Ex Post Efficient Auction,” *Econometrica*, 70, 1199–1212.
- RAITH, M. (1996): “A General Model of Information Sharing in Oligopoly,” *Journal of Economic Theory*, 71, 260–288.
- RAYO, L., AND I. SEGAL (2010): “Optimal Information Disclosure,” *Journal of Political Economy*, 118, 949–987.
- SWEETING, A. (2009): “The Strategic Timing of Radio Commercials: An Empirical Analysis Using Multiple Equilibria,” *RAND Journal of Economics*, 40, 710–742.
- VIVES, X. (1984): “Duopoly Information Equilibrium: Cournot and Bertrand,” *Journal of Economic Theory*, 34, 71–94.
- (1988): “Aggregation of Information in Large Cournot Markets,” *Econometrica*, 56, 851–876.
- (1990): “Nash Equilibrium with Strategic Complementarities,” *Journal of Mathematical Economics*, 19, 305–321.
- (2011): “A Large Market Rational Expectations Equilibrium Model,” Discussion paper, IESE Business School.