

**ROBUST PREDICTIONS
IN GAMES WITH INCOMPLETE INFORMATION**

By

Dirk Bergemann and Stephen Morris

**September 2011
Revised December 2011
Revised October 2012**

COWLES FOUNDATION DISCUSSION PAPER NO. 1821RR



**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[‡]

September 26, 2012

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. The present paper formed the basis of the Fisher-Schultz Lecture given by Stephen Morris at the European Econometric Society Meeting 2012. We benefitted from comments of Steve Berry, Vincent Crawford, Matthew Gentzkow, Phil Haile, Emir Kamenica, Marc Henry, Arthur Lewbel, Larry Samuelson, Elie Tamer, and Xavier Vives, as well as 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

Suppose that some economic agents each have a set of feasible actions that they can take and their payoffs depend on the actions that they all take and a payoff state with a known distribution. Call this scenario the *payoff environment*. To analyze behavior in this setting, we also have to specify what agents believe about the payoff states, what others believe, and so on. Call this the *belief environment*. A standard incomplete information game consists of a combination of a payoff environment and a belief environment. Instead of asking what happens in a fixed incomplete information game, in this paper we will characterize what may happen in a given payoff environment for any given belief environment. In particular, we identify which outcomes, i.e., probability distributions over action profiles and payoff states, could arise in a Bayes Nash equilibrium for a fixed payoff environment and for some belief environment.

There are a number of reasons why this exercise is both tractable and interesting.

The belief environment will generally be very hard to observe, as it is in the agents' minds and does not necessarily have an observable counterpart. We know that outcomes are very sensitive to the belief environment (Rubinstein (1989), Kajii and Morris (1997) and Weinstein and Yildiz (2007)). If we can characterize equilibrium outcomes independent of the information structure, we can identify *robust predictions* for a given payoff environment which are independent of - and in that sense robust to - the specification of the belief environment.

Conversely, if we are able to identify a mapping from payoff environments to outcomes which does not depend on the belief environment, then we can also study the inverse of the map, seeing which payoff environments are consistent with an observed outcome. This mapping gives us a framework for partially identifying the payoff environment without assumptions about the belief environment. Thus we can carry out *robust identification* of the payoff environment.

Characterizing the set of all equilibria for all belief environments sounds daunting, but it turns out that it is often easier to characterize what happens for all or many information structures at once than it is for a fixed information structure. Suppose that instead of explicitly modelling the belief environment, we use the classical game theoretic metaphor of a mediator who makes private action recommendations to the agents. In particular, suppose that a mediator was able to make private, perhaps correlated, action recommendations to the agents as a function of the payoff state. If the agents have an incentive to follow the mediator's recommendation, we say that the resulting joint distribution of payoffs and actions is a *Bayes correlated equilibrium*. We can show that the set of Bayes correlated equilibria for a given payoff environment correspond one to one with the set of Bayes Nash equilibria that could arise for any belief environment.

The Bayes correlated equilibrium characterization can also be used to analyze the *strategic value of*

information. Economists are often interested in analyzing which information structure is best for some welfare measure or some subset of agents in a given setting. In analyzing such problems, it is usual to focus on a low dimensional parameterized set of information structures because working with all information structures seems intractable. Our results suggest an alternative approach: one can find the Bayes correlated equilibrium that maximizes some objective, and then reverse engineer the information structure that generates that distribution as a Bayes Nash equilibrium.

In Bergemann and Morris (2011a), we pursue this research agenda for general games. In this paper, we examine these issues in a tractable payoff environment with a continuum of players, symmetric payoff functions, linear best response functions and normally distributed uncertainty. We consider a tractable belief environment, consisting of a noisy private and a noisy public signal of the payoff relevant state. This combination of tractable payoff environment and belief environment is widely studied, see Morris and Shin (2002) and Angeletos and Pavan (2007) among many others. The analysis in this paper provides a powerful illustration of the logic and usefulness of the more general approach, as well as providing new results about an important economic environment that is widely used in economic applications. Symmetry and normality assumptions are maintained throughout the analysis.

Bayes correlated equilibria in this environment are symmetric normal distributions of the state and the actions in the continuum population with the "obedience" property that a player with no information beyond the action that he is to play would not have an incentive to choose a different action. We compare Bayes correlated equilibria with Bayes Nash equilibria for every belief environment in our bivariate class of information structures. Integrating out the agents' signals, we show that each information structure and its Bayes Nash equilibrium gives rise to a Bayes correlated equilibrium. Conversely, each Bayes correlated equilibrium corresponds to the unique Bayes Nash equilibrium for some information structure in the bivariate class. This result illustrates the more general equivalence in Bergemann and Morris (2011a), within the class of symmetric normal distributions. Because Bayes correlated equilibria are two dimensional in this environment, a simple two dimensional class of information structures is large enough to reach all Bayes correlated equilibria.

We can identify robust predictions in terms of restrictions on the first and second moments of the joint distribution over actions and state. 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. 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 mean of the individual actions (and of the population action) 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.

We show how our approach can be used to analyze the strategic value of information by considering 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.

Our benchmark model contrasts two extremes: either nothing is known about the belief environment, or it is perfectly known. For both robust prediction and robust identification results, it is natural to consider intermediate cases where there is partial information about the belief environment. In particular, we 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.

We use our characterization of what happens in intermediate belief environments to analyze the robust identification question in depth. 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. With no restrictions on the belief 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.¹ However, we also

¹The 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).

identify conditions on the belief environment under which we are able to identify the sign of the interaction parameter.

This present work examines how the analysis of fixed games can be made robust to informational assumptions. This work parallels work in robust mechanism design, where games are designed so that equilibrium outcomes are robust to informational assumptions (our own work in this area beginning with Bergemann and Morris (2005) is collected in Bergemann and Morris (2012)). While the endogeneity of the game design makes the issues in the robust mechanism design literature quite different, in both literatures informational robustness can be studied with richer, more global, perturbations of the informational environment and more local ones. This paper is very permissive in allowing for a rich class of belief environments but less permissive in restricting attention to common prior belief environments.

The remainder of the paper is organized as follows. Section 2 defines the payoff environment, a class of quadratic games with normally distributed uncertainty, and the belief environment. We also define the relevant solution concepts, namely Bayes Nash equilibrium and Bayes correlated equilibrium. Section 3 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 4 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 5 we consider the optimal sharing of information among firms. In Section 6 we analyze how prior restrictions about the belief environment can further restrict the equilibrium predictions. 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 7, we turn from prediction to the issue of identification. Section 8 discusses some possible extensions and offers concluding remarks. The Appendix collects the proofs from the main body of the text.

2 Set-Up

2.1 Payoff Environment

There is a continuum of players and an individual player is indexed by $i \in [0, 1]$. Each player chooses an action $a_i \in \mathbb{R}$. The average action of all players is represented by $A \in \mathbb{R}$ and is the integral:

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

There is a payoff relevant state $\theta \in \Theta$ with a prior distribution $\psi \in \Delta(\Theta)$. All players have the same payoff function

$$u : \mathbb{R} \times \mathbb{R} \times \Theta \rightarrow \mathbb{R}, \quad (2)$$

where $u(a, A, \theta)$ is a player's payoff if she chooses action a , the average (or population) action is A and the state is θ .

A payoff environment is thus parameterized by (u, ψ) .

We will restrict attention to 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 A and the payoff relevant state θ :

$$a_i = r\mathbb{E}_i[A] + s\mathbb{E}_i[\theta] + k, \quad (3)$$

where $r, s, k \in \mathbb{R}$ are the parameters of the best response function and are assumed to be identical across players. 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 to the state of the world θ , but is not concerned about his interaction with the other players.

The parameter s represents the informational sensitivity of player i , the responsiveness to the state θ , 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

$$r \in (-\infty, 1). \quad (4)$$

Under this assumption,² there is a unique Nash equilibrium of the game with complete information given by:

$$a_i(\theta) = \frac{k}{1-r} + \frac{s}{1-r}\theta, \text{ for all } i \text{ and } \theta. \quad (5)$$

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

$$\theta \sim N(\mu_\theta, \sigma_\theta^2). \quad (6)$$

The present environment of linear best response and normally distributed uncertainty encompasses a wide class of interesting economic environments. The following four applications are prominent examples and we shall return to them throughout the paper to illustrate some of the results.

²If $r > 1$, the Nash equilibrium is unstable and, if actions sets were bounded, there would be multiple Nash equilibria.

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:

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

The weight r reflects concern for the average action A taken in the population. Morris and Shin (2002) analyze the Bayes Nash equilibrium for the belief environment analyzed in this paper. In terms of our notation, the beauty contest model set $s = 1 - r$ and $k = 0$ 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 $p(A) = s\theta + rA + k$ 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

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

In an alternative interpretation, we can have a common cost shock, so the demand curve is $p(A) = rA + k$ 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 $u(a_i, A, \theta)$, which depends on the individual action a_i , the average action A and the payoff state $\theta \in \mathbb{R}$:

$$u(a_i, A, \theta) \triangleq \frac{1}{2} \begin{pmatrix} a_i \\ A \\ \theta \end{pmatrix}' \begin{pmatrix} U_{aa} & U_{aA} & U_{a\theta} \\ U_{aA} & U_{AA} & U_{A\theta} \\ U_{a\theta} & U_{A\theta} & U_{\theta\theta} \end{pmatrix} \begin{pmatrix} a_i \\ A \\ \theta \end{pmatrix}, \quad (7)$$

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”):

$$-U_{aA}/U_{aa} < 1 \Leftrightarrow U_{aa} + U_{aA} < 0. \quad (8)$$

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

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

Their restriction (8) is equivalent to the present restriction (4). 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 highlighted in Angeletos and Pavan (2009)), but for the welfare analysis discussed in this paper they are not and can be uniformly set to zero.

Example 4 (Quadratic Economies with a Finite Number of Agents) In the case of a finite number I of players, the average action of all players but i is represented by the sum:

$$A \triangleq \frac{1}{I-1} \sum_{j \neq i} a_j. \tag{9}$$

With the linear best response (3), 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.

2.2 Belief Environment

Starting with the payoff environment (u, ψ) described in the previous subsection, we add a description of the belief environment, i.e., what players know about the state and others' beliefs. We first introduce the relevant belief (or signal) environment for this continuum agent economy. Each player is assumed to observe a two-dimensional signal. In the first dimension, the signal x_i is privately observed and idiosyncratic to the player i , whereas in the second dimension, the signal y is publicly observed and common to all the players:

$$x_i = \theta + \varepsilon_i, \quad y = \theta + \varepsilon. \tag{10}$$

The respective random variables ε_i and ε are assumed to be 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 . The type of each player is therefore given by the pair of signals: (x, y_i) . 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, σ_x^2 and σ_y^2 . 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.

2.3 Solution Concepts

We first describe the standard approach to analyze games of incomplete information by means of a fixed belief environment (or type space) and the associated Bayes Nash equilibria. We then introduce the notion of Bayes correlated equilibrium which will be defined as a function of the payoff environment alone.

2.3.1 Bayes Nash Equilibrium

The payoff environment and the belief environment together define a game of incomplete information. A symmetric pure strategy in the game is then defined by $s : \mathbb{R}^2 \rightarrow \mathbb{R}$.

Definition 1 (Bayes Nash Equilibrium)

A (symmetric) pure strategy s is a Bayes Nash equilibrium (BNE) if

$$\mathbb{E}[u(s(x_i, y), A, \theta) | x_i, y] \geq \mathbb{E}[u(a', A, \theta) | x_i, y],$$

for all $x_i, y \in \mathbb{R}$ and $a' \in \mathbb{R}$.

2.3.2 Bayes Correlated Equilibrium

The Bayes correlated equilibrium depends on the payoff environment alone, and does not refer to the belief environment. A Bayes correlated equilibrium is defined to be a joint distribution over states and players' actions which has the property that a player who knows only what action he is supposed to play has no incentive to choose a different action. In addition, in this paper, we maintain the assumptions of normality and symmetry. Each player chooses an action $a \in \mathbb{R}$ and there will then be a realized average or population action A . There is a payoff relevant state $\theta \in \Theta$. We are interested in probability distributions $\mu \in \Delta(\mathbb{R} \times \mathbb{R} \times \Theta)$ with the interpretation that μ is the joint distribution of the individual, the average action and the state θ . For any such μ , we write $\mu(\cdot|a)$ for the conditional probability distribution on $\mathbb{R} \times \Theta$ if $(a, A, \theta) \in \mathbb{R} \times \mathbb{R} \times \Theta$.

Definition 2 (Bayes Correlated Equilibrium)

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

$$\mathbb{E}_{\mu(\cdot|a)} [u(a, A, \theta) | a] \geq \mathbb{E}_{\mu(\cdot|a)} [u(a', A, \theta) | a], \tag{11}$$

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

$$\text{marg}_{\Theta} \mu = \psi. \tag{12}$$

The condition (11) states that whenever a player is asked to choose a , he cannot profitably deviate by choosing any different action a' . This is the obedience condition, by direct analogy with the best response condition in the correlated equilibrium for games of complete information in Aumann (1987). The condition (12) states that the marginal of the Bayes correlated equilibrium distribution over the payoff state space Θ has to be consistent with the common prior distribution ψ .

This definition is a special case of a concept introduced in Bergemann and Morris (2011a). The definition here is written for the particular continuum player games studied in this paper, maintaining symmetry and normality, and with players conditioning on their actions only and not on any additional information. The general definition of Bayes correlated equilibrium in Bergemann and Morris (2011a) is defined for general games and allows for players to have some information about the state - giving rise to additional incentive constraints - but allowing for arbitrary additional correlation in their actions. In Section 6 of this paper, we analogously analyze how prior knowledge of the belief environment can refine the set of equilibrium predictions, maintaining our restriction to symmetric normal distributions.

There is a significant literature on alternative definitions of correlated equilibrium in incomplete information environments, with Forges (1993) providing a classic taxonomy. As we discuss in Bergemann and Morris (2011a), our definition of Bayes correlated equilibrium is weaker than the weakest definition in the literature and Forges (1993), intuitively because we allow the mediator to know the payoff state which no individual player knows. While this assumption may seem extreme when defining solution concepts de novo, we will see how it precisely delivers the solution concept that captures the entire set of possible equilibrium outcomes for all possible belief environments.

3 Bayes Nash Equilibrium

We first report, as a benchmark, the standard approach to analyzing this class of games of incomplete information. We 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.

We consider the bivariate normal information structure given by a private and a public signal for each agent i :

$$x_i = \theta + \varepsilon_i, \quad y = \theta + \varepsilon. \quad (13)$$

The respective random variables ε_i and ε are assumed to be normally distributed with zero mean and variance given by σ_x^2 and σ_y^2 . It is at times convenient to express the variance of the random variables in terms of the precision:

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

and 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:

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

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 (5), that the resulting Nash equilibrium strategy is given by:

$$a^*(\theta) \triangleq \frac{k}{1-r} + \frac{s}{1-r}\theta. \quad (14)$$

We refer to the terms in equilibrium strategy (14), $k/(1-r)$ 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 (13). 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 :

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

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, k) and the information structure (τ_x, τ_y) .

Proposition 1 (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:*

$$\alpha_0^* = \frac{k}{1-r} + \frac{s}{1-r} \frac{\mu_\theta \tau_\theta}{\tau - r\tau_x}, \quad \alpha_x^* = s \frac{\tau_x}{\tau - r\tau_x}, \quad \alpha_y^* = \frac{s}{1-r} \frac{\tau_y}{\tau - r\tau_x}. \quad (15)$$

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 (1) and (9), 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:

$$\frac{\alpha_y^*}{\alpha_x^*} = \frac{\tau_y}{\tau_x} \frac{1}{1-r}. \quad (16)$$

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, (14) and (15), 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 2 (Attenuation)

The mean of the individual action in equilibrium is:

$$\mathbb{E}[a] = \alpha_0^* + \alpha_x^* \mu_\theta + \alpha_y^* \mu_\theta = \frac{k + s\mu_\theta}{1 - r}, \quad (17)$$

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. Moreover, with a continuum of agents, by the law of large numbers, the realized average action A always satisfies the equality (17) for every realization of the state θ , or:

$$A = \alpha_0^* + \alpha_x^* \theta + \alpha_y^* \theta = \frac{k + s\theta}{1 - r}, \quad \forall \theta,$$

and hence the realized average action is also independent of the information structure.

Now, if we ask how the joint distribution of the Bayes Nash equilibrium varies with the information structure, then Proposition 2 established that it is sufficient to consider the higher moments of the equilibrium distribution. But given the normality of the equilibrium distribution, it follows that it is sufficient to consider the second moments, that is the variance-covariance matrix. The variance-covariance matrix of the equilibrium joint distribution over individual actions a_i, a_j , and state θ is given by:

$$\Sigma_{a_i, a_j, \theta} = \begin{bmatrix} \sigma_a^2 & \rho_a \sigma_a^2 & \rho_{a\theta} \sigma_a \sigma_\theta \\ \rho_a \sigma_a^2 & \sigma_a^2 & \rho_{a\theta} \sigma_a \sigma_\theta \\ \rho_{a\theta} \sigma_a \sigma_\theta & \rho_{a\theta} \sigma_a \sigma_\theta & \sigma_\theta^2 \end{bmatrix}. \quad (18)$$

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 $\sigma_A^2 = \rho_a \sigma_a^2$.³ 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 1, 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 $(\theta, \varepsilon_i, \varepsilon)$:

$$\Sigma_{a_i, A, \theta} = \begin{bmatrix} \alpha_x^2 \sigma_x^2 + \alpha_y^2 \sigma_y^2 + \sigma_\theta^2 (\alpha_x + \alpha_y)^2 & \alpha_y^2 \sigma_y^2 + \sigma_\theta^2 (\alpha_x + \alpha_y)^2 & \sigma_\theta^2 (\alpha_x + \alpha_y) \\ \alpha_y^2 \sigma_y^2 + \sigma_\theta^2 (\alpha_x + \alpha_y)^2 & \alpha_y^2 \sigma_y^2 + \sigma_\theta^2 (\alpha_x + \alpha_y)^2 & \sigma_\theta^2 (\alpha_x + \alpha_y) \\ \sigma_\theta^2 (\alpha_x + \alpha_y) & \sigma_\theta^2 (\alpha_x + \alpha_y) & \sigma_\theta^2 \end{bmatrix}. \quad (19)$$

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:

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

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:

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

³With 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 (20) are indeed the equilibrium coefficients of the associated Bayes Nash equilibrium strategy.

Proposition 3 (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 $(\rho_a, \rho_{a\theta})$:

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

and

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

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 each individual receives more information about the state θ , an increase in precision always leads to an increase in the correlation with the 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 the actions of the other agents. 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

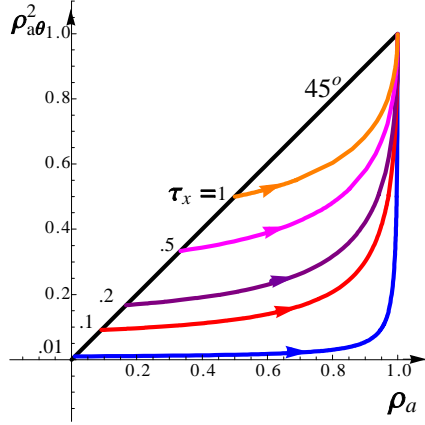


Figure 1: Bayes Nash equilibrium of beauty contest, $r = 1/4$, with varying degree of precision τ_x of private signal.

the change leads to an increase in the correlation across actions. In fact for a given correlation between the individual actions, represented by ρ_a , an increase in the precision of the public signal leads to the elimination of Bayes Nash equilibria with very low *and* with very high correlation between the state of the world and the individual action.

4 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, $a_i - A$. With a continuum of agents, the distribution around the realized average action A can be taken to represent 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.

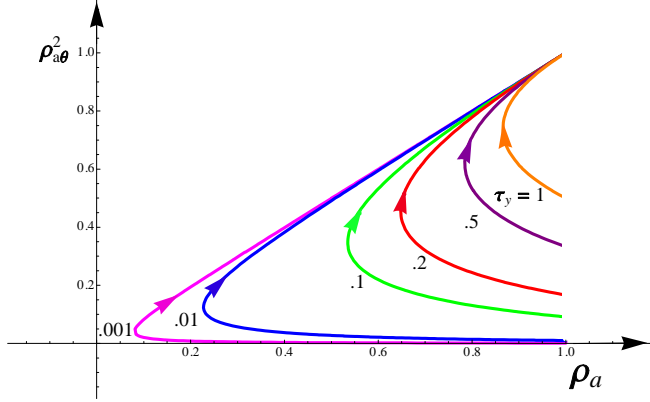


Figure 2: Bayes Nash equilibrium of beauty contest, $r = 1/4$, with varying degree of precision τ_x of public signal.

4.1 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:

$$\begin{pmatrix} \theta \\ A \\ a \end{pmatrix} \sim N \left(\begin{pmatrix} \mu_\theta \\ \mu_A \\ \mu_a \end{pmatrix}, \begin{pmatrix} \sigma_\theta^2 & \rho_{A\theta}\sigma_A\sigma_\theta & \rho_{A\theta}\sigma_A\sigma_\theta \\ \rho_{A\theta}\sigma_A\sigma_\theta & \sigma_A^2 & \sigma_A^2 \\ \rho_{A\theta}\sigma_A\sigma_\theta & \sigma_A^2 & \sigma_A^2 + \sigma_\eta^2 \end{pmatrix} \right). \quad (22)$$

The analysis of the Bayes correlated equilibrium proceeds by deriving restrictions on the joint equilibrium distribution (22). 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 (22). The equilibrium restrictions arise from two sources: (i) the best response conditions of the individual agents:

$$a_i = r\mathbb{E}[A|a_i] + s\mathbb{E}[\theta|a_i] + k, \quad \text{for all } i \text{ and } a_i \in \mathbb{R}, \quad (23)$$

and (ii) the consistency condition of Definition 2, namely that the marginal distribution over θ is equal to the common prior over θ , is satisfied by construction of the joint equilibrium distribution (22). The best

response condition (23) 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 (23) uses the expectation of the individual agent, it is convenient to introduce the following change of variable for the equilibrium random variables. 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:

$$\rho_a\sigma_a^2 = \sigma_A^2. \quad (24)$$

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

$$\rho_a = \frac{\sigma_A^2}{\sigma_A^2 + \sigma_\eta^2}, \quad (25)$$

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

$$\rho_{a\theta} = \rho_{A\theta} \frac{\sigma_A}{\sigma_a}. \quad (26)$$

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:

$$\begin{pmatrix} \theta \\ A \\ a \end{pmatrix} \sim N \left(\begin{pmatrix} \mu_\theta \\ \mu_a \\ \mu_a \end{pmatrix}, \begin{pmatrix} \sigma_\theta^2 & \rho_{a\theta}\sigma_a\sigma_\theta & \rho_{a\theta}\sigma_a\sigma_\theta \\ \rho_{a\theta}\sigma_a\sigma_\theta & \rho_a\sigma_a^2 & \rho_a\sigma_a^2 \\ \rho_{a\theta}\sigma_a\sigma_\theta & \rho_a\sigma_a^2 & \sigma_a^2 \end{pmatrix} \right). \quad (27)$$

With the joint equilibrium distribution described by (27), we now use the best response property (23), to completely characterize the moments of the equilibrium distribution.

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

$$\mathbb{E}[a_i] = k + s\mathbb{E}[\mathbb{E}[\theta | a_i]] + r\mathbb{E}[\mathbb{E}[A | a_i]]. \quad (28)$$

By symmetry, the expected action of each agent is equal to expected average action A , and hence we can use (28) to solve for the mean of the individual action and the average action:

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

It follows that the mean of the individual action and the mean of the average action is uniquely determined by the mean value μ_θ of the state of the world and the parameters (r, s, k) across *all* correlated equilibria.

The complete description of the set of correlated equilibria then rests on the description of the second moments of the multivariate distribution. The characterization of the second moments of the equilibrium distribution again uses the best response property of the individual action, see (23). But, now we use the property of the conditional expectation, rather than the iterated expectation to derive 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 (27), the conditional expectations $\mathbb{E}[\theta | a_i]$ and $\mathbb{E}[A | a_i]$ are linear in a_i and given by

$$\mathbb{E}[\theta | a_i] = \left(1 - \frac{\rho_{a\theta}\sigma_\theta}{\sigma_a} \frac{s}{1-r} \right) \mu_\theta + \frac{\rho_{a\theta}\sigma_\theta}{\sigma_a} \left(a_i - \frac{k}{1-r} \right), \quad (30)$$

and

$$\mathbb{E}[A | a_i] = \frac{k + s\mu_\theta}{1-r} (1 - \rho_a) + \rho_a a_i. \quad (31)$$

The optimality of the best response property can then be expressed, using (30) and (31) 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

$$\sigma_a = \frac{\sigma_\theta s \rho_{a\theta}}{1 - \rho_a r}. \quad (32)$$

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 (24)-(26), where

$$\rho_{a\theta}^2 = \rho_{A\theta}^2 \frac{\sigma_A^2}{\sigma_a^2} = \rho_{A\theta}^2 \rho_a \leq \rho_a. \quad (33)$$

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 (33) 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.

Proposition 4 (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:

$$\mathbb{E}[a_i] = \frac{k}{1-r} + \mu_\theta \frac{s}{1-r}; \quad (34)$$

2. the standard deviation of the individual action is:

$$\sigma_a = \frac{s\rho_{a\theta}}{1-\rho_a r} \sigma_\theta; \quad \text{and} \quad (35)$$

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

$$\rho_{a\theta}^2 \leq \rho_a \quad \text{and} \quad s \cdot \rho_{a\theta} \geq 0. \quad (36)$$

Thus the robust predictions of the linear model are that the mean of the individual action μ_a is pinned down by the parameters of the model and there is a one dimensional restriction on the remaining free endogenous variables $(\sigma_a, \rho_a, \rho_{a\theta})$. Notably, there are (almost) no robust predictions about the correlation coefficients. The sign of $\rho_{a\theta}$ is pinned down by the sign of s , and there is a statistical requirement that $\rho_{a\theta}^2 \leq \rho_a$, but beyond these restrictions, any correlation coefficients are consistent with any values of the parameters of the model and, in particular, with any value of the interaction parameter r . In Section 7, we analyze the issue of robust identification in the model. In particular, we will argue formally in Proposition 12 that any value of the interaction parameter $r \in (-\infty, 1)$ is consistent with any given observed first and second moments of the state $(\mu_\theta, \sigma_\theta)$ and the endogenous variables $(\mu_a, \sigma_a, \rho_a, \rho_{a\theta})$.

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 (24)-(26). 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 (27) is indeed a legitimate variance-covariance matrix, i.e. that it is a nonnegative definite matrix. A necessary condition for the nonnegativity of the matrix is that the determinant of the variance-covariance matrix is nonnegative, or,

$$\sigma_\theta^6 \rho_{a\theta}^4 s^4 (1-\rho_a) \frac{\rho_a - \rho_{a\theta}^2}{(1-\rho_a r)^4} \geq 0 \quad \Rightarrow \quad \rho_{a\theta}^2 \leq \rho_a. \quad (37)$$

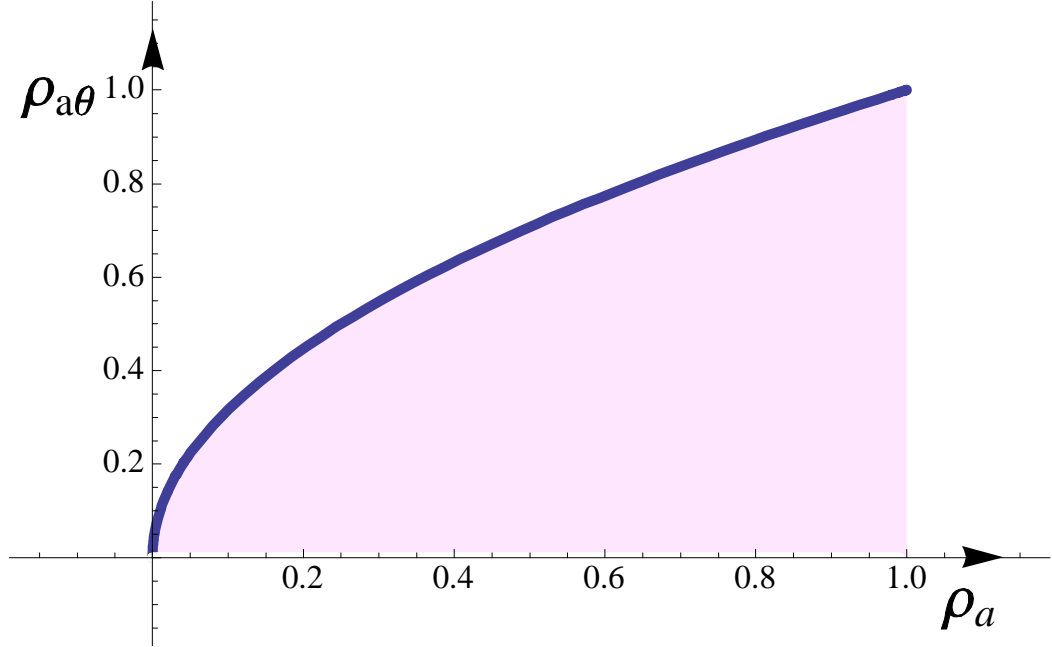


Figure 3: Set of Bayes correlated equilibrium in terms of correlation coefficients ρ_a and $|\rho_{a\theta}|$

In addition, due to the special structure of the present matrix, namely $\sigma_A^2 = \rho_a \sigma_a^2$, the above inequality is also a sufficient condition for the nonnegative definiteness of the matrix.

Later, we extend the analysis from the pure common value environment analyzed here, to an interdependent value environment (in Section 4.4) and to prior restrictions on the private information of the agents (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 variance of the individual action σ_a^2 has to be equal to zero by (32), and hence if the individual actions do not display any correlation with the payoff state θ , then the individual action and hence the average actions must be constant. 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, $\mathbb{E}[\theta]$.

At this point, it is useful 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 2 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 (34) and (35), are unaffected by the number, in particular the finiteness, of the players. The only modification arises with the change

of variable, see (24)-(26), 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 (34), (35), and the correlation coefficients ρ_a and $\rho_{a\theta}$ satisfy the inequalities:

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

It is immediate to verify that the restrictions of the correlation structure in (38) converge towards the one in (36) as $I \rightarrow \infty$. We observe that the restrictions in (38) 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 (32), actually follows the same logic as the condition on the mean of the individual action, given by (29). 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 (32) by using the law of total variance and covariance. More precisely, we could require, using the equality (23), 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 (32). 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.

4.2 Volatility and Dispersion

Proposition 4 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 5 (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

$$\rho_a = |\rho_{a\theta}|^2 = \min \left\{ -\frac{1}{r}, 1 \right\}. \quad (39)$$

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 6 (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 5, the volatility, is increasing in the correlation coefficient ρ_a for a larger range of strategic interaction parameters, including moderate strategic substitutes.

4.3 Matching Bayes Correlated and Nash Equilibria

What is the relation between the set of Bayes correlated equilibria identified in this section, and the set of Bayes Nash equilibria defined over different belief environments in the previous section? For each fixed belief environment and Bayes Nash equilibrium, we identified the implied joint distribution of actions and the payoff state in Proposition 3. It is straightforward to verify directly that these distributions satisfy the conditions characterizing Bayes correlated equilibrium. In particular, this is true for the class of bivariate normal belief environments that we explicitly studied. However, it would also have been true if we had taken any other symmetric normal belief environment (perhaps with more than two dimensional signals), the implied action state distributions would still correspond to some Bayes correlated equilibria.

In this section, we will show the converse result. For every Bayes correlated equilibrium, there is a belief environment such that the Bayes Nash equilibrium for that belief environment gives rise to the Bayes correlated equilibrium action state distributions. It turns out to be sufficient to use the class of bivariate belief environments.

This result illustrates the more general connection between Bayes correlated equilibrium and Bayes Nash equilibrium for all information structures studied in Bergemann and Morris (2011a). In particular, it shows that the connection continues to hold when we impose normality and symmetry assumptions on the analysis. And we get a bonus result - in this special environment - in that we can find a set of belief environments with low dimensionality that are sufficient to establish the connection.

To see the connection concretely, 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 . Moreover, from Proposition 3 and 4, we already know that for every pair of correlation coefficients $(\rho_a, \rho_{a\theta})$ that form a Bayes correlated equilibrium, see (36), we can find a unique information structure (τ_x, τ_y) that generates the same second moments as a Bayes Nash equilibrium. We therefore have the following result.

Corollary 2 (Matching BCE and BNE)

For every interaction structure (r, s, k) , there is a bijection between Bayes correlated and Bayes Nash equilibrium.

Finally we observe that for a given strictly positive and 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 every Bayes Nash equilibrium with finite precision 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 distributions of the payoff state θ . 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.

There is a twist to the matching of Bayes correlated and Bayes Nash equilibrium in the finite player version of the model. In this case, it is possible for the individual actions to become negatively correlated. This cannot arise with the bivariate belief environment. To match Bayes correlated equilibria in the finite player version of the model, we would have to allow for negatively correlated (and thus conditionally dependent) private signals.

4.4 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. Here 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

we are in a canonical model of interdependent values.

The analysis of the Bayes correlated equilibrium can proceed as in Section 4.1. 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:

$$a = r\mathbb{E}[A|a] + s\mathbb{E}[\theta + \nu|a] + k. \quad (40)$$

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 (40), 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:

$$\rho_a - \rho_{a\theta}^2 \geq 0, \quad 1 - \rho_{a\nu}^2 - \rho_a \geq 0. \quad (41)$$

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 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 restrictions regarding the information structure on the equilibrium set. There a natural case of a prior restriction 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.

4.5 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.

5 The Strategic Value of Information

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 observe a conditionally independent noisy signal of the aggregate of the others' information.

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 restricted the information policy to disclose the same, common signal to all the firms. In the present perspective an information policy is a *pair* of information transmission and information disclosure policies. The analysis of the Bayes correlated equilibria now allows us to substantially modify the earlier insights. Interestingly, Proposition 7 establishes that it is with substantial loss in generality to restrict attention to a common and hence perfectly correlated 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 7 (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:

$$\rho_a^* = -\frac{1}{r} < 1 \text{ and } \rho_{a\theta}^* = \sqrt{\rho_a^*} < 1. \tag{42}$$

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 precision τ_x . If all the information were to be publicly shared, then we would 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, i.e. the demand is inelastic. 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 the slope of the demand curve is sufficiently large, and hence demand is elastic, 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 5, these are the correlation coefficients which maximizes the variance of the individual action, i.e. the individual supply decisions. Now if the initial private signals were sufficiently accurate in terms of τ_x , then the induced 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 4, 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 independent across agents, given the state θ . 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 8

1. *If $r \geq -1$, then full disclosure is firm optimal.*
2. *If $r < -1$ and*
 - (a) *if $-\frac{1}{r} > \tau_x / (\tau_x + \tau_\theta)$, 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 $-\frac{1}{r} \leq \tau_x / (\tau_x + \tau_\theta)$, 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 the 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 1. 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 does not reveal all the aggregate information of the agents, as they would otherwise achieve complete correlation in their actions.

The very last result of Proposition 8 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:

$$\rho_a^* = -\frac{1}{r} < \tau_x / (\tau_x + \tau_\theta), \quad (43)$$

says that if the profit maximizing level of correlation ρ_a^* is below the level already induced by the private information conveyed through the signal x with precision 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 2. One 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. Another modification is that with a continuum of firms, the correlation of individual actions cannot be negative. But with finite firms, we noted in Corollary 1 that the correlation can go negative. In this case, when r is sufficiently negative, the optimal information structure requires negative correlation in the firms' actions generated by negatively correlated private signals.

6 Prior Restrictions on the Belief Environment

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 showed 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 contain little information about the state of the world, then the action of each agent does not vary much in the realization of the signal. On the other hand, with precise signals about the state of the world, the best response of each agent does vary substantially with the realized signal and hence displays a larger

variance in equilibrium. We began, in the spirit of the robust analysis, without any assumptions about the private information that the agents may have. But in many circumstances, the analyst may have prior knowledge about some aspects of the private information of the agents, which would allow the analyst to impose *prior restrictions* on the belief environment. In this context, a natural restriction is a lower bound on the precision of the private information of the agents. In other words, the analyst knows that the agent have at least as much information as given by the lower bound, but possibly their information is even more precise. We can then ask how the prediction of the equilibrium behavior can be refined in the presence of these prior restrictions on the private information of the agents. Thus we are interested how the equilibrium set and the equilibrium predictions are affected by *intermediate* restrictions on the belief environment. By intermediate restriction, we mean restrictions which are in between the extremes in which the analyst either has zero information about the belief environment (beyond the common prior about the state of the world) or has perfect information about the belief environment.

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 3 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 (13).

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 a 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 $\Sigma_{\theta, x, y, a, A}$:

$$\begin{bmatrix} \sigma_\theta^2 & \sigma_\theta^2 & \sigma_\theta^2 & \rho_{a\theta}\sigma_a\sigma_\theta & \rho_{a\theta}\sigma_a\sigma_\theta \\ \sigma_\theta^2 & \sigma_\theta^2 + \sigma_x^2 & \sigma_\theta^2 & \sigma_a\sigma_x\rho_{ax} + \sigma_a\sigma_\theta\rho_{a\theta} & \sigma_a\sigma_\theta\rho_{a\theta} \\ \sigma_\theta^2 & \sigma_\theta^2 & \sigma_\theta^2 + \sigma_y^2 & \sigma_a\sigma_y\rho_{ay} + \sigma_a\sigma_\theta\rho_{a\theta} & \sigma_a\sigma_y\rho_{ay} + \sigma_a\sigma_\theta\rho_{a\theta} \\ \rho_{a\theta}\sigma_a\sigma_\theta & \sigma_a\sigma_x\rho_{ax} + \sigma_a\sigma_\theta\rho_{a\theta} & \sigma_a\sigma_y\rho_{ay} + \sigma_a\sigma_\theta\rho_{a\theta} & \sigma_a^2 & \rho_a\sigma_a^2 \\ \rho_{a\theta}\sigma_a\sigma_\theta & \sigma_a\sigma_\theta\rho_{a\theta} & \sigma_a\sigma_y\rho_{ay} + \sigma_a\sigma_\theta\rho_{a\theta} & \rho_a\sigma_a^2 & \rho_a\sigma_a^2 \end{bmatrix}. \quad (44)$$

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:

$$a = r\mathbb{E}[A | a, x, y] + s\mathbb{E}[\theta | a, x, y] + k, \quad \forall a, x, y. \quad (45)$$

In contrast to the analysis of the Bayes correlated equilibrium without prior restrictions, 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 (45) for all realizations of a, x, y . The complete characterization of the set of Bayes correlated equilibria with prior restrictions 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 (45), 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:

$$\sigma_a = \frac{\sigma_\theta s \rho_{a\theta}}{1 - \rho_a r}, \quad \rho_{ax} = \frac{\sigma_\theta ((1 - \rho_a r) - \rho_{a\theta}^2 (1 - r))}{\sigma_x \rho_{a\theta}}, \quad \rho_{ay} = \frac{\sigma_\theta}{\sigma_y \rho_{a\theta}} \left(\frac{1 - \rho_a r}{1 - r} - \rho_{a\theta}^2 \right). \quad (46)$$

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 (46) 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 nonnegative definite matrix. As before in the absence of a priori restrictions, see (37), a necessary but with the structure of the interaction also sufficient condition is that the matrix has a nonnegative determinant or:

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

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

$$1 - \rho_a - \rho_{ax}^2 \geq 0, \quad (47)$$

$$\rho_a - \rho_{a\theta}^2 - \rho_{ay}^2 \geq 0. \quad (48)$$

We encountered the above inequalities before, see Proposition 4.3, 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 (46) with (47), or with (48), 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 (46), 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 restriction τ_x on the private signal as the *private equilibrium set* $C_x(\tau_x, r)$:

$$C_x(\tau_x, r) \triangleq \{(\rho_a, \rho_{a\theta}) \in [0, 1] \times [-1, 1] \mid (\rho_a, \rho_{a\theta}, \rho_{ax}) \text{ satisfy (36), (46), (47)}\}.$$

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

$$C_y(\tau_y, r) \triangleq \{(\rho_a, \rho_{a\theta}) \in [0, 1] \times [-1, 1] \mid (\rho_a, \rho_{a\theta}, \rho_{ay}) \text{ satisfy (36), (46), (48)}\}.$$

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

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

The above description of the Bayes correlated equilibria in terms of the correlation coefficients suggests that we might be interested in the comparative statics of the equilibrium set with respect to the prior restrictions $\tau = (\tau_x, \tau_y)$ and with respect to the nature of the strategic interaction r . The first exercise in comparative statics is central for the robustness of equilibrium predictions, whereas the second is central for robust identification of the structural parameters of the game to which we turn in the next section.

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 (46) and (47) as an equality, given a lower bound on the precision τ_x of the private information. Similarly, each backward bending curve traces out the set of correlation coefficients $(\rho_a, \rho_{a\theta})$ which solve (46) and (48) 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.

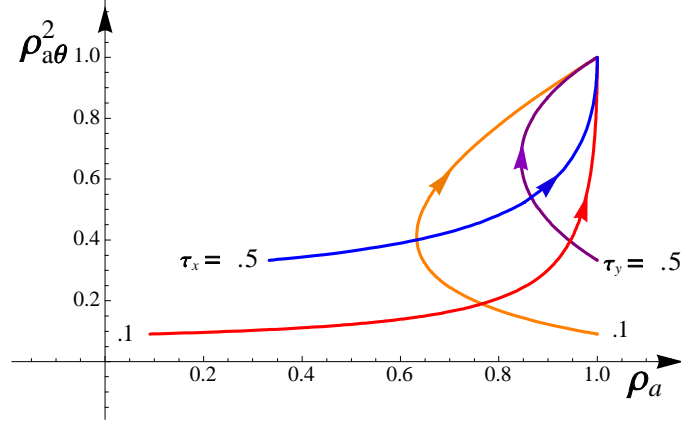


Figure 4: Set of BCE with given public and private information

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 9 (Prior Restrictions and Equilibrium Set)

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

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

Thus, as the precision of the prior restriction 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 restriction 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 Proposition 9 we described the set of possible equilibrium coefficients $(\rho_a, \rho_{a\theta})$ as a function of the prior restrictions τ and the interaction parameter r . Now, suppose we were observe the equilibrium outcomes, and in particular the equilibrium correlation coefficients $(\rho_a, \rho_{a\theta})$, we may then ask which values of the interaction parameter r are consistent with the observed data. To this end we need to know the set of possible equilibrium correlation coefficients $(\rho_a, \rho_{a\theta})$ varies with the interaction parameter r of the game. To this end we need to know the set of possible equilibrium correlation coefficients $(\rho_a, \rho_{a\theta})$ varies with the interaction parameter r of the game.

Proposition 10 (Interaction and Equilibrium Set)

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

1. $C_x(\tau_x, r)$ is increasing in r ;
2. $C_y(\tau_y, r)$ is decreasing in r .

The comparative static results in the interaction parameter r are straightforward. The information in the private signal x leads each agent to choose an action which is less correlated with the average action than the same information contained in the public signal y . Now, as the interaction in the game tends towards strategic substitutability, each agent tends to rely more heavily on the private signal relative to the public signal. Thus, for every level of correlation with the state θ , expressed in terms of $\rho_{a\theta}$, there will be less correlation across actions, expressed in terms of ρ_a . The behavior of the equilibrium set $C_x(\tau_x, r)$ with respect to the interaction parameter r is illustrated in Figure 5.

Conversely, the restrictions imposed by the public information, represented by the set $C_y(\tau_y, r)$ become weaker as the game is moving from strategic complements to strategic substitutes. After all, the public information correlates the agent's action because they rely on the same information. If we decrease the propensity to coordinate, and hence correlate, then all equilibria will display less correlation across actions, for a given correlation with respect to the state θ . The behavior of the equilibrium set $C_y(\tau_y, r)$ with respect to the interaction parameter r is illustrated in Figure 6.

We thus find that the comparative static results with respect to the strategic interaction are pointing in the opposite direction for the equilibrium sets $C_x(\tau_x, r)$ and $C_y(\tau_y, r)$, respectively. In consequence, the equilibrium set $C(\tau, r)$ formed by the intersection of the private and public equilibrium sets, $C(\tau, r) = C_x(\tau_x, r) \cap C_y(\tau_y, r)$, does not display a monotone behavior in r in terms of set inclusion.

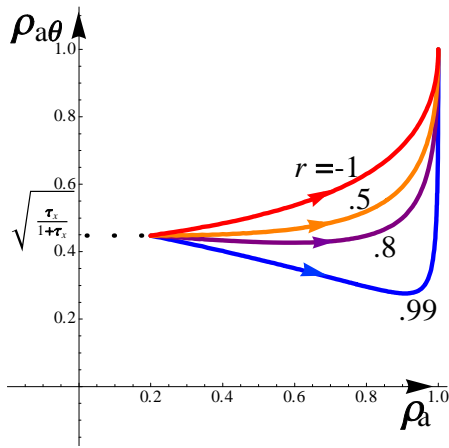


Figure 5: Bayes correlated equilibrium set with precision τ_x of prior private information

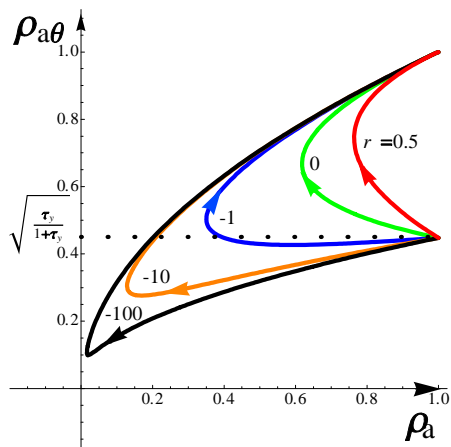


Figure 6: Bayes correlated equilibrium set with precision τ_y of prior public information

Finally, we note that we extended the analysis of the Bayes correlated equilibrium from an environment with pure common values to an environment with interdependent values in Section 4.4. Similarly, we could extend the present analysis of the impact of prior restrictions, pursued here in some detail for the environment with pure common values to the one with interdependent values.

7 Robust Partial 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, k) , by the observed variables? We are particularly interested in how the identification of the structural parameters is influenced by the solution concept, and hence the specification of the private information of the agents as known to the analyst.

Now, identification 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 .⁴ In other words, the econometrician learns the first and second *moments* of the joint equilibrium distribution over actions and state: $m \triangleq (\mu_a, \sigma_a, \rho_a, \sigma_\theta, \rho_{a\theta})$. We begin the identification analysis under the hypothesis of Bayes Nash equilibrium and a given information structure $\tau = (\tau_x, \tau_y)$ of the agents.

For a given information structure and observed moments of the Bayes Nash equilibrium distribution 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 (20). 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 (16):

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

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

⁴In 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.

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

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:

$$\mu_a = \frac{k + \mu_\theta s}{1 - r}, \quad \sigma_a = \frac{\sigma_\theta s \rho_{a\theta}}{1 - \rho_a r}. \quad (49)$$

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

$$\frac{k}{1 - r} = \mu_a - \frac{\sigma_a \mu_\theta (1 - \rho_a r)}{\sigma_\theta (1 - r) \rho_{a\theta}}, \quad \frac{s}{1 - r} = \frac{\sigma_a}{\rho_{a\theta} \sigma_\theta} \frac{1 - \rho_a r}{1 - r}. \quad (50)$$

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,

$$\frac{k}{1 - r} \in \begin{cases} \left(-\infty, \mu_a - \frac{\mu_\theta \rho_a \sigma_a}{\rho_{a\theta} \sigma_\theta} \right) & \text{if } \frac{\mu_\theta}{\rho_{a\theta}} > 0; \\ \left(\mu_a - \frac{\mu_\theta \rho_a \sigma_a}{\rho_{a\theta} \sigma_\theta}, \infty \right) & \text{if } \frac{\mu_\theta}{\rho_{a\theta}} < 0; \end{cases} \quad (51)$$

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

$$\frac{s}{1-r} \in \begin{cases} \left(\frac{\rho_a \sigma_a}{\rho_{a\theta} \sigma_\theta}, \infty \right) & \text{if } \rho_{a\theta} > 0; \\ \left(-\infty, \frac{\rho_a \sigma_a}{\rho_{a\theta} \sigma_\theta} \right) & \text{if } \rho_{a\theta} < 0. \end{cases} \quad (52)$$

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

Proposition 12 (Partial 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 (52), 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 restrictions can improve the identification of the structural parameters, just as prior restrictions could improve the equilibrium prediction. In Section 6, we showed that the prior restriction (τ_x, τ_y) on the belief environment systematically restricted the equilibrium predictions. Now, as we consider the identification of the structural parameters, we might use the knowledge of the prior restrictions (τ_x, τ_y) together with the data to identify the set of structural parameters consistent with the data and the prior restrictions. The content of the subsequent propositions indeed establishes that the set-identification improves with the prior restrictions. In particular, we ask how the identification of the sign of r and the set identification of the equilibrium slope $\frac{s}{1-r}$ is affected by the nature of the prior restrictions $\tau = (\tau_x, \tau_y)$. We denote the lower and upper bound of the identified set for the *equilibrium slope* $\frac{s}{1-r}$ by $\underline{s}(\tau, m)$ and $\bar{s}(\tau, m)$, respectively. The bounds depend naturally on the prior restrictions and the observed data. Similarly, we denote the lower and upper bound of the identified set for the strategic interaction r by $\underline{r}(\tau, m)$ and $\bar{r}(\tau, m)$, respectively.

We observed earlier that the identification of the equilibrium slope in the Bayes correlated equilibrium, see (50):

$$-\frac{s}{1-r} = \frac{\sigma_a}{\rho_{a\theta}\sigma_\theta} \frac{1 - \rho_a r}{1 - r}, \quad (53)$$

relies on a ratio of the observed data $\sigma_a/\rho_{a\theta}\sigma_\theta$ and a ratio

$$\frac{1 - \rho_a r}{1 - r}, \quad (54)$$

which involves the unknown structural parameters and the data. For the set identification of the equilibrium slope, $s/(1-r) \in [\underline{s}(\tau, m), \bar{s}(\tau, m)]$, we have to ask what is the range of the ratio $(1 - \rho_a r)/(1 - r)$ consistent with data. For $\rho_a < 1$, the value of this ratio is increasing in r , and hence the greatest possible value of r , consistent with the data, provides the upper bound for the above ratio. Now, as a consequence of Proposition 10 the upper bound on r is given by the restrictions of the public equilibrium set, and similarly the lower bound is given by the restrictions of the private equilibrium set. To wit, as we increase the interaction parameter r , the set of correlation coefficients $(\rho_a, \rho_{a\theta})$ consistent with a given prior restriction τ is shrinking in the public equilibrium set, and hence any given data point $(\rho_a, \rho_{a\theta})$ is eventually eliminated. becomes inconsistent. We therefore improve the identification with an increase in the precision of the prior restrictions.

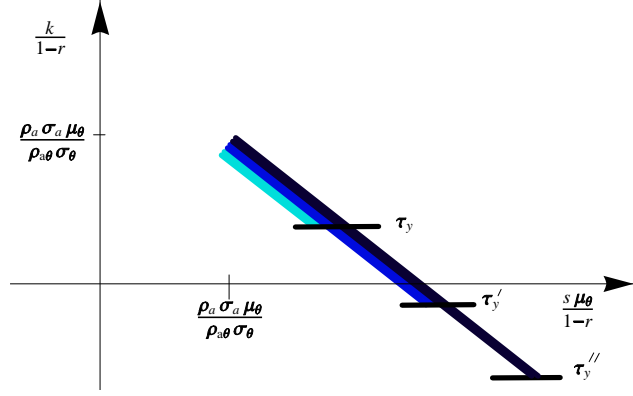


Figure 7: Set Identification and Prior Restrictions: $\tau_y > \tau_y' > \tau_y''$

Proposition 13 (Slope Identification and Prior Restrictions)

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

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

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

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

Proposition 14 (Sign Identification and Prior Restrictions)

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

$$0 < \underline{r}((\tau_x, \tau_y), m) < \bar{r}((\tau_x, \tau_y), m),$$

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

$$\underline{r}((\tau_x, \tau_y), m) < \bar{r}((\tau_x, \tau_y), m) < 0.$$

Thus, as the lower bound on the private and the public signal increases, eventually the sign of the strategic interaction can be identified as well. Hence, sufficiently precise restrictions on the belief environment reestablish the possibility of identification in case of the strategic interaction as well.

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 (2012). 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 (2012), 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 (2012) 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.

8 Conclusion

It was the objective of this paper to derive robust equilibrium predictions for a large class of games. 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 implications 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 partial identification. The current perspective, namely to analyze the set of Bayes 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 Bergemann, Brooks, and Morris (2012) 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 restrictions of the policy maker about the belief environment of the agents.

9 Appendix

Proof of Proposition 1. 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). The uniqueness of the Bayes Nash equilibrium for the present general environment is established in Ui and Yoshizawa (2012). ■

Proof of Proposition 3. 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:

$$\rho_{a\theta} = \pm \frac{\sigma_\theta (\alpha_x + \alpha_y)}{\sqrt{\alpha_x^2 \sigma_x^2 + \alpha_y^2 \sigma_y^2 + \sigma_\theta^2 (\alpha_x + \alpha_y)^2}}, \quad (55)$$

and

$$\rho_a = \frac{\alpha_y^2 \sigma_y^2 + \sigma_\theta^2 (\alpha_x + \alpha_y)^2}{\alpha_x^2 \sigma_x^2 + \alpha_y^2 \sigma_y^2 + \sigma_\theta^2 (\alpha_x + \alpha_y)^2}. \quad (56)$$

It now follows immediately from (55) - (56), and the formulae of α_x^* and α_y^* , see (15), 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 5. The variance σ_a^2 is given by (32), 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 6. (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 latter 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 (35), 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 4, 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 9. We form the conditional expectation using (44) and the equilibrium conditions for the Bayes correlated equilibrium are then given by (45) and the solution to these equations is given by (46).

(1.) The equilibrium set is described as the set which satisfies the inequalities (47) and (48), where the correlation coefficients ρ_{ax}^2 and ρ_{ay}^2 appear separately. By determination of (46), 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 (47) and (48) 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 (48) 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 7. 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

$$\frac{\sigma_\theta^2 s^2 \rho_{a\theta}^2}{2(1 - r\rho_a)^2} \quad (57)$$

The remaining restriction of the Bayes correlated equilibrium, see Proposition 4, is that $\rho_{a\theta}^2 \leq \rho_a$, and hence (57) 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 8. By Proposition 7, 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 1, 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

(19), 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 10. The comparative static results follow directly from the description of the correlation coefficients ρ_{ax} and ρ_{ay} given by (46). These correlation coefficients are a function of r only. We insert their solutions into the inequalities (47) and (48) and solve for the relation between ρ_a and $\rho_{a\theta}$ as we restrict the the inequalities (47) and (48) to be equalities. ■

Proof of Proposition 11. (1.) Given the knowledge of σ_θ^2 , σ_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 (18), say:

$$\alpha_x^2 = \frac{\sigma_a^2 - \sigma_A^2}{\sigma_x^2}, \quad \alpha_y^2 = \frac{\sigma_A^2 (1 - \rho_{A\theta}^2)}{\sigma_y^2}. \quad (58)$$

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 (15).

(2.) We have from the description of the Bayes Nash equilibrium in Proposition 1 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 $k/(1-r)$ 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 (15), we have the values of μ_a, α_x and α_y . This allows us to solve for r, s, k as a function of $\mu_a, \alpha_x, \alpha_y$:

$$\begin{aligned} k &= -\frac{\alpha_x^2 \sigma_x^2 \mu_\theta - \alpha_x \mu_a \sigma_x^2 - \alpha_x \alpha_y \sigma_y^2 \mu_\theta + \alpha_x \alpha_y \tau \sigma_x^2 \sigma_y^2 \mu_\theta}{\alpha_y \sigma_y^2}, \\ r &= \frac{\alpha_y \sigma_y^2 - \alpha_x \sigma_x^2}{\alpha_y \sigma_y^2}, \\ s &= \frac{\alpha_x^2 \sigma_x^2 - \alpha_x \alpha_y \sigma_y^2 + \alpha_x \alpha_y \tau \sigma_x^2 \sigma_y^2}{\alpha_y \sigma_y^2}. \end{aligned} \quad (59)$$

If we form the ratios $k/(1-r)$ and $s/(1-r)$ with the expressions on the rhs of (59), then we obtain expressions which do only depend on the observable data, and are hence point identified, and in particular

$$\frac{k}{1-r} = -\frac{-\mu_a \sigma_x^2 + \alpha_x \sigma_x^2 \mu_\theta - \alpha_y \sigma_y^2 \mu_\theta + \tau \sigma_x^2 \alpha_y \sigma_y^2 \mu_\theta}{\sigma_x^2}, \quad (60)$$

and

$$\frac{s}{1-r} = \frac{\alpha_x \sigma_x^2 - \alpha_y \sigma_y^2 + \tau \alpha_y \sigma_x^2 \sigma_y^2}{\sigma_x^2}, \quad (61)$$

which completes the proof of identification. We observe that, using (58), we could express the ratios (60) and (61) 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 (49). 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 (49) in the presence of the linear return term k , 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 (49). We thus have two equations to identify the three unknown structural parameters (r, s, k) . We can solve for (s, k) in terms of the remaining unknown r to obtain:

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

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

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

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 (51) and (52) which describe the respective sets into which each ratio can be identified. ■

Proof of Proposition 13. We know from Proposition 12 that the interaction ratio is a function of the observed data and the unobserved interaction parameters:

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

The prior restrictions on the private and public information restricts the possible values of r , and hence the values that the above interaction ratio can attain.

We begin the argument with the public equilibrium set which will provide an upper bound on the ratio $(1 - \rho_a r) / (1 - r)$ which appears in the correlation coefficient ρ_{ay} as described in (46). The value of the ratio is maximized when the inequality constraint (48) of the public equilibrium set holds as an equality, and thus

$$\frac{1 - \rho_a r}{1 - r} = \rho_{a\theta}^2 + \frac{\sigma_y}{\sigma_\theta} |\rho_{a\theta}| \sqrt{\rho_a - \rho_{a\theta}^2}, \quad (62)$$

and hence

$$\bar{s}(\tau, m) = \frac{\sigma_a}{\rho_{a\theta}\sigma_\theta} \left(\rho_{a\theta}^2 + \frac{1}{\sqrt{\tau_y}\sigma_\theta} |\rho_{a\theta}| \sqrt{\rho_a - \rho_{a\theta}^2} \right).$$

It follows that if the precision τ_y increases, then the largest value of the above ratio decreases and as τ_y increases to infinity, the value of the ratio tends to $\rho_{a\theta}^2$, which implies that

$$\lim_{\tau_y \uparrow \infty} \bar{s}(\tau, m) = \frac{\sigma_a |\rho_{a\theta}|}{\sigma_\theta}.$$

Now, consider the private equilibrium set. We ask what feasible pair in the private equilibrium set minimizes the ratio (54). The minimal ratio is achieved by an r which solves the inequality (47) as an equality. We find that the minimal value of the ratio (54) is given by:

$$\frac{1 - \rho_a r}{1 - r} = \frac{(\rho_a - \rho_{a\theta}^2) \sigma_\theta + \rho_a \left(|\rho_{a\theta}| \sqrt{\frac{1 - \rho_a}{\tau_x}} - (1 - \rho_{a\theta}^2) \sigma_\theta \right)}{(\rho_a - \rho_{a\theta}^2) \sigma_\theta + \left(|\rho_{a\theta}| \sqrt{\frac{1 - \rho_a}{\tau_x}} - (1 - \rho_{a\theta}^2) \sigma_\theta \right)}, \quad (63)$$

and hence

$$\underline{s}(\tau, m) = \frac{\sigma_a}{|\rho_{a\theta}| \sigma_\theta} \frac{(\rho_a - \rho_{a\theta}^2) \sigma_\theta + \rho_a \left(|\rho_{a\theta}| \sqrt{\frac{1 - \rho_a}{\tau_x}} - (1 - \rho_{a\theta}^2) \sigma_\theta \right)}{(\rho_a - \rho_{a\theta}^2) \sigma_\theta + \left(|\rho_{a\theta}| \sqrt{\frac{1 - \rho_a}{\tau_x}} - (1 - \rho_{a\theta}^2) \sigma_\theta \right)}.$$

It is immediate to verify that $\underline{s}(\tau, m) \leq \bar{s}(\tau, m)$ for all τ and m . It follows from the determination of $\underline{s}(\tau, m)$ that as a function of τ_x , $\underline{s}(\tau, m)$ is increasing in τ_x . Moreover as τ_x increases, the value of the ratio tends to $\rho_{a\theta}^2$ and hence

$$\lim_{\tau_x \uparrow \infty} \underline{s}(\tau, m) = \frac{\sigma_a |\rho_{a\theta}|}{\sigma_\theta},$$

which concludes the proof. ■

Proof of Proposition 14. For any given ρ_a , with $0 \leq \rho_a < 1$, the ratio

$$\frac{1 - \rho_a r}{1 - r}$$

is larger than 1 if and only $r > 0$. Thus we can identify the sign of r if we can establish that the ratio on the rhs of (63), which determined the lower bound on the equilibrium slope, is larger than 1. Now, clearly if $\rho_a < 1$, and if for some $\bar{\sigma}_x$:

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

we have

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

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

Similarly, if for given data the expression on the rhs of (62) is smaller than 1 for some $\bar{\sigma}_y$:

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

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

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.
- BERGEMANN, D., B. BROOKS, AND S. MORRIS (2012): “Extremal Information Structures in First Price Auctions,” 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.
- (2012): *Robust Mechanism Design*. World Scientific Publishing, Singapore.
- BERGEMANN, D., AND M. PESENDORFER (2007): “Information Structures in Optimal Auctions,” *Journal of Economic Theory*, 137, 580–609.
- BRILLINGER, D. (1969): “The Calculation of Cumulants Via Conditioning,” *Annals of the Institute of Statistical Mathematics*, 21, 215–218.

- CLARKE, R. (1983): “Collusion and the Incentives for Information Sharing,” *Bell Journal of Economics*, 14, 383–394.
- 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.
- 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. SONNENSCHNEIN (1982): “Fulfilled Expectations Cournot Duopoly with Information Acquisition and Release,” *Bell Journal of Economics*, 13, 214–218.
- PAULA, A. D., AND X. TANG (2012): “Inference of Signs of Interaction Effects in Simultaneous Games with Incomplete Information,” *Econometrica*, 80, 143–172.
- 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.
- RUBINSTEIN, A. (1989): “The Electronic Mail Game: Strategic Behavior under ‘Almost Common Knowledge’,” *American Economic Review*, 79, 385–391.
- SWEETING, A. (2009): “The Strategic Timing of Radio Commercials: An Empirical Analysis Using Multiple Equilibria,” *RAND Journal of Economics*, 40, 710–742.

- UI, T., AND Y. YOSHIZAWA (2012): “On the Uniqueness of Equilibrium in Quadratic Bayesian Games with a Continuum of Players,” Discussion paper, Yokohama National University.
- 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.
- WEINSTEIN, J., AND M. YILDIZ (2007): “A Structure Theorem for Rationalizability with Applications to Robust Predictions of Refinements,” *Econometrica*, 75, 365–400.