

Robust Predictions
in
Games with Incomplete Information

Dirk Bergemann and Stephen Morris

Northwestern University
March, 30th, 2011

- game theoretic predictions are very sensitive to "higher order beliefs" or (equivalently) information structure
- higher order beliefs are rarely observed
- what predictions can we make and analysis can we do if we do not observe higher order beliefs?

Agenda: Robust Predictions

- fix "payoff relevant environment"
 - = action sets, payoff-relevant variables ("states"), payoff functions, distribution over states
 - = incomplete information game without higher order beliefs about states
- analyze what could happen for all possible higher order beliefs (maintaining common prior assumption and equilibrium assumptions)
- make set valued predictions about joint distribution of actions and states

Agenda: Robust Identification

- having identified mapping from "payoff relevant environment" to action-state distributions, we can analyze its inverse:
- assume payoff relevant environment is observed by the econometrician
- given knowledge of the action-state distribution distribution, or some moments of it, what can be deduced about the payoff relevant environment?
- partial identification / set identification

Agenda: Robust Comparative Statics / Policy Analysis

- what can say about the impact of changes in parameters of the payoff relevant environment (for example, policy choices) for the set of possible outcomes?
- robust information policy
- robust taxation policy

Partial Information about Information Structure

- perhaps you don't observe all higher order beliefs but you are sure of some aspects of the information structure:
 - you are sure that bidders know their private values of an object in an auction, but you have no idea what their beliefs and higher order beliefs about others' private values are...
 - you are sure that oligopolists know their own costs, but you have no idea what beliefs and higher order beliefs about demand and others' private values.
- what can you say then?

- ① General Approach
- ② Illustration with Continuum Player, Symmetric, Linear Best Response, Normal Distribution Games

① General Approach

- Set valued prediction is set of "Bayes Correlated Equilibria"
- Partial information monotonically reduces the set of "Bayes Correlated Equilibria"

② Illustration with Continuum Player, Symmetric, Linear Best Response, Normal Distribution Games

- These sets are tractable and intuitive
- Cannot distinguish strategic substitutes and complements

- players $i = 1, \dots, I$
- (payoff relevant) states Θ
- actions $(A_i)_{i=1}^I$
- utility functions $(u_i)_{i=1}^I$, each

$$u_i : A \times \Theta \rightarrow \mathbb{R}$$

- common prior state distribution $\psi \in \Delta(\Theta)$
- "basic game", "belief-free game"

$$G = \left((A_i, u_i)_{i=1}^I, \psi \right)$$

- signals (types) $(T_i)_{i=1}^I$
- signal distribution

$$\pi : \Theta \rightarrow \Delta(T_1 \times T_2 \times \dots \times T_I)$$

- "higher order beliefs", "type space," "signal space"

$$S = \left((T_i)_{i=1}^I, \pi \right)$$

- a standard Bayesian game is described by (u, ψ, \mathcal{T})
- a behavior strategy of player i is defined by:

$$\sigma_i : T_i \rightarrow \Delta(\mathcal{A}_i)$$

Definition (Bayes Nash Equilibrium (BNE))

A strategy profile σ is a Bayes Nash equilibrium of (u, ψ, \mathcal{T}) if

$$\begin{aligned} & \sum_{t_{-i}, \theta} u_i((\sigma_i(t_i), \sigma_{-i}(t_i)), \theta) \psi(\theta) \pi[t_i, t_{-i}](\theta) \\ & \geq \sum_{t_{-i}, \theta} u_i((a_i, \sigma_{-i}(t_{-i})), \theta) \psi(\theta) \pi[t_i, t_{-i}](\theta). \end{aligned}$$

for each i , t_i and a_i .

Bayes Nash Equilibrium Distribution

- given a Bayesian game (u, ψ, \mathcal{T}) , a BNE σ generates a joint probability distribution μ_σ over outcomes and states $\mathcal{A} \times \Theta$,

$$\mu_\sigma(a, \theta) = \psi(\theta) \sum_t \pi[t](\theta) \left(\prod_{i=1}^I \sigma_i(a_i | t_i) \right)$$

- equilibrium distribution $\mu_\sigma(a, \theta)$ is specified without reference to information structure \mathcal{T} which gives rise to $\mu_\sigma(a, \theta)$

Definition (Bayes Nash Equilibrium Distribution)

A probability distribution $\mu \in \Delta(\mathcal{A} \times \Theta)$ is a Bayes Nash equilibrium distribution (over action and states) of (u, ψ, \mathcal{T}) if there exists a BNE σ of (u, ψ, \mathcal{T}) such that

$$\mu = \mu_\sigma.$$

- recall the original equilibrium conditions on (u, ψ, \mathcal{T}) :

$$\begin{aligned} & \sum_{t_{-i}, \theta} u_i ((\sigma_i(t_i), \sigma_{-i}(t_{-i})), \theta) \psi(\theta) \pi[t_i, t_{-i}](\theta) \\ & \geq \sum_{t_{-i}, \theta} u_i ((a_i, \sigma_{-i}(t_{-i})), \theta) \psi(\theta) \pi[t_i, t_{-i}](\theta). \end{aligned}$$

- with the equilibrium distribution

$$\mu_\sigma(a, \theta) = \psi(\theta) \sum_t \pi[t](\theta) \left(\prod_{i=1}^I \sigma_i(a_i | t_i) \right)$$

- an implication of BNE of (u, ψ, \mathcal{T}) : for all $a_i \in \text{supp } \mu_\sigma(a, \theta)$:

$$\sum_{a_{-i}, \theta} u_i((a_i, a_{-i}), \theta) \mu_\sigma(a, \theta) \geq \sum_{a_{-i}, \theta} u_i((a'_i, a_{-i}), \theta) \mu_\sigma(a, \theta);$$

Definition (Bayes Correlated Equilibrium (BCE))

An action state distribution $\mu \in \Delta(A \times \Theta)$ is a Bayes Correlated Equilibrium (BCE) of G if is *obedient*, i.e., for each i , a_i and a'_i ,

$$\sum_{a_{-i} \in A_{-i}, \theta \in \Theta} (u_i((a_i, a_{-i}), \theta) - u_i((a'_i, a_{-i}), \theta)) \mu((a_i, a_{-i}), \theta) \geq 0$$

and *consistent*, i.e., for each θ

$$\sum_{a \in A} \mu(a, \theta) = \psi(\theta).$$

- no restrictions (or lower bound) on the private information of the agents; in particular, zero private information is possible
- BCE is defined in terms of the payoff environment and without reference to type spaces
- BCE is defined on “small” payoff space and characterized as

- now given (u, ψ) , what is the set of equilibrium distributions μ across all possible information structures \mathcal{T}

Theorem (Equivalence)

A probability distribution $\mu \in \Delta(\mathcal{A} \times \Theta)$ is a Bayes correlated equilibrium of (u, ψ) if and only if it is a Bayes Nash Equilibrium distribution of (u, ψ, \mathcal{T}) for some information system \mathcal{T} .

- *BCE \Rightarrow BNE* uses the richness of the possible information structure to complete the equivalence result
- Aumann (1987) established the above characterization result for complete information games

- Forges (1993): "Five Legitimate Definitions of Correlated Equilibrium in Games with Incomplete Information"; Forges (2006) gives #6
- our definition is "illegitimate" because it fails "join feasibility"

Definition

Action type state distribution ν is join feasible for (G, S) if there exists $f : T \rightarrow \Delta(A)$ such that

$$\nu(a, t, \theta) = \psi(\theta) \pi(t|\theta) f(a|t)$$

for each a, t, θ .

- BCE fails join feasibility, Forges' weakest definition (Bayesian solution) is BCE satisfying join feasibility
- in companion paper, "Correlated Equilibrium in Games with Incomplete Information" we relate it to earlier definitions and establish comparative results with respect to information environments

Trivial One Player Example

- $I = 1$
- $\Theta = \{\theta, \theta'\}$
- $\psi(\theta) = \psi(\theta') = \frac{1}{2}$
- Payoffs u_1

	θ	θ'
a_1	2	-1
a'_1	0	0

- unique Bayesian solution: $\mu(a_1, \theta) = \mu(a_1, \theta') = \frac{1}{2}$
- a BCE: $\mu(a_1, \theta) = \mu(a'_1, \theta') = \frac{1}{2}$

Payoff Environment: Quadratic Payoffs

- utility of each agent i is given by quadratic payoff function:
- determined by individual action $a_i \in \mathbb{R}$, state of the world $\theta \in \mathbb{R}$, and average action $A \in \mathbb{R}$:

$$A = \int_0^1 a_i di$$

and thus $u_i(a_i, A, \theta) =$

$$\begin{pmatrix} \lambda_a \\ \lambda_A \\ \lambda_\theta \end{pmatrix}' \begin{pmatrix} a_i \\ A \\ \theta \end{pmatrix} + \begin{pmatrix} a_i \\ A \\ \theta \end{pmatrix}' \begin{pmatrix} \gamma_a & \gamma_{aA} & \gamma_{a\theta} \\ \gamma_{aA} & \gamma_A & \gamma_{A\theta} \\ \gamma_{a\theta} & \gamma_{A\theta} & \gamma_\theta \end{pmatrix} \begin{pmatrix} a_i \\ A \\ \theta \end{pmatrix}$$

- game is completely described by linear returns λ and interaction matrix $\Gamma = \{\gamma_{ij}\}$

Payoff Environment: Normal Payoffs

- the state of the world θ is normally distributed

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

with mean $\mu_\theta \in \mathbb{R}$ and variance $\sigma_\theta^2 \in \mathbb{R}_+$

- the distribution of the state of the world, θ , is commonly known common prior

- given the interaction matrix Γ , complete information game is a potential game (Monderer and Shapley (1996)):

$$\Gamma = \begin{pmatrix} \gamma_a & \gamma_{aA} & \gamma_{a\theta} \\ \gamma_{aA} & \gamma_A & \gamma_{A\theta} \\ \gamma_{a\theta} & \gamma_{A\theta} & \gamma_\theta \end{pmatrix}$$

- diagonal entries: $\gamma_a, \gamma_A, \gamma_\theta$ describe “own effects”
- off-diagonal entries: $\gamma_{a\theta}, \gamma_{A\theta}, \gamma_{aA}$ “interaction effects”
- fundamentals matter, “return shocks”:

$$\gamma_{a\theta} \neq 0;$$

- strategic complements and strategic substitutes:

$$\gamma_{aA} > 0 \quad \text{vs.} \quad \gamma_{aA} < 0$$

- concavity at the individual level (well-defined best response):

$$\gamma_a < 0$$

- concavity at the aggregate level (existence of an interior equilibrium)

$$\gamma_a + \gamma_{aA} < 0$$

- concave payoffs imply that the complete information game has unique Nash **and** unique correlated equilibrium (Neyman (1997))

Example 1: Beauty Contest

- continuum of agents: $i \in [0, 1]$
- action (= message): $a \in \mathbb{R}$
- state of the world: $\theta \in \mathbb{R}$
- payoff function

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

with $r \in (0, 1)$

- see Morris and Shin (2002), Angeletos and Pavan (2007)

Example 2: Competitive Market

- action (= quantity): $a_i \in \mathbb{R}$
- cost of production $c(a_i) = \frac{1}{2}\gamma_a (a_i)^2$
- state of the world (= demand intercept): $\theta \in \mathbb{R}$
- inverse demand (= price):

$$p(A) = \gamma_{a\theta}\theta - \gamma_{aA}A$$

where A is average supply:

$$A = \int_0^1 a_i di$$

- see Guesnerie (1992) and Vives (2008)

fix an information system:

- 1 every agent i observes a public signal y about θ :

$$y \sim N(\theta, \sigma_y^2)$$

- 2 every agent i observes a private signal x_i about θ :

$$x_i \sim N(\theta, \sigma_x^2)$$

- the best response of each agent is:

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

- suppose the equilibrium strategy is given by a linear function:

$$a(x, y) = \alpha_0 + \alpha_x x + \alpha_y y,$$

- denote the sum of the precisions:

$$\sigma^{-2} = \sigma_{\theta}^{-2} + \sigma_x^{-2} + \sigma_y^{-2}$$

Theorem

The unique Bayesian Nash equilibrium (given the bivariate information structure) is a linear equilibrium,

$$\alpha_0^* + \alpha_x^* x + \alpha_y^* y,$$

with

$$\alpha_x^* = -\frac{\gamma_{a\theta}\sigma_x^{-2}}{\gamma_{Aa}\sigma_x^{-2} + \gamma_a\sigma^{-2}},$$

and

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

Joint Action State Distribution

- there is an implied joint distribution of (a, A, θ)

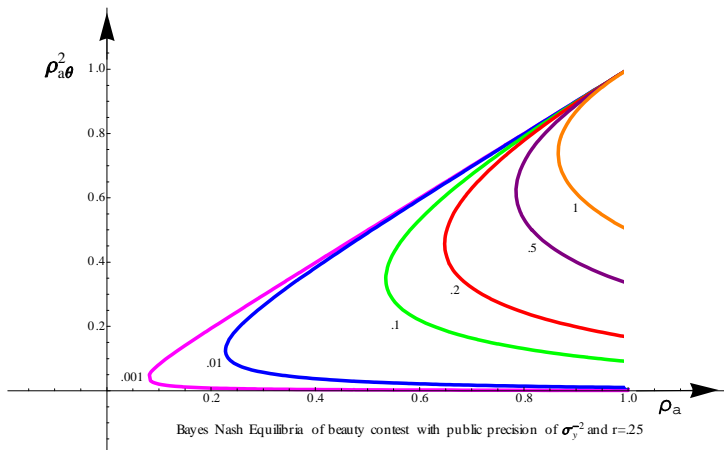
$$\begin{pmatrix} a_i \\ A \\ \theta \end{pmatrix} \sim N \left(\begin{pmatrix} \mu_a \\ \mu_A \\ \mu_\theta \end{pmatrix}, \begin{pmatrix} \sigma_a^2 & \rho_{aA}\sigma_a\sigma_A & \rho_{a\theta}\sigma_a\sigma_\theta \\ \rho_{aA}\sigma_a\sigma_A & \sigma_A^2 & \rho_{A\theta}\sigma_A\sigma_\theta \\ \rho_{a\theta}\sigma_a\sigma_\theta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

- and in terms of the equilibrium coefficients:


$$\begin{array}{ccc} \alpha_x^2\sigma_x^2 + \alpha_y^2\sigma_y^2 + \sigma_\theta^2(\alpha_x + \alpha_y)^2 & \alpha_y^2\sigma_y^2 + \sigma_\theta^2(\alpha_x + \alpha_y)^2 & \sigma_\theta^2(\alpha_x + \alpha_y) \\ \alpha_y^2\sigma_y^2 + \sigma_\theta^2(\alpha_x + \alpha_y)^2 & \alpha_x^2\sigma_x^2 + \sigma_\theta^2(\alpha_x + \alpha_y)^2 & \sigma_\theta^2(\alpha_x + \alpha_y) \\ \sigma_\theta^2(\alpha_x + \alpha_y) & \sigma_\theta^2(\alpha_x + \alpha_y) & \sigma_\theta^2 \end{array}$$

Given Public Information

- movements along level curve are variations in σ_x^{-2} given σ_y^{-2} :



- movements along level curve are variations in σ_y^{-2} given



- the object of analysis: joint distribution over actions and states:

$$\mu(a, A, \theta)$$

- characterize the set of (normally distributed) BCE:

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

- σ_A^2 is the aggregate volatility (common variation)
- $\sigma_a^2 - \sigma_A^2$ is the cross-section dispersion (idiosyncratic variation)
- statistical representation of equilibrium in terms of first and second order moments

Symmetric Bayes Correlated Equilibria

- with focus on symmetric equilibria:

$$\mu_A = \mu_a, \quad \sigma_A^2 = \rho_a \sigma_a^2, \quad \rho_{aA} \sigma_a \sigma_A = \rho_a \sigma_a^2$$

where ρ_a is the correlation coefficient across individual actions

- the first and second moments of the correlated equilibria are:

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

- correlated equilibria are characterized by:

$$\{\mu_a, \sigma_a, \rho_a, \rho_{a\theta}\}$$

- in the complete information game, the best response is:

$$a = -\theta \frac{\gamma_{a\theta}}{\gamma_a} - A \frac{\gamma_{Aa}}{\gamma_a}$$

- best response is weighted linear combination of fundamental θ and average action A relative to the cost of action:

$$\gamma_{a\theta}/\gamma_a, \gamma_{Aa}/\gamma_a$$

- in the incomplete information game, θ and A are uncertain:

$$\mathbb{E}[\theta], \quad \mathbb{E}[A]$$

- given the correlated equilibrium distribution $\mu(a, \theta)$ we can use the conditional expectations:

$$\mathbb{E}_\mu[\theta | a], \quad \mathbb{E}_\mu[A | a]$$

- in the incomplete information game, the best response is:

$$a = -\mathbb{E}_\mu [\theta | a] \frac{\gamma_{a\theta}}{\gamma_a} - \mathbb{E}_\mu [A | a] \frac{\gamma_{Aa}}{\gamma_a}$$

- best response property has to hold for all $a \in \text{supp } \mu(a, \theta)$
- a fortiori, the best response property has to hold in expectations over all a :

$$\mathbb{E}_\mu [a] = \mathbb{E}_\mu \left[- \left(\mathbb{E}_\mu [\theta | a] \frac{\gamma_{a\theta}}{\gamma_a} + \mathbb{E}_\mu [A | a] \frac{\gamma_{Aa}}{\gamma_a} \right) \right]$$

- by the law of iterated expectation, or law of total expectation:

$$\mathbb{E}_\mu [\mathbb{E}_\mu [\theta | a]] = \mu_\theta, \quad \mathbb{E}_\mu [\mathbb{E}_\mu [A | a]] = \mathbb{E}_\mu [A] = \mathbb{E}_\mu [a],$$

- the best response property implies that for all $\mu(a, \theta)$:

$$\mathbb{E}_\mu [a] = \mathbb{E}_\mu \left[- \left(\mathbb{E}_\mu [\theta | a] \frac{\gamma_{a\theta}}{\gamma_a} + \mathbb{E}_\mu [A | a] \frac{\gamma_{Aa}}{\gamma_a} \right) \right]$$

or by the law of iterated expectation:

$$\mu_a = -\mu_\theta \frac{\gamma_{a\theta}}{\gamma_a} - \mu_a \frac{\gamma_{Aa}}{\gamma_a}$$

Theorem (First Moment)

In all Bayes correlated equilibria, the mean action is given by:

$$\mathbb{E}[a] = -\mu_\theta \frac{\gamma_{a\theta}}{\gamma_a + \gamma_{aA}}.$$

- result about “mean action” is independent of symmetry or normal distribution

Equilibrium Moments: Variance

- in any correlated equilibrium $\mu(a, \theta)$, best response demands

$$a = - \left(\mathbb{E}[\theta | a] \frac{\gamma_{a\theta}}{\gamma_a} + \mathbb{E}[A | a] \frac{\gamma_{Aa}}{\gamma_a} \right), \quad \forall a \in \text{supp } \mu(a, \theta)$$

- or varying in a

$$1 = - \left(\frac{\partial \mathbb{E}[\theta | a]}{\partial a} \frac{\gamma_{a\theta}}{\gamma_a} + \frac{\partial \mathbb{E}[A | a]}{\partial a} \frac{\gamma_{Aa}}{\gamma_a} \right),$$

- the change in the conditional expectation

$$\frac{\partial \mathbb{E}[\theta | a]}{\partial a}, \quad \frac{\partial \mathbb{E}[A | a]}{\partial a}$$

is a statement about the correlation between a, A, θ

Equilibrium Moment Restrictions

- the best response condition **and** the condition that $\Sigma_{a,A,\theta}$ forms a multivariate distribution, meaning that the variance-covariance matrix has to be positive definite
- we need to determine:

$$\{\sigma_a, \rho_a, \rho_{a\theta}\}$$

Theorem (Second Moment)

The triple $(\sigma_a, \rho_a, \rho_{a\theta})$ forms a Bayes correlated equilibrium iff:


$$\rho_a - \rho_{a\theta}^2 \geq 0,$$

and

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

Moment Restrictions: Correlation Coefficients

- the equilibrium set is characterized by inequality $\rho_a - \rho_{a\theta}^2 \geq 0$
- ρ_a : correlation of actions across agents; $\rho_{a\theta}$: correlation of actions and fundamental



Equivalence between BCE and BNE

- bivariate information structure which generates volatility (common signal) and dispersion (idiosyncratic signal)

Theorem

There is BCE with $(\rho_a, \rho_{a\theta})$ if and only if there is a BNE with (σ_x^2, σ_y^2) .

- a public and a private signal are sufficient to generate the entire set of correlated equilibria...
- but a given BCE does not uniquely identify the information environment of a BNE

- the analyst may not know how much private information the agents have, yet he may have a lower bound on how much information the agents have
- how does the set of BCE change with the lower bound
- assume that all agents observe a public signal y :

$$y = \theta + \varepsilon$$

and a private signal x_i

$$x_i = \theta + \varepsilon_i$$

with

$$\begin{pmatrix} \varepsilon \\ \varepsilon_i \end{pmatrix} \sim N \left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma_y^2 & 0 \\ 0 & \sigma_x^2 \end{pmatrix} \right)$$

Information Bounds and Correlated Equilibrium

- the equilibrium conditions are now augmented from for all “ a ” to “for all a, x, y ”:

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

- we determine $\sigma_a, \rho_{ax}, \rho_{ay}$ in terms of $\rho_a, \rho_{a\theta}$, e.g.:


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

- set of correlated equilibria is given by the inequalities:

$$\begin{aligned} \rho_a - \rho_{a\theta}^2 - \rho_{ay}^2 &\geq 0, \\ 1 - \rho_a - \rho_{ax} &\geq 0, \end{aligned}$$


Given Public Information

- describe the equilibrium set $C(\sigma) \in [0, 1]^2$ in terms of the noise of the signal pair $\sigma = (\sigma_x, \sigma_y)$
- the interior of each level curve describes the correlated equilibria with a given amount of public correlation
- movements along level curve are variations in σ_x^{-2} given σ_y^{-2}



Given Private and Public Information

- the interior intersection of the level curves generates the corresponding equilibrium set
- more information reduces the set of possible outcomes, because it adds incentive constraints but does not remove any correlation possibilities



Given Information and the Equilibrium Set

- describe the equilibrium set $C(\sigma) \in [0, 1]^2$ in terms of the noise of the signal pair $\sigma = (\sigma_s, \sigma_i)$

Theorem (Information Bounds)

- For all $\sigma < \sigma'$, $C(\sigma_s) \subset C(\sigma'_s)$;
- For all $\sigma < \sigma'$:

$$\min_{\rho_a \in C(\sigma)} \rho_a > \min_{\rho_a \in C(\sigma')} \rho_a;$$

- For all $\sigma < \sigma'$:

$$\min_{\rho_{a\theta} \in C(\sigma)} \rho_{a\theta} > \min_{\rho_{a\theta} \in C(\sigma')} \rho_{a\theta}.$$

- more private information shrinks the equilibrium set and makes predictions sharper
- general result in "Correlated Equilibrium in Games of Incomplete Information"

1 Predictions:

What restrictions are imposed by the structural model (u, ψ) on the observable endogenous statistics $(\mu_a, \sigma_a, \rho_a, \rho_{a\theta})$?

2 Identification:

What restrictions can be imposed/inferred on the structural model (u) by the observations of the outcome variables $(\mu_a, \mu_\theta, \sigma_a, \sigma_\theta, \rho_a, \rho_{a\theta})$?

- can we identify sign and size of interaction?
- can we identify the nature of the informational externality $\gamma_{a\theta}$ and the strategic externality γ_{aA}

Benchmark: Complete Information

- the state θ is observable by all agents
- the econometrician observes: $(\mu_a, \mu_\theta, \sigma_a, \sigma_\theta, \rho_a, \rho_{a\theta})$
- the information about the mean:

$$\mu_\alpha = \frac{\lambda_a + \mu_\theta \gamma_{a\theta}}{\gamma_a + \gamma_{aA}}$$

and covariance

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

permits identification of slope and intercept of best response:

$$\frac{\gamma_{a\theta}}{\gamma_a + \gamma_{aA}}, \frac{\lambda_a}{\gamma_a + \gamma_{aA}}$$

Proposition (Sign Identification)

1. Nash equilibrium identifies the sign of $\gamma_{a\theta}$, but not the sign of γ_{aA} .
2. Nash equilibrium point identifies the ratios:

$$\frac{\gamma_{a\theta}}{\gamma_a + \gamma_{aA}} \quad \text{and} \quad \frac{\lambda_a}{\gamma_a + \gamma_{aA}}.$$

- the sign of the strategic interaction, γ_{aA} , cannot be identified, due to the perfect correlation of the agents' action in the complete information game

Identification with Incomplete Information: Bayes Nash

- the information structure (σ_x^2, σ_y^2) of the Bayesian game is assumed to be known
- the identification, given the hypothesis of Bayes Nash equilibrium, uses the variance-covariance matrix of actions and states:

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

- the relationship between equilibrium coefficients

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

lends information about the sign of γ_{aA}

Identification with Incomplete Information: Bayes Nash

Proposition (Sign Identification)

If $0 < \sigma_x^2, \sigma_y^2 < \infty$, then:

1. BNE identifies the sign of $\gamma_{a\theta}$ **and** of γ_{aA} .
2. BNE point identifies the ratios:

$$\frac{\gamma_{a\theta}}{\gamma_a + \gamma_{aA}} \quad \text{and} \quad \frac{\lambda_a}{\gamma_a + \gamma_{aA}}.$$

- Bayes Nash equilibrium can recover sign of strategic interaction as opposed to Nash equilibrium

Identification with Incomplete Information: Bayes Correlated

- we can recover from the equilibrium conditions mean and variance:

$$\mu_a = \frac{\lambda_a + \gamma_{a\theta}\mu_\theta}{\gamma_a + \gamma_{aA}}$$

and

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

as well ρ_a and $\rho_{a\theta}$

- notice the rate at which γ_{aA} and γ_a substitute differ in mean and standard deviation, provided $\rho_a < 1$
- source of partial identification

Proposition (Partial Identification)

If $\rho_a < 1$, the interaction ratios are partially identified in BCE:


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

and

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

Moreover, the sign of $\gamma_{a\theta}$, but not the sign of γ_{aA} is identified.

- discontinuity of identified set at $\rho_a = 1$



- failure to identify the strategic nature of the game, strategic complements or strategic substitutes

Prior Information and Identification

- earlier, we analyzed how information bounds in terms of private and public information, represented by σ_x^2 and σ_y^2 , sharpen the prediction
- similarly, the information bounds sharpen the identification of the interaction ratios

Proposition (Prior Information and Identification)

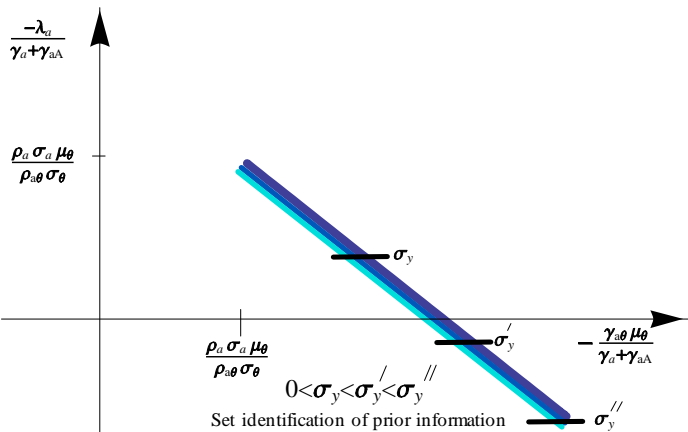
With prior information, the interaction ratios are sharper identified with:

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

with

$$\frac{\partial r(\sigma_x, \sigma_y)}{\partial \sigma_x} < 0, \quad \frac{\partial R(\sigma_x, \sigma_y)}{\partial \sigma_y} > 0.$$

Prior Information and Set Identification



- the private information set provides the lower bound on γ_{aA}
- the public information set provides the upper bound on γ_{aA}
- eventually, as $\sigma_x^2, \sigma_y^2 \rightarrow 0$, the sign of the strategic interaction is identified

Aggregate Information and Identification

- suppose we could only observe the aggregate (average) action, but not idiosyncratic action profile
- would we lose (much) in identification power
- in Bayes Nash equilibrium, we would not lose any identification result,
- in Bayes Correlated equilibrium, we lose information, as remaining restriction is weak:

$$\rho_a \geq \rho_{a\theta}^2$$

Demand and Supply Identification

- demand is given by:

$$P_d = d_0 + d_1 Q + d_2 \theta_d$$

- supply is given by:

$$P_s = s_0 + s_1 Q + s_2 \theta_s$$

- the exogenous random variables are θ_d and θ_s are demand and supply shocks ("demand, supply shifters")
- **complete information:** each firm observes shocks (θ_d, θ_s) and makes supply decisions accordingly
- econometrician only observes realized aggregate variables, P and Q , but not individual choices

- each firm observes a vector of signals:

$$y_i = (y_d, y_{di}, y_s, y_{si})$$

regarding the true cost and demand shocks:

$$y_s = \theta_s + \varepsilon_s, y_{si} = \theta_s + \varepsilon_{si}, y_d = \theta_d + \varepsilon_d, y_{di} = \theta_d + \varepsilon_{di}$$

before its supply decision

- maintain normality and independence of $(\theta_d, \theta_s, \varepsilon_d, \varepsilon_{id}, \varepsilon_s, \varepsilon_{is})$
- the variance of the random variables:

$$\mathcal{I} = (\lambda_d^2, \lambda_s^2, \sigma_d^2, \sigma_{id}^2, \sigma_s^2, \sigma_{is}^2)$$

represents the information structure \mathcal{I} in the economy

Competitive Equilibrium with Incomplete Information

- in Bayes Nash equilibrium of competitive economy firm i supplies

$$q_i(y_i)$$

on the basis of its private, but noisy, information

- equilibrium price is given by equilibrium condition:

$$\begin{aligned} P_d &= d_0 + d_1 \int q_i(y_i) di + d_2 \theta_d \\ &= d_0 + d_1 Q + d_2 \theta_d \end{aligned}$$

with respect to the realized demand shock θ_d

Identification with Incomplete Information

- can we identify and estimate the slope of supply and demand function in the presence of incomplete information?
- for every information structure \mathcal{I} each firm observes a noisy signal of the true cost and demand shock and makes a supply decision on the basis of the noisy information

Theorem (Point Identification)

For every information structure \mathcal{I} , the demand and supply functions are **point identified** if the firms have noisy information about their cost:

$$\min \{ \sigma_s^2, \sigma_{is}^2 \} < \infty.$$

- asymmetry arises as realized price varies with realized demand

- now we ask can identification be accomplished for all (or a subset of) information structures, i.e. can we achieve “robust identification”
- demand function remains point identified as aggregate quantity, aggregate price and demand shock are observed

Theorem (Set Identification)

- ① *For every information bound, the supply function (s_1, s_2) is **set identified**.*
 - ② *If the information bounds increase, then the **identified set** decreases.*
- concern for robustness weakens the ability to identify the structural parameter to allow for **partial identification** only

- Bayes correlated equilibrium encodes concern for robustness to information environment
- identification: complete vs. incomplete information
 - demand and supply identification: the market participants have complete information, but the analyst has noisy information (uses instrument to recover the information)
 - auction identification: each bidder has private information about his valuation, but the bidders have the same information about each other as the analyst
 - presumably neither informational assumption is valid, suggesting a role for robust identification
- robust welfare improving policy
 - how responsive is robust policy to informational conditions