Robust Predictions in Games with Incomplete Information

Dirk Bergemann and Stephen Morris

April 2011

- in games of incomplete information, private information represents information about:
  payoff environment and strategic environment
- first order beliefs vs. higher order beliefs
- game theoretic predictions are very sensitive to specification of the strategic environment, the “higher order beliefs", or (equivalently) the information structure
- email game, coordination games in game theory
- revenue equivalence, surplus extraction in mechanism design

The Role of Higher Order Beliefs: An Example

- first price sealed bid auction with independent, private values
- identical, independent common prior over valuations of bidders
- first order beliefs of bidder i is about:
  - valuation of bidder i himself; valuations of bidders j, k
- second order beliefs of bidder i is about:
  - the first order belief of bidder j about the valuation of bidder i
- in first price sealed bid auction, bid of i depends on his first-order, second-order and higher-order beliefs
- revenue equivalence - assuming second-order belief of i puts probability 1 on j's first order belief being equal to common prior
- i.e. private information of bidder i doesn't contain any additional information about bidder j beyond common prior; revenue equivalence fails without this restrictive assumption
- what predictions can we offer independent of the higher order beliefs?

Payoff Environment and Strategic Environment

- fix "payoff relevant environment"
  - = action sets, payoff-relevant variables ("states"), payoff functions, distribution over states
  - = incomplete information game without higher order beliefs about states
- there are many strategic environments consistent with fixed payoff environment
- consistent: after integrating over the higher-order beliefs, the marginal over payoff relevant states coincides with common prior over payoff relevant states
- the possible information environments vary widely:
  from "zero information", where every agent knows nothing beyond common prior to "complete information", where every agents knows realization of payoff relevant state

Agenda: Robust Predictions

- analyze what could happen for all possible higher order beliefs (maintaining common prior assumption and equilibrium assumptions)
- each specific information environment generates specific predictions regarding equilibrium behavior
- given that the different strategic environments share the same payoff environment, does the predicted behavior display common features?
- can analyst make predictions which are robust to the exact specification of the strategic environment?
- make set valued predictions about joint distribution of actions and states

Agenda: Prior Information and Prediction

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

Agenda: Robust Identification

- the observable outcomes of the game are the actions and the payoff relevant states
- the chosen action reveals the preference of the agent given his higher-order belief, but typically does not reveal his higher order belief
- 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 , or some moments of it, what can be deduced about the payoff relevant environment?
- partial identification / set identification

This Talk

1. General Approach
   - set valued prediction is set of "Bayes Correlated equilibria"
   - epistemic result linking Bayes Nash and Bayes Correlated equilibrium
   - partial prior information monotonically reduces the set of "Bayes Correlated equilibria"
2. Illustration with Continuum Player, Symmetric, Linear Best Response, Normal Distribution Games
   - the resulting equilibrium sets are tractable and intuitive
   - cannot distinguish between games with strategic substitutes and complements

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

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

- 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

$$\pi : \Theta \to \Delta(T_1 \times T_2 \times ... \times T_I)$$

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

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

- common prior $\psi(\theta)$ and conditional distribution

$$\pi[t_i, t_{-i}](\theta)$$

allow agent i to hold private information t_i in terms of

- posterior beliefs about the payoff relevant state θ
- posterior beliefs about the beliefs t_{-i} of the other agents

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

$$\sigma_i : T_i \to \Delta(A_i)$$

Definition (Bayes Nash Equilibrium (BNE))

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

$$\sum_{t_{-i},\theta} u_i\left(\left(\sigma_i(t_i), \sigma_{-i}(t_i)\right), \theta\right) \pi[t_i, t_{-i}](\theta)\, \psi(\theta)$$
$$\geq \sum_{t_{-i},\theta} u_i\left(\left(a_i, \sigma_{-i}(t_{-i})\right), \theta\right) \pi[t_i, t_{-i}](\theta)\, \psi(\theta).$$

for each i, t_i and a_i.

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

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

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

Definition (Bayes Nash Equilibrium Distribution)

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

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

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

$$\sum_{t_{-i},\theta} u_i\left(\left(\sigma_i\left(t_i\right), \sigma_{-i}\left(t_i\right)\right), \theta\right) \psi\left(\theta\right) \pi\left[t_i, t_{-i}\right]\left(\theta\right)$$

$$\geq \sum_{t_{-i},\theta} u_i\left(\left(a_i, \sigma_{-i}\left(t_{-i}\right)\right), \theta\right) \psi\left(\theta\right) \pi\left[t_i, t_{-i}\right]\left(\theta\right).$$

- with the equilibrium distribution

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

- an implication of BNE of $(G, \mathcal{T})$: for all $i, a_i \in \operatorname{supp} \mu_\sigma\left(a, \theta\right)$:

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

- joint distribution over action-state $\mu(a, \theta)$ describes choices, but is silent about reason (information) for choice

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',

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

and *consistent*, i.e., for each θ

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

- no restrictions on private information beyond $\psi(\theta)$; zero information and complete information of θ are possible

- BCE is defined in terms of the payoff environment and without reference to type space:

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

- without uncertainty, $\Theta = \{\theta\}$, Bayes correlated equilibrium reduces to correlated equilibrium (Aumann (1974))
- earlier definitions of correlated equilibrium for games of incomplete information, see Forges (1993), (2006), define solution concepts for $(G, \mathcal{T})$, integrate out payoff relevant states θ, and describe correlated equilibrium as action type distributions $\Delta(A \times T)$
- we work with a basic game G, describe correlated equilibrium as action state distributions $\Delta(A \times \Theta)$
- in companion paper, "Correlated Equilibrium in Games with Incomplete Information" we relate definitions and establish comparative results wrt information environments

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

Theorem (Equivalence)

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

- $BCE \Rightarrow BNE$ uses the richness of the possible information structure to complete the equivalence result:

$$\forall i,\ \forall a_i:\ t_i = \mu(\theta, a_{-i} | a_i),\ \sigma_i(t_i) = a_i$$

- Aumann (1987) established the above characterization result for complete information games

- 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

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

  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

- the matrix Γ defines the nature of the interaction:

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

- diagonal entries: $\gamma_a, \gamma_A, \gamma_\theta$ describe "own effects"
- off-diagonal entries: $\gamma_{a\theta}, \gamma_{A\theta}, \gamma_{aA}$ "interaction effects"
- best response of agent i depends on $\gamma_a, \gamma_{aA}, \gamma_{a\theta}$:
  - cost of adjustment: γ_a
  - informational externality: $\gamma_{a\theta}$
  - strategic externality: γ_{aA}; strategic complements and strategic substitutes, $\gamma_{aA} > 0$ vs. $\gamma_{aA} < 0$

- suppose θ is commonly known
- best response of agent i :

$$a = -\frac{\lambda_a + \gamma_{a\theta} \cdot \theta + \gamma_{aA} \cdot .A}{\gamma_a}$$

- equilibrium response of agent i:

$$a^* (\theta) = -\overbrace{\frac{\lambda_a}{\gamma_a + \gamma_{aA}}}^{\text{intercept}} - \overbrace{\frac{\gamma_{a\theta}}{\gamma_a + \gamma_{aA}}}^{\text{slope}} \cdot \theta$$

- maintained assumptions:
  - concavity at individual level (well-defined best response): $\gamma_a < 0$
  - concavity at aggregate level (existence of interior equilibrium): $\gamma_a + \gamma_{aA} < 0$
- concave payoffs imply that complete information game has unique Nash **and** correlated equilibrium (Neyman (1997))

- 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)

Example 2: Beauty Contest, Macroeconomic Coordination Games

- 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)

Standard Approach

- fix an arbitrary information system:

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

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

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

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

- binary information structure with a private and a public component
- every pair $\left(\sigma_x^2, \sigma_y^2\right)$ generates a different information system, type space

- the best response of each agent is:

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

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

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

- 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:

$$s^*(x,y) = \alpha_0^* + \alpha_x^* x + \alpha_y^* y$$

with

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

and

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

- the linear coefficients α_x^* and α_y^* satisfy the relationship:

$$\frac{\alpha_y^*}{\alpha_x^*} = \frac{\sigma_y^{-2}}{\sigma_x^{-2}} \frac{\gamma_a}{\gamma_a + \gamma_{aA}}, \text{ w/o strategic interaction } \gamma_{aA} = 0 : \frac{\alpha_y^*}{\alpha_x^*} = \frac{\sigma_y^{-2}}{\sigma_x^{-2}}$$

- linear coefficients α_x^* and α_y^* respond to signals x_i and y:

$$x_i = \theta + \varepsilon_i \text{ and } y = \theta + \varepsilon,$$

  and hence to state of the world θ
- sum of linear coefficients α_x^* and α_y^* display "attenuation":

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

  yet mean action:

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

  for all σ_x^2, σ_y^2
- first equilibrium moment is constant across information structures
- second equilibrium moments vary across information structures.

Joint Action State Distribution

- given normally distributed random variables $\left(\sigma_\theta^2, \sigma_x^2, \sigma_y^2\right)$ and linear strategy, the joint distribution of action and state $\mu\left(a_i, a_j, \theta\right)$ is multivariate normal as well:

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

- and in terms of equilibrium coefficients:

$$\begin{pmatrix} \sigma_a^2 \\ \rho_a\sigma_a^2 \\ \rho_{a\theta}\sigma_a\sigma_\theta \end{pmatrix} = \begin{pmatrix} \alpha_x^2\sigma_x^2 + \alpha_y^2\sigma_y^2 + \sigma_\theta^2\left(\alpha_x + \alpha_y\right)^2 \\ \alpha_y^2\sigma_y^2 + \sigma_\theta^2\left(\alpha_x + \alpha_y\right)^2 \\ \sigma_\theta^2\left(\alpha_x + \alpha_y\right) \end{pmatrix}$$

- represent equilibrium in terms of correlation coefficients $\rho_a, \rho_{a\theta}$

Equilibrium and Private Information

- equilibrium covariances via correlation coefficients $\rho_a, \rho_{a\theta}$
- consider given precision σ_x^{-2} of private information

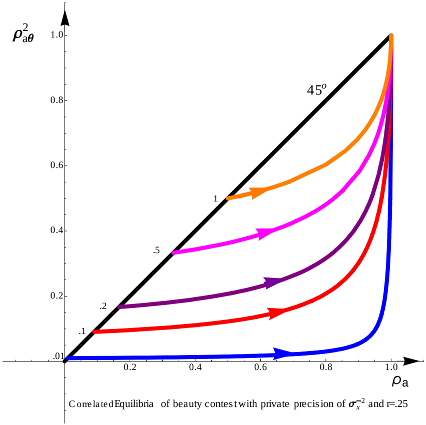

- increase in precision of public information σ_y^{-2} is upward movement along level curve

Equilibrium and Public Information

- equilibrium covariances via correlation coefficients $\rho_a, \rho_{a\theta}$
- consider given precision σ_y^{-2} of public information

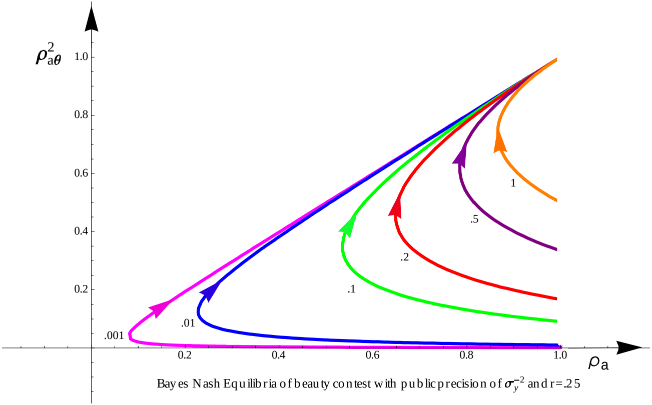

Bayes Nash Equilibria of beauty contest with public precision of σ_y^{-2} and r=.25

- increase in precision of private information σ_x^{-2} is upward movement along level curve

Multitude of Information Environments

- every type t_i of agent i could contain many pieces of information

  $$t_i = (s, s_i, s_{ij}, s_{ijk}, ...)$$

  every agent i may observe a public (common) signal s centered around the state of the world θ:

  $$s \sim N\left(\theta, \sigma_s^2\right)$$

- every agent i may observe a private signal s_i centered around the state of the world θ:

  $$s_i \sim N\left(\theta, \sigma_i^2\right)$$

- every agent i may observe a private signal $s_{i,j}$ about the signal of agent j :

  $$s_{i,j} \sim N\left(s_j, \sigma_{i,j}^2\right)$$

- every agent i may observe a private signal $s_{i,j,k}$ about ....:

  $$s_{i,j,k} \sim N\left(s_{j,k}, \sigma_{i,j,k}^2\right)$$

- object of analysis: joint distribution over actions and states

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

- for exposition, focus on symmetric Bayes correlated equilibria:

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

which are normally distributed:

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

- 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

- with focus on symmetric equilibria:

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

- 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 characterized by first and second moments:

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

- 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: $\gamma_{a\theta}/\gamma_a$, γ_{Aa}/γ_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:

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

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

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

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

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

- 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

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

- 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
- "law of total variance"

- 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

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

- correlation of actions across agents: ρ_a
- correlation of actions and fundamentals: $\rho_{a\theta}$

Moment Restrictions: Correlation Coefficients

- the equilibrium set is completely characterized by inequality:

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

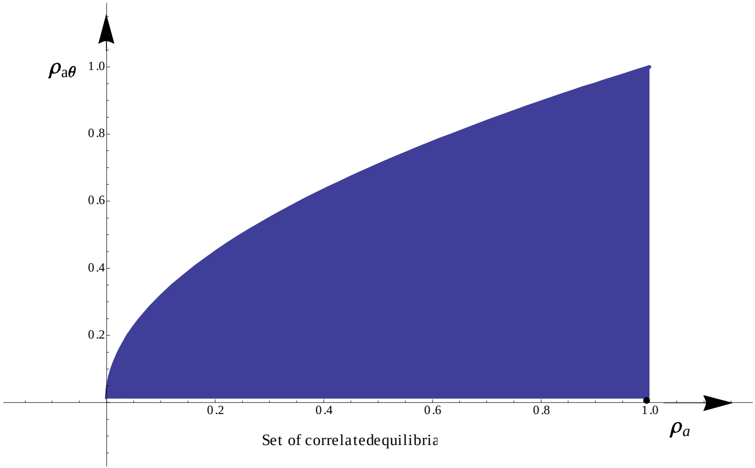

Equivalence between BCE and BNE

- 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 in 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" as additional "incentive constraints"

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

- we determine σ_a, ρ_{ax}, ρ_{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:

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

- movements along level curve are variations in σ_x^{-2} given σ_y^{-2}
- the interior of each level curve describes the correlated equilibria for a given lower bound on public information

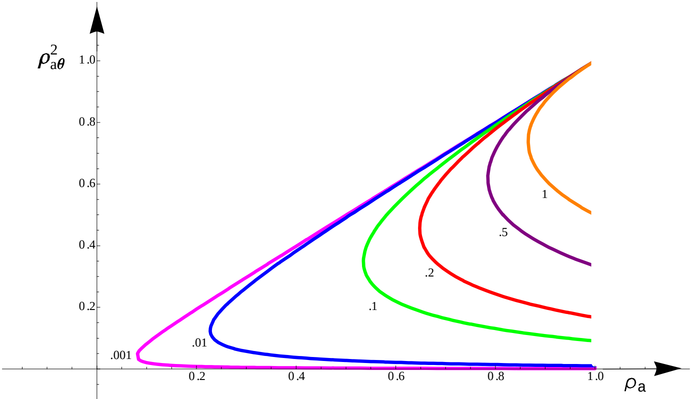

Correlated Equilibria of beauty contest with minimal precsion of σ_y^{-2} and r=.25

Lower Bounds on Private and Public Information

- interior of intersection of level curves is the set of Bayes correlated equilibrium subject to lower bounds on private and public information

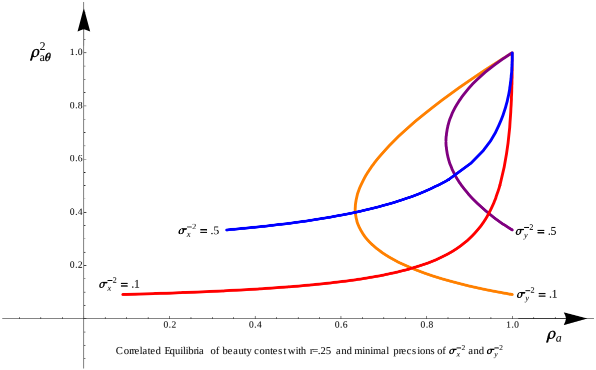

- more information reduces set of possible distributions, as it adds incentive constraints but does not remove correlation possibilities

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}

Identification and Information: Single Agent

- zero interaction: $\gamma_{aA} = 0$,
- normalize cost of adjustment: $\gamma_a = -1$
- complete information: the state θ is observable to the agent
- the best response of the agent is given by:

$$a(\theta) = \gamma_{a\theta} \cdot \theta + u,$$

  where u is an error, observable to the agent, but unobservable to the econometrician
- unobservable error u, $\mathbb{E}[u] = u$, $u \perp \theta$
- the econometrician observes a and θ but not u
- the covariance between a and θ point identifies $\gamma_{a\theta}$

Identification and Noisy Information

- noisy information: the state θ is unobservable, the agent only observes a noisy signal

$$s = \theta + \varepsilon$$

with $\varepsilon \sim N\left(0, \sigma_\varepsilon^2\right)$

- the agent chooses a upon observing s and u,

$$a\left(s\right) = \gamma_{a\theta} \cdot \frac{s\sigma_\varepsilon^{-2} + \mu_\theta \sigma_\theta^{-2}}{\sigma_\varepsilon^{-2} + \sigma_\theta^{-2}} + u$$

- the econometrician observes a and θ but neither s nor u
- the best response of the agent has the slope:

$$\gamma_{a\theta} \cdot \frac{\sigma_\varepsilon^{-2}}{\sigma_\varepsilon^{-2} + \sigma_\theta^{-2}} < \gamma_{a\theta},$$

- the best response is attenuated by the signal to noise ratio

- the covariance between a and θ point identifies:

$$\widehat{\gamma}_{a\theta} = \gamma_{a\theta} \cdot \frac{\sigma_{\varepsilon}^{-2}}{\sigma_{\varepsilon}^{-2} + \sigma_{\theta}^{-2}},$$

- how to identify $\gamma_{a\theta}$ without knowing the information of the agent, i.e. without knowing the variance σ_{ε}^{2} of ε :
- without knowing σ_{ε}^{2}, there is only set identification:

$$\gamma_{a\theta} = \widehat{\gamma}_{a\theta} \cdot \frac{\sigma_{\varepsilon}^{-2} + \sigma_{\theta}^{-2}}{\sigma_{\varepsilon}^{-2}} \quad \Rightarrow \gamma_{a\theta} \in [\widehat{\gamma}_{a\theta}, \ \infty)$$

- noise in the predictor variable induces a bias: "attenuation bias", "regression dilution"
- in the presence of interaction, "attenuation bias" impacts identification of other structural parameters, here the strategic externality $\gamma_{a\theta}$

Identification and Information: Many Agents

- non-zero interaction: $\gamma_{aA} \neq 0$
- information structure $\left(\sigma_x^2, \sigma_y^2\right)$ of Bayesian game is assumed to be known
- the identification, given the hypothesis of BNE, uses variance-covariance matrix of actions and states:

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

- the relationship between equilibrium coefficients

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

lends information about the sign of γ_{aA}

- information structure $\left(\sigma_x^2, \sigma_y^2\right)$ of Bayesian game is assumed to be known

Proposition (Sign and Point 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 slope of the equilibrium response:*

$$\frac{\gamma_{a\theta}}{\gamma_a + \gamma_{aA}}.$$

- information structure of Bayesian game is assumed to be unknown
- we can use observed mean, variance, and covariance of (a, θ) and equilibrium conditions:

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

and

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

Proposition (Partial Identification)

1. *The sign of $\gamma_{a\theta}$ is, but the sign of γ_{aA} is not identified.*
2. *If $\rho_a < 1$, the slope of the equilibrium response is partially identified:*

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

- failure to identify the strategic nature of the game, the sign of γ_{aA} describes whether the game is one of strategic complements or strategic substitutes

- 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:

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

and

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

Prior Information and Set Identification

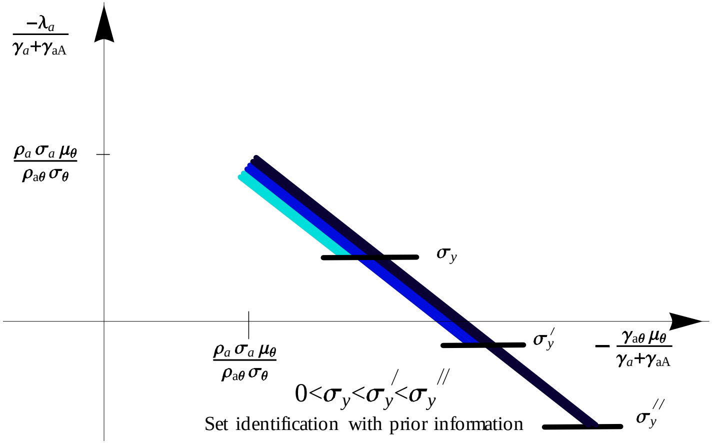

- the private information set provides the lower bound on γ_{aA}
- the public information set provides the upper bound on γ_{aA}
- as $\sigma_x^2, \sigma_y^2 \longrightarrow 0$, sign of strategic interaction is identified

Demand and Supply Identification

- linear inverse demand is given by:

$$P_d = \lambda_d + \gamma_{aA} Q + \gamma_{a\theta_d} \theta_d$$

- linear inverse supply is given by:

$$P_s = \lambda_s + \gamma_a Q + \gamma_{a\theta_s} \theta_s$$

- two-dimensional uncertainty: θ_d and θ_s are demand and supply shocks ("demand, supply shifters")
- firm i makes supply decision with noisy information about (θ_d, θ_s)
- only aggregate data is observed: aggregate quantity and price
- in Bayes Nash equilibrium $(\gamma_a, \gamma_{aA}, \gamma_{a\theta_d}, \gamma_{a\theta_s})$ are point identified
- in Bayes Correlated equilibrium γ_a and $\gamma_{a\theta_s}$ are only set identified

Discussion

- Bayes correlated equilibrium encodes concern for robustness to strategic information environment
- next items on the agenda: robust comparative statics, robust 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