Learning to Coordinate

Very preliminary - Comments welcome

Edouard Schaal[1] Mathieu Taschereau-Dumouchel[2]

[1]New York University

[2]Wharton School
University of Pennsylvania

Introduction

- We want to understand how agents learn to coordinate in a dynamic environment
- In the global game approach to coordination, information determines how agents coordinate
  - In most models, information comes from various exogenous signals
  - In reality, agents learn from endogenous sources (prices, aggregates, social interactions, ...)
    - Informativeness of endogenous sources depends on agents' decisions
- We find that the interaction of coordination and learning generates interesting dynamics
  - The mechanism dampens the impact of small shocks...
  - ...but amplifies and propagates large shocks

Overview of the Mechanism

- Dynamic coordination game
  - Payoff of action depends on actions of others and on unobserved fundamental θ
  - Agents use private and public information about θ
  - Observables (output,...) aggregate individual decisions
- These observables are non-linear aggregators of private information
  - When public information is very good or very bad, agents rely less on their private information
  - The observables becomes less informative
  - Learning is impeded and the economy can deviate from fundamental for a long time

Roadmap

- Stylized game-theoretic framework
  - Characterize equilibria and derive conditions for uniqueness
  - Explore relationship between decisions and information
  - Study the planner's problem
  - Provide numerical examples and simulations along the way

Abbreviated Literature Review

- Learning from endogenous variables
  - Angeletos and Werning (2004); Hellwig, Mukherji and Tsyvinksi (2005): static, linear-Gaussian framework (constant informativeness)
  - Angeletos, Hellwig and Pavan (2007): dynamic environment, non-linear learning, fixed fundamental, stylized cannot be generalized
  - Chamley (1999): stylized model with cycles, learning from actions of others, public signal is fully revealing upon regime change and uninformative otherwise

Model

- Infinite horizon model in discrete time
- Mass 1 of risk-neutral agents indexed by $i \in [0, 1]$
- Agents live for one period and are then replaced by new entrant
- Each agent has a project that can either be undertaken or not

Model

Realizing the project pays

$$\pi_{it} = (1-\beta)\,\theta_t + \beta m_t - c$$

where:

- θ_t is the **fundamental** of the economy
  - Two-state Markov process $\theta_t \in \{\theta_l, \theta_h\}$, $\theta_h > \theta_l$ with

$$P(\theta_t = \theta_j | \theta_{t-1} = \theta_i) = P_{ij} \text{ and } P_{ii} > \frac{1}{2}$$

- m_t is the mass of undertaken projects plus some noise
- β determines the degree of complementarity in the agents payoff
- $c > 0$ is a fixed cost of undertaking the project

Information

Agents do not observe θ directly but have access to several sources of information

1. A private signal v_{it}
   - Drawn from cdf G_θ for $\theta \in \{\theta_l, \theta_h\}$ with support $v \in [a, b]$
   - G_θ are continuously differentiable with pdf g_θ
   - Monotone likelihood ratio property: $g_h(v)/g_l(v)$ is increasing
2. An exogenous public signal z_t drawn from cdf F^z_θ and pdf f^z_θ
3. An endogenous public signal m_t
   - Agents observe the mass of projects realized with some additive noise ν_t

$$m_t(\theta, \hat{v}) = \text{mass of projects realized} + \nu_t$$

   - $\nu_t \sim$ iid cdf F^ν with associated pdf f^ν
   - Assume without loss of generality that F^ν has mean 0

Timing

Agents start with the knowledge of past public signals z_t and m_t

1. θ_t is realized
2. Private signals v_{it} are observed
3. Decisions are made
4. Public signals m_t and z_t are observed

Information

Information sets:

- At time t, the public information is

$$\mathcal{F}_t = \left\{ m^{t-1}, z^{t-1} \right\}$$

- Agent i's information is

$$\mathcal{F}_{it} = \{ v_{it} \} \cup \mathcal{F}_t$$

Beliefs:

- Beliefs of agent i about the state of the world

$$p_{it} = P\left(\theta = \theta_h | \mathcal{F}_{it}\right)$$

- Beliefs of an **outside observer** without private information

$$p_t = P\left(\theta = \theta_h | \mathcal{F}_t\right)$$

Agent's Problem

Agents i realizes the project if its expected value is positive

$$E\left[(1-\beta)\theta_t + \beta m_t - c|\mathcal{F}_{it}\right] > 0$$

For now, restrict attention to **monotone strategy equilibria**:

- There is a threshold $\hat{v}_t$ such that

$$\text{Agent } i \text{undertakes his project } \Leftrightarrow v_{it} \geq \hat{v}_t$$

- Later, we show that all equilibria have this form
- With this threshold strategy, the endogenous public signal is

$$m_t = \underbrace{1 - G_\theta\left(\hat{v}_t\right)}_{\text{signal}} + \underbrace{\nu_t}_{\text{noise}}$$

Dynamics of Information

- For a given signal s_t, beliefs are updated using the likelihood ratio

$$LR_{it} = \frac{P\left(s_t \mid \theta_h, \mathcal{F}_{it}\right)}{P\left(s_t \mid \theta_l, \mathcal{F}_{it}\right)}$$

- Using Bayes' rule, we have the following updating rule

$$P\left(\theta_h \mid p_{it}, s_t\right) = \frac{1}{1 + \frac{1-p_{it}}{p_{it}} LR_{it}^{-1}} := \mathcal{L}\left(p_{it}, LR_{it}\right)$$

Dynamics of Information

- At the beginning of every period, the individual beliefs are given by

$$p_{it}\left(p_t, v_{it}\right) = \mathcal{L}\left(p_t, \frac{g_h\left(v_{it}\right)}{g_l\left(v_{it}\right)}\right)$$

- By the end of the period, public beliefs p_t are updated according to

$$p_t^{end} = \mathcal{L}\left(p_t, \frac{f_h^z\left(z_t\right)}{f_l^z\left(z_t\right)} \frac{P\left(m_t | \theta_h, \mathcal{F}_t\right)}{P\left(m_t | \theta_l, \mathcal{F}_t\right)}\right)$$

- Moving to the next period,

$$p_{t+1} = p_t^{end} P_{hh} + \left(1 - p_t^{end}\right) P_{lh}$$

Full expression for dynamic of p

Dynamics of Information

Lemma 1

The distribution of individual beliefs is entirely described by (θ, p):

$$P\left(p_i \leq \tilde{p} | \theta, p\right) = \int \mathbb{1}\left(\frac{1}{1 + \frac{1-p}{p}\frac{g_l(v_i)}{g_h(v_i)}} \leq \tilde{p}\right) dG_\theta\left(v_i\right).$$

- Conditional on θ agents know that all signals come from G_θ
- From G_θ and p they can construct the distribution of beliefs
- Rich structure of higher-order beliefs in the background

Monotone Strategy Equilibrium

Definition

A *monotone strategy equilibrium* is a threshold function $\hat{v}(p)$ and an endogenous public signal m such that

1. Agent i realizes his project if and only if his v_i is higher than $\hat{v}(p)$
2. The public signal m is defined by $m = 1 - G_\theta\left(\hat{v}\left(p\right)\right) + \nu$
3. Public and private beliefs are consistent with Bayesian learning

Given the payoff function

$$\pi\left(v_i; \hat{v}, p\right) = E\left[\left(1-\beta\right)\theta + \beta\left(1 - G_\theta\left(\hat{v}\right)\right) - c \mid p, v_i\right]$$

the threshold function $\hat{v}(p)$ satisfies

$$\pi\left(\hat{v}(p); \hat{v}(p), p\right) = 0$$

for every p.

Equilibrium Characterization: **Complete** Information

Lemma 2 (Complete info)

If $\beta \geq c - (1-\beta)\,\theta \geq 0$*, the economy admits multiple equilibria under complete information.*

In particular, there is an equilibrium in which all projects are undertaken and one equilibrium in which no projects are undertaken.

Equilibrium Characterization: **Incomplete** Information

Assumption 1

The likelihood ratio $\frac{g_h}{g_l}$ *is differentiable and there exists* $\underline{\rho} > 0$ *such that*

$$\left| \left(\frac{g_h}{g_l} \right)' \right| \geq \underline{\rho}.$$

Proposition 1 (Incomplete info)

Under assumption 1,

1. If $\frac{\beta}{1-\beta} \leq \theta_h - \theta_l$, all equilibria are monotone,
2. If $\frac{\beta}{1-\beta} \leq \frac{\underline{\rho} P_{hl} P_{lh}}{\max\{\|g_h\|,\|g_l\|\}^3}$, there exists a unique equilibrium.

Uniqueness requires:

1. an upper bound on β; Role of β
2. enough beliefs dispersion. Role of dispersion

Endogenous vs Exogenous Information

Sample path with only **exogenous** information:

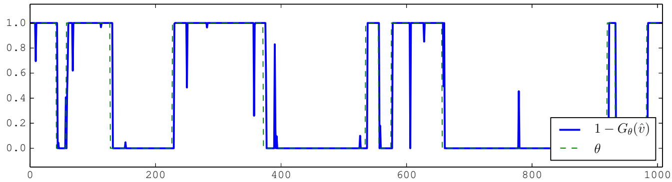

Sample path with only **endogenous** information:

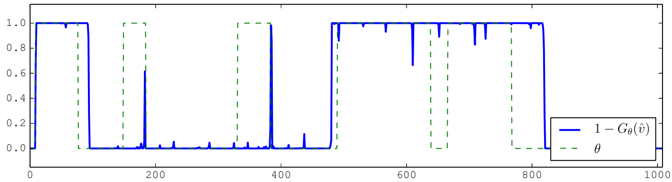

From now on, focus on endogenous public signal only: $\text{Var}(z_t) \to \infty$

Endogenous Information

Lemma 3

If $F^{\nu} \sim \mathcal{N}(0, \sigma_{\nu}^{2})$, then the mutual information between θ and m is

$$I(\theta; m) = p\,(1-p)\,\frac{\Delta^2}{2\sigma_{\nu}^{2}} + O\left(\Delta^3\right)$$

where $\Delta = G_l\left(\hat{v}\right) - G_h\left(\hat{v}\right) \geq 0$.

Version of the Lemma with general F^{ν}: ▸ General Lemma

The informativeness of the public signal depends on:

1. The current beliefs p
2. The amount of noise σ_{ν} added to the signal
3. The difference between $G_l\left(\hat{v}\right)$ and $G_h\left(\hat{v}\right)$

Point 3 is the source of **endogenous information.** ▸ Definition of mutual information

Signal vs. Noise

Example 1: Normal case with different means $\mu_h > \mu_l$

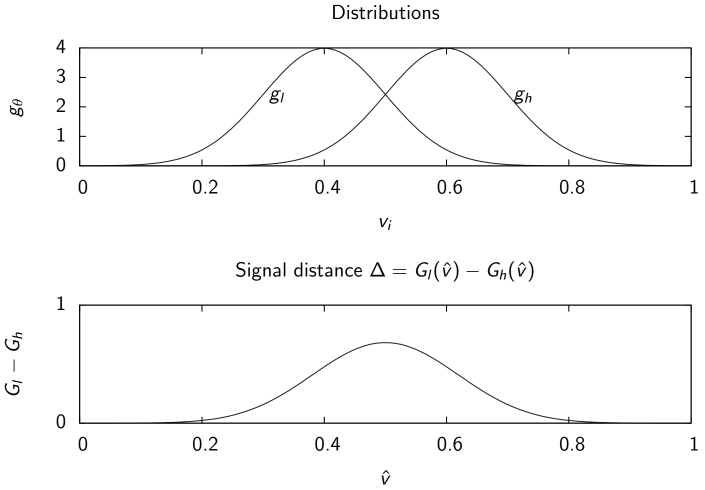

Result: more information when $\hat{v} = \frac{\mu_h + \mu_l}{2}$, i.e., $0 \ll m \ll 1$. ▸ Alt. signals

Inference from Endogenous Signal

$$m_t = \underbrace{1 - G_\theta(\hat{v}_t)}_{\text{signal}} + \underbrace{\nu_t}_{\text{noise}}$$

Example 1: Normal case with different means $\mu_h > \mu_l$

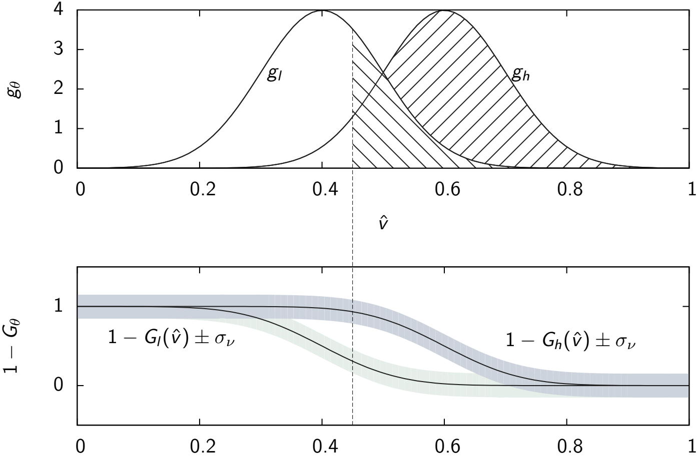

Signal vs. Noise

Example 2: Information contained in m under the **equilibrium** $\hat{v}$

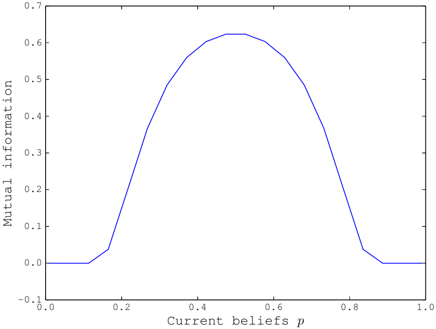

Result: in the extremes of the state-space, the endogenous signal reveals no information ▸ Parameters

Proposition 2 (Coordination traps)

Under the conditions of proposition 1,

1. *If* $(1-\beta)\,\theta_l \leq c \leq (1-\beta)\,\theta_h$*, there exists* $\underline{p} \in [0,1]$*, such that for all* $p \leq \underline{p}$*,* $\hat{v}\,(p) = b$*, i.e., nobody undertakes the project;*
2. *If* $(1-\beta)\,\theta_l + \beta \leq c \leq (1-\beta)\,\theta_h + \beta$*, there exists* $\overline{p} \in [0,1]$*, such that for all* $p \geq \overline{p}$*,* $\hat{v}\,(p) = a$*, i.e., everyone undertakes the project;*
3. *For* $p \leq \underline{p}$ *and* $p \geq \overline{p}$*,* m *contains no information about* θ*.*

Furthermore, the regions with no and full activity widen with the degree of complementarity β*:*

$$\overline{p}'\,(\beta) < 0 \text{ and } \underline{p}'\,(\beta) > 0.$$

We refer to the set $[0, \underline{p}] \cup [\bar{p}, 1]$ has the **no-learning zone.** Details

- Agents disregard their private information and all act together
- m is independent of the true state of the world

Signal vs. Noise: Role of β

Example 2: Information contained in m under the **equilibrium** $\hat{v}$

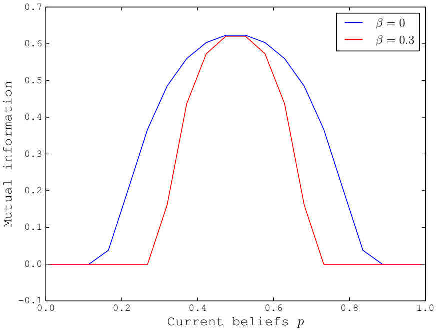

Result: the complementarity lowers informativeness and widens the no-learning zones ▸ Parameters ▸ Details

- for $p > \frac{1}{2}$, higher β implies more projects realized ($\hat{v} \to a$)
- for $p < \frac{1}{2}$, higher β implies fewer projects realized ($\hat{v} \to b$)

Complementarity and the Persistence of Recession

To summarize:

- Higher complementarity reduces informativeness of public signals in the extremes of the state space
- In the no-learning zone, agents get no information from public signal

As a result, an economy with high complementarity might

- resist well to brief shocks;
- magnify the duration of booms/recessions after a lengthier shock.

Persistence of Recession

The economy with high complementarity resists well to **brief** shocks...

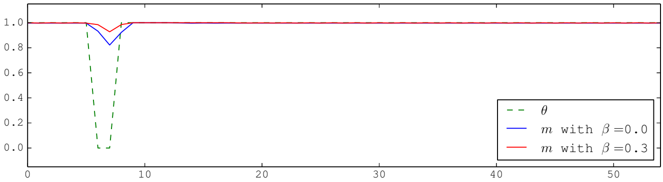

...but recovers slowly after **lengthy** shocks.

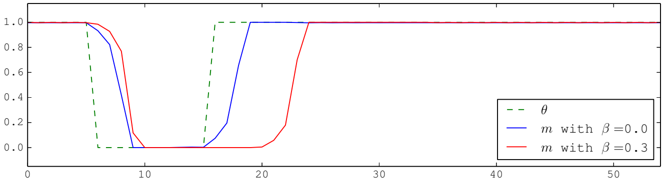

"Bubble-like" Behavior

The complementarity makes the response to ν shocks highly non-linear.

$2 \times \sigma^f$ positive shock to ν:

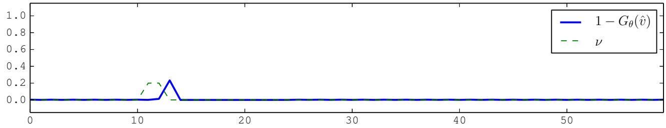

$3 \times \sigma^f$ positive shock to ν:

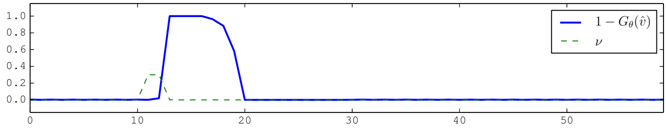

$4 \times \sigma^f$ positive shock to ν:

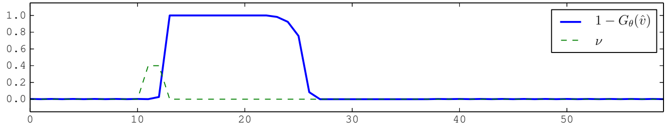

Efficiency

Agents don't internalize the impact of their decision on m.

There are two externalities:

1. **Complementarity**: a higher m increases the payoff of others
2. **Information**: m influences the amount of information revealed

We adopt the formulation of Angeletos and Pavan (2007):

- Planner cannot aggregate the information dispersed across agents
- He maximizes the ex-ante welfare of agents according to their own individual beliefs

$$V(p) = \max_{\hat{v}} E_{\theta,\nu} \left[\int_{\hat{v}}^{b} \underbrace{E_{\theta,\nu}\left[\pi_{it}(\theta,\hat{v})|\mathcal{F}_{it}\right]}_{\text{Agent } i\text{'s expected payoff}} + \gamma V(p') \middle| \mathcal{F}_t \right]$$

subject to the same law of motion for the public beliefs: $p'(p,\hat{v})$.

Dynamics in the Efficient Allocation

Response to shock in the **efficient allocation** vs **equilibrium**

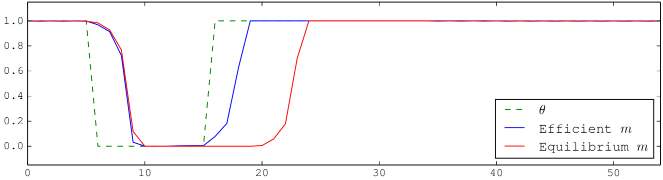

Planner's decision compared to equilibrium:

	Complementarity	Information externality
p low	more agents act	more agents act
p high	more agents act	less agents act

The planner responds to **recessions** more than to **booms**.

Conclusion

Summary

- We have built a model in which the interaction of coordination motives and endogenous information generates persistent episodes of expansions and contractions.
- Optimal government intervention reduces the length of recessions while keeping the expansions mostly unchanged.
  - Large government spending multiplier?

Extensions

- Generalized payoff function and endogenous public signal
- Intensive margin and unbounded distributions
- Long-lived agents with dynamic decision

Applications

- Unemployment fluctuations, investment dynamics, currency attacks, bank runs, asset pricing, etc.

Dynamic of Information

The public beliefs evolve according to

$$p' = \frac{P_{hh} p f_h^z(z) f(m-1+G_h(\hat{v})) + P_{lh}(1-p) f_l^z(z) f(m-1+G_l(\hat{v}))}{p f_h^z(z) f(m-1+G_h(\hat{v})) + (1-p) f_l^z(z) f(m-1+G_l(\hat{v}))}$$

Details

General Statement of Mutual Information Lemma

Lemma 4

The mutual information between θ and m is

$$I(\theta; m) = p\,(1-p)\,\Delta^2\Gamma + O\left(\Delta^3\right)$$

where $\Delta = G_l\,(\hat{v}) - G_h\,(\hat{v}) \geq 0$ *and*

$$\Gamma = \int \left[-\frac{d^2 f^\nu}{d\nu^2} + \frac{1}{2f^\nu}\left(\frac{df^\nu}{d\nu}\right)^2\right] d\nu.$$

If $F^\nu \sim \mathcal{N}(0, \sigma_\nu^2)$*, then* $\Gamma = (2\sigma_\nu^2)^{-1}$*.*

Return

Mutual Information

Definition 1

The mutual information between θ and m is

$$I(\theta; m) = H(\theta) - H(\theta|m) = \sum_{\theta \in \{\theta_L, \theta_H\}} \int_m P(\theta, m) \log\left(\frac{P(\theta, m)}{P(\theta)P(m)}\right) dm$$

where H denotes the entropy.

Return

Numerical Example

Description	Value
Low fundamental value	$\theta_L = 0$
High fundamental value	$\theta_H = 1$
Persistence of fundamental	$q = 0.99$
Cost of investment	$c = 0.5$
Time discount	$\gamma = 0.5$
Private signal in state H	$G_H \sim \mathcal{N}(0.8, 0.4)$ truncated on $[0, 1]$
Private signal in state L	$G_L \sim \mathcal{N}(0.2, 0.4)$ truncated on $[0, 1]$
Noise in public signal	$F \sim \mathcal{N}(0, 0.1)$

Return

Signal vs. Noise

Example 1.1: Truncated normals case with different variances $\sigma_h < \sigma_l$:

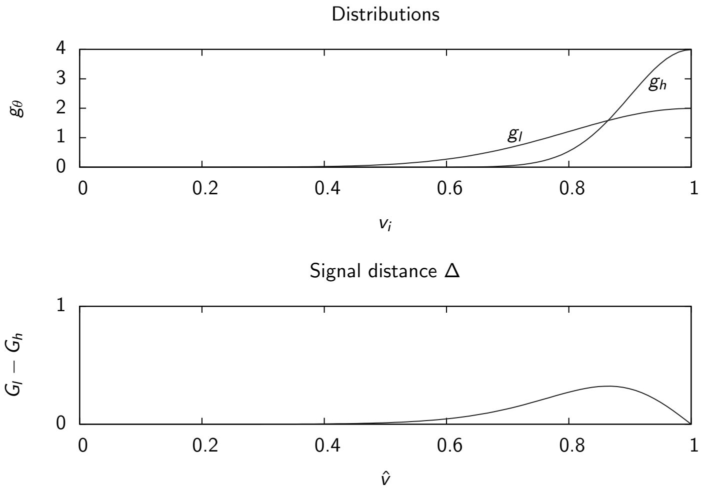

Result: informativeness of signal depends on underlying distributions

Uniqueness: Intuition

Recall the payoff function:

$$\pi\left(v_i; \hat{v}, p\right) = \underbrace{(1-\beta)\,\mathrm{E}_i\left[\theta\right]}_{\text{Fundamental}} + \underbrace{\beta \mathrm{E}_i\left[1 - G_\theta\left(\hat{v}\right)\right]}_{\text{Complementarity}} - c$$

we're looking for

$$\pi\left(\hat{v}; \hat{v}, p\right) = (1-\beta)\,\mathrm{E}\left[\theta|\hat{v}\right] + \beta \mathrm{E}\left[1 - G_\theta\left(\hat{v}\right)|\hat{v}\right] - c$$

Example: normal case with different means $\mu_h > \mu_l$

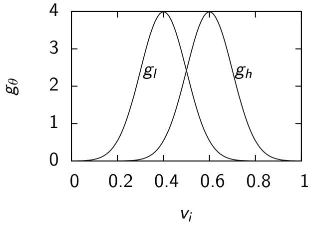

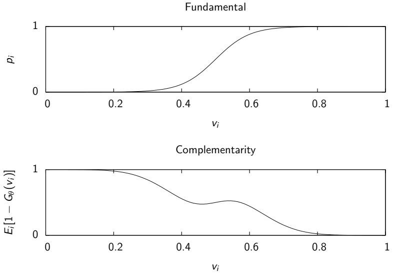

Role of complementarity β

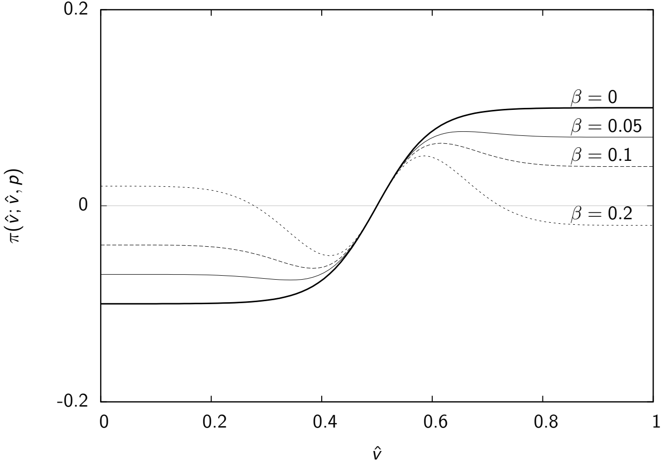

Result: Uniqueness requires upper bound on complementarity ◂ Return

Role of belief dispersion

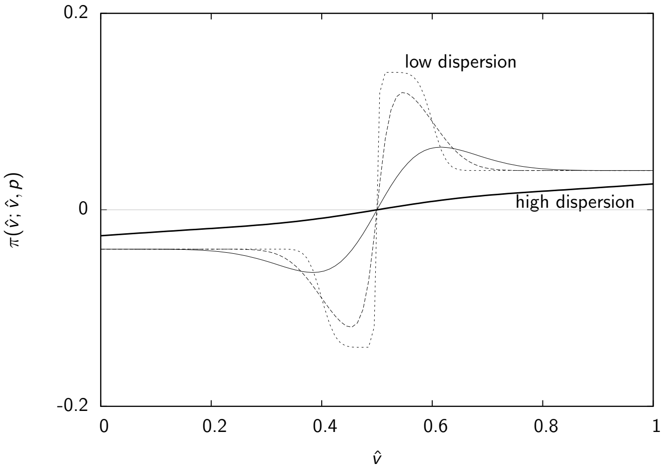

Result: Uniqueness requires enough belief dispersion ◂ Return

- Distributions g_h, g_l sufficiently dispersed
- Fundamental sufficiently volatile (P_{hl} and P_{lh} high enough)

Coordination Traps

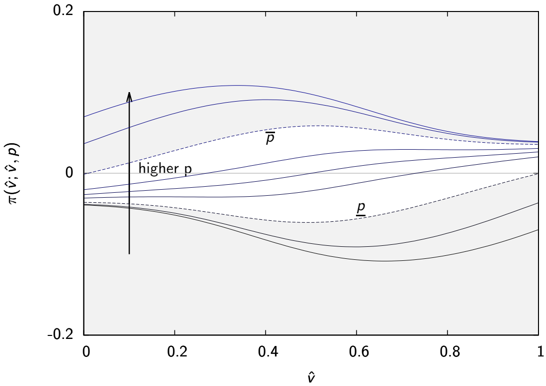

Result: endogenous channel uninformative for extreme values of p

- for $p < \underline{p}$, no project realized: $\hat{v} = b$, θ_l and θ_h are indistinguishable
  $1 - G_h(b) = 1 - G_l(b) = 0$
- for $p > \overline{p}$, all projects realized: $\hat{v} = a$, θ_l and θ_h are indistinguishable
  $1 - G_h(a) = 1 - G_l(a) = 1$

Return

Signal vs. Noise: Role of β

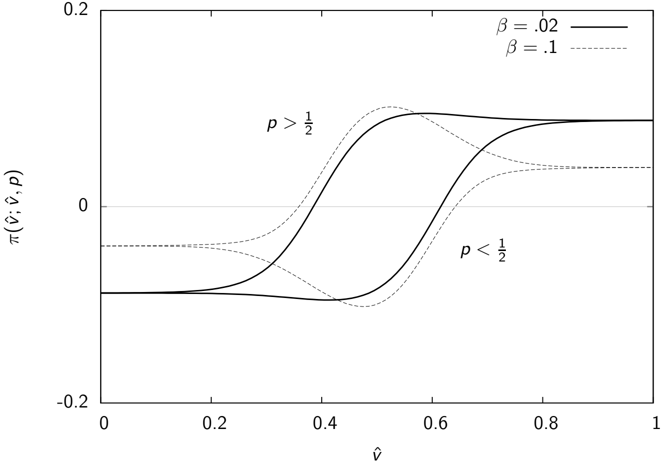

Result: high complementarity induces convergence in strategies

- for $p > \frac{1}{2}$, higher β implies more projects realized ($\hat{v} \to a$)
- for $p < \frac{1}{2}$, higher β implies fewer projects realized ($\hat{v} \to b$)