Composition of Markets with Conflicting Incentives

Alex Peysakhovich & Mikkel Plagborg-Moller

October 20, 2010

What if traders aren't myopic?

- Logic of market scoring rules relies on assumption that traders don't consider how their current report might influence outside payoffs.
- But what if traders try to game their opponents?

What if traders aren't myopic?

- Logic of market scoring rules relies on assumption that traders don't consider how their current report might influence outside payoffs.
- But what if traders try to game their opponents?
- We need a theory of manipulation incentives.

What goes wrong?

1. Outcome manipulation: Traders may be able to influence the variable they're predicting through non-market channels.
   - Ottaviani & Sørensen (2007): Corporate prediction market where employees can sabotage the outcome—moral hazard.
   - Shi, Conitzer & Guo (2009): Principal-agent framework.

What goes wrong?

1. Outcome manipulation: Traders may be able to influence the variable they're predicting through non-market channels.
   - Ottaviani & Sørensen (2007): Corporate prediction market where employees can sabotage the outcome—moral hazard.
   - Shi, Conitzer & Guo (2009): Principal-agent framework.
2. Information manipulation: Private information is your golden egg in an MSR. Participants have incentive to guard it jealously—this can produce inefficiency.

Outline

Example: Set-up (1)

A very simple scenario:

- States of the world: $\{\omega_1, \omega_2, \omega_3, \omega_4\}$ with common prior $p(\omega_i) = \frac{1}{4}$ for each i.
- Two players: Alice and Bob with type spaces $T_a = \{a, a'\}$ and $T_b = \{b, b'\}$.

Example: Set-up (1)

A very simple scenario:

- States of the world: $\{\omega_1, \omega_2, \omega_3, \omega_4\}$ with common prior $p(\omega_i) = \frac{1}{4}$ for each i.
- Two players: Alice and Bob with type spaces $T_a = \{a, a'\}$ and $T_b = \{b, b'\}$.
- Information structure:

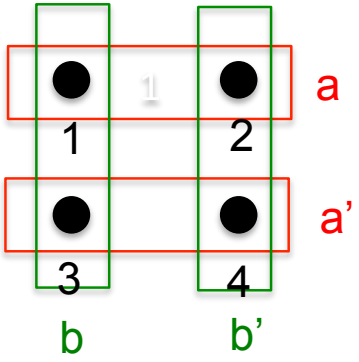

Example: Set-up (1)

A very simple scenario:

- States of the world: $\{\omega_1, \omega_2, \omega_3, \omega_4\}$ with common prior $p(\omega_i) = \frac{1}{4}$ for each i.
- Two players: Alice and Bob with type spaces $T_a = \{a, a'\}$ and $T_b = \{b, b'\}$.
- Information structure:

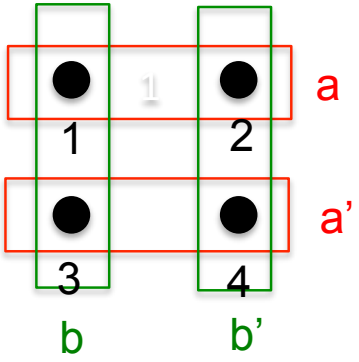

- What is type b of Bob's posterior? What happens if we know both types?

Set-up (2)

We'll call this game Ω and it works as follows:

- Nature picks a state $\omega \in \Omega$ and players receive their types.
- Bob can move the market probability distribution from p_0 to p_b.
- Alice can move the distribution from p_b to p_a. The report is public.
- The state of the world ω is revealed.
- Payments are computed.

Set-up (3)

- The payments are:
  - Bob is paid according to $\log(p_b(\omega_i)) - \log(p_0(\omega_i))$.
  - Alice is paid according to $\log(p_a(\omega_i)) - \log(p_b(\omega_i))$.
- Optimal strategy: report true p_i.

Set-up (3)

- The payments are:
  - Bob is paid according to $\log(p_b(\omega_i)) - \log(p_0(\omega_i))$.
  - Alice is paid according to $\log(p_a(\omega_i)) - \log(p_b(\omega_i))$.
- Optimal strategy: report true p_i.
- If signals are distinguishable (and they are in this example), Bob's strategy is an invertible mapping from types to actions.
- Therefore if Alice believes Bob is rational, $p_a = p(\omega_i \mid t_a, t_b)$.

Set-up (3)

- The payments are:
  - Bob is paid according to $\log(p_b(\omega_i)) - \log(p_0(\omega_i))$.
  - Alice is paid according to $\log(p_a(\omega_i)) - \log(p_b(\omega_i))$.
- Optimal strategy: report true p_i.
- If signals are distinguishable (and they are in this example), Bob's strategy is an invertible mapping from types to actions.
- Therefore if Alice believes Bob is rational, $p_a = p(\omega_i \mid t_a, t_b)$.
- Notice this is a weaker epistemic condition than equilibrium (only need rationality + 1 level of knowledge of rationality, not common knowledge).
- By a backward backward induction argument all information will always be revealed in this market.

Solution

The payments in equilibrium are:

- Bob will get $\log(\frac{1}{2}) - \log(\frac{1}{4})$.
- Alice will get $\log(1) - \log(\frac{1}{2})$.

Making it more interesting

What if Alice's type also includes information about something else? Now before we play game Ω we'll play game Γ.

- Suppose we have more states we are interested in$\{\gamma_1, \gamma_2\}$ (so 'true' state space is $\Gamma \times \Omega$).
- Alice gets type a if the state is γ_1 and a' if the state is γ_2.
- Alice is asked to publicly report her assessment p'_a on Γ (true state will be revealed later when ω is revealed).
- Alice is paid $\lambda \log(p'_a(\gamma))$ at the end of game Ω.

$2 \times$ incentive compatible $\neq$ incentive compatible

Claim

In the compound game induced by playing Γ then Ω we can set λ such that there is no PBE where Alice reports p'_a truthfully. However, there is always full revelation in Ω.

$2 \times$ incentive compatible $\neq$ incentive compatible

Claim

In the compound game induced by playing Γ then Ω we can set λ such that there is no PBE where Alice reports p'_a truthfully. However, there is always full revelation in Ω.

- Heuristic: Suppose Alice reports p'_a truthfully in equilibrium, she gets $\lambda \log(1)$ for sure. However, then $p_b(\omega_i) = 1$ for some i and Alice loses $\log(1) - \log(\frac{1}{2})$. For the game Ω the argument from before goes through.

$2 \times$ incentive compatible $\neq$ incentive compatible

Claim

In the compound game induced by playing Γ then Ω we can set λ such that there is no PBE where Alice reports p'_a truthfully. However, there is always full revelation in Ω.

- Heuristic: Suppose Alice reports p'_a truthfully in equilibrium, she gets $\lambda \log(1)$ for sure. However, then $p_b(\omega_i) = 1$ for some i and Alice loses $\log(1) - \log(\frac{1}{2})$. For the game Ω the argument from before goes through.
- Intuition: Since there is total revelation in Γ, MSR induces split of total information payoff pie in proportion to information each trader adds. Alice gains in Γ at the expense of losing some pie in Ω.

Discussion

- Questions about the example?

Dimitrov & Sami (2010): Basics

- Formalize incentive conflicts through (scoring rule) prices.
- Consider two privately informed traders: Alice and Bob.
- Before they each trade in a market V, Alice trades in a market U. Thus, Bob may extract information about Alice's private signal.

Dimitrov & Sami (2010): Basics

- Formalize incentive conflicts through (scoring rule) prices.
- Consider two privately informed traders: Alice and Bob.
- Before they each trade in a market V, Alice trades in a market U. Thus, Bob may extract information about Alice's private signal.
- Alice's tradeoff: If Bob backs out her signal, she will not be able to make a profit in market V.

Set-up (1)

- Two binary random variables that we want forecasted: $U, V \in \{0, 1\}$. Give rise to two markets.
- Private signals: Alice receives $X \in \mathcal{X}$ (realization x), Bob receives $Y \in \mathcal{Y}$ (realization y).

Set-up (2)

- Assume common knowledge of joint distribution. Priors:

$$f_0 = \Pr(U = 1), \quad g_0 = \Pr(V = 1).$$

- Probability aggregates:

$$f_x = \Pr(U = 1 \mid X = x),$$

$$f_y = \Pr(U = 1 \mid Y = y),$$

$$f_{xy} = \Pr(U = 1 \mid X = x, Y = y).$$

- Define g_x, g_y and g_{xy} similarly for V.

Set-up (2)

- Assume common knowledge of joint distribution. Priors:

$$f_0 = \Pr(U = 1), \quad g_0 = \Pr(V = 1).$$

- Probability aggregates:

$$f_x = \Pr(U = 1 \mid X = x),$$

$$f_y = \Pr(U = 1 \mid Y = y),$$

$$f_{xy} = \Pr(U = 1 \mid X = x, Y = y).$$

- Define g_x, g_y and g_{xy} similarly for V.
- Assume V-distinguishability: Conditional expectation of V given y is different for different values of y.

Market scoring rule

- Logarithmic market scoring rule: If Alice changes the predicted probability in the U market from r to q, she earns

$$\pi_U(u, r \to q) = \begin{cases} \lambda_U[\log q - \log r] & \text{if } u = 1 \\ \lambda_U[\log(1-q) - \log(1-r)] & \text{if } u = 0 \end{cases},$$

  where u is the realization of U.
- $\lambda_U > 0$ is a weight on the U-payoff.

Market scoring rule

- Logarithmic market scoring rule: If Alice changes the predicted probability in the U market from r to q, she earns

$$\pi_U(u, r \to q) = \begin{cases} \lambda_U[\log q - \log r] & \text{if } u = 1 \\ \lambda_U[\log(1-q) - \log(1-r)] & \text{if } u = 0 \end{cases},$$

  where u is the realization of U.
- $\lambda_U > 0$ is a weight on the U-payoff.
- Define similarly the V market payoff $\pi_V(v, r \to q)$ with weight λ_V.

Sequence

- Stage 1: Alice trades in the U market (*Bob observes Alice's trade*).
- Stage 2: Bob trades in the V market (*Alice observes Bob's trade*).
- Stage 3: Alice trades in the V market.

Strategies

- Solution concept: Weak PBE.

Strategies

- Solution concept: Weak PBE.
- "Backward induction:"
  - In Stage 3 Alice will back out Bob's signal (using V-distinguishability) and trade according to her honest updated prediction.

Strategies

- Solution concept: Weak PBE.
- "Backward induction:"
  - In Stage 3 Alice will back out Bob's signal (using V-distinguishability) and trade according to her honest updated prediction.
  - In Stage 2 Bob honestly trades according to his belief about Alice's signal.

Strategies

- Solution concept: Weak PBE.
- "Backward induction:"
  - In Stage 3 Alice will back out Bob's signal (using V-distinguishability) and trade according to her honest updated prediction.
  - In Stage 2 Bob honestly trades according to his belief about Alice's signal.
  - In Stage 1 Alice employs a mixed strategy $\sigma : \mathcal{X} \to \mathcal{D}([0,1])$. Assumption: $\sigma(x)$ has support in finite set $\mathcal{R}$. Let r denote a generic realized prediction.

Strategies

- Solution concept: Weak PBE.
- "Backward induction:"
  - In Stage 3 Alice will back out Bob's signal (using V-distinguishability) and trade according to her honest updated prediction.
  - In Stage 2 Bob honestly trades according to his belief about Alice's signal.
  - In Stage 1 Alice employs a mixed strategy $\sigma : \mathcal{X} \to \mathcal{D}([0,1])$. Assumption: $\sigma(x)$ has support in finite set $\mathcal{R}$. Let r denote a generic realized prediction.
- Denote Bob's Stage 2 beliefs by $\mu_r(x)$. Due to honesty in equilibrium, the authors identify his strategy with these beliefs.

Strategies

- Solution concept: Weak PBE.
- "Backward induction:"
  - In Stage 3 Alice will back out Bob's signal (using V-distinguishability) and trade according to her honest updated prediction.
  - In Stage 2 Bob honestly trades according to his belief about Alice's signal.
  - In Stage 1 Alice employs a mixed strategy $\sigma : \mathcal{X} \to \mathcal{D}([0,1])$. Assumption: $\sigma(x)$ has support in finite set $\mathcal{R}$. Let r denote a generic realized prediction.
- Denote Bob's Stage 2 beliefs by $\mu_r(x)$. Due to honesty in equilibrium, the authors identify his strategy with these beliefs.
- Strategy pair: (σ, μ).

Payoffs

- Efficiency in markets:

$$\mathrm{Eff}_U(\sigma) = \sum_{x \in \mathcal{X}} \Pr(X = x) \sum_{r \in \mathcal{R}} \sigma_x(r) \pi_U(f_x, f_0 \to r),$$

$$\mathrm{Eff}_V = \sum_{x \in \mathcal{X}} \sum_{y \in \mathcal{Y}} \Pr(X = x, Y = y) \pi_V(g_{xy}, g_0 \to g_{xy}),$$

where e.g.
$\pi_U(f_x, f_0 \to r) = f_x \pi_U(1, f_0 \to r) + (1 - f_x) \pi_U(0, f_0 \to r).$

Payoffs

- Efficiency in markets:

$$\mathrm{Eff}_U(\sigma) = \sum_{x \in \mathcal{X}} \Pr(X = x) \sum_{r \in \mathcal{R}} \sigma_x(r) \pi_U(f_x, f_0 \to r),$$

$$\mathrm{Eff}_V = \sum_{x \in \mathcal{X}} \sum_{y \in \mathcal{Y}} \Pr(X = x, Y = y) \pi_V(g_{xy}, g_0 \to g_{xy}),$$

where e.g.
$\pi_U(f_x, f_0 \to r) = f_x \pi_U(1, f_0 \to r) + (1 - f_x) \pi_U(0, f_0 \to r)$.

- If Alice learns Bob's μ in Stage 3, her payoff can be written

$$\pi_A(\sigma, \mu) = \mathrm{Eff}_U(\sigma) + \mathrm{Eff}_V - \pi_B(\sigma, \mu).$$

- Looks almost like constant-sum game!

Minimax (1)

- Game isn't exactly constant-sum: Total payoff depends on how far Alice moves market U.
- Constant-sum games may be analyzed using minimax strategies (goes back to birth of game theory).
- Turns out we can use this concept for our purposes.

Minimax (2)

Definition

*A **maximin pair** (σ, μ) solves*

$$\sigma \in \operatorname*{argmax}_{\sigma'} \min_{\mu'} \pi_A(\sigma', \mu'),$$

$$\mu \in \operatorname*{argmin}_{\mu'} \pi_A(\sigma, \mu').$$

Interpretation: Alice maximizes over worst possible payoffs; Bob minimizes given this.

Minimax (3)

Definition

*A **minimax pair** (σ, μ) solves*

$$\mu \in \operatorname*{argmin}_{\mu'} \max_{\sigma'} \pi_A(\sigma', \mu'),$$

$$\sigma \in \operatorname*{argmax}_{\sigma'} \pi_A(\sigma', \mu).$$

Interpretation: Bob minimizes over Alice's best payoffs; Alice maximizes given this.

Main results (1)

Lemma (1)

The minimax and maximin values of π_A coincide, and there exists a pair (σ, μ) for which this holds.

Main results (1)

Lemma (1)

The minimax and maximin values of π_A coincide, and there exists a pair (σ, μ) for which this holds.

Proof outline: Show that π_A is concave in σ and convex in μ—follows readily from properties of the logarithmic function. Use result from the literature.

Main results (2)

Lemma (3)

Any two minimax strategy pairs (σ, μ) and (σ', μ') lead to the same payoff for every report by Alice and Bob.

Main results (2)

Lemma (3)

Any two minimax strategy pairs (σ, μ) and (σ', μ') lead to the same payoff for every report by Alice and Bob.

Proof outline: Check that (σ, μ') is also minimax. Consider two different convex combinations $(0.6\sigma + 0.4\sigma', \mu')$ and $(0.4\sigma + 0.6\sigma', \mu')$ and suppose they lead to different Eff_U.

Main results (2)

Lemma (3)

Any two minimax strategy pairs (σ, μ) and (σ', μ') lead to the same payoff for every report by Alice and Bob.

Proof outline: Check that (σ, μ') is also minimax. Consider two different convex combinations $(0.6\sigma + 0.4\sigma', \mu')$ and $(0.4\sigma + 0.6\sigma', \mu')$ and suppose they lead to different Eff_U. These pairs are minimax by linearity of π_A, so they lead to the same total payoff for Alice.

Main results (2)

Lemma (3)

Any two minimax strategy pairs (σ, μ) and (σ', μ') lead to the same payoff for every report by Alice and Bob.

Proof outline: Check that (σ, μ') is also minimax. Consider two different convex combinations $(0.6\sigma + 0.4\sigma', \mu')$ and $(0.4\sigma + 0.6\sigma', \mu')$ and suppose they lead to different Eff_U.
These pairs are minimax by linearity of π_A, so they lead to the same total payoff for Alice.
Then Bob's payoffs must be different, but since the support of Alice's Stage 1 strategy is the same, this requires the posterior probability of x to change.

Main results (2)

Lemma (3)

Any two minimax strategy pairs (σ, μ) and (σ', μ') lead to the same payoff for every report by Alice and Bob.

Proof outline: Check that (σ, μ') is also minimax. Consider two different convex combinations $(0.6\sigma + 0.4\sigma', \mu')$ and $(0.4\sigma + 0.6\sigma', \mu')$ and suppose they lead to different Eff_U.
These pairs are minimax by linearity of π_A, so they lead to the same total payoff for Alice.
Then Bob's payoffs must be different, but since the support of Alice's Stage 1 strategy is the same, this requires the posterior probability of x to change.
This is impossible since Bob's beliefs must equal the posterior for any minimax pair (Lemma 2).

Main results (3)

Theorem (5)

There exists a minimax pair such that Bob's beliefs are consistent (Lemma 4). Such a pair represents a Weak PBE.

Main results (3)

Theorem (5)

There exists a minimax pair such that Bob's beliefs are consistent (Lemma 4). Such a pair represents a Weak PBE.

Proof outline: Follows readily from optimality conditions of a minimax pair.

Main results (3)

Theorem (5)

There exists a minimax pair such that Bob's beliefs are consistent (Lemma 4). Such a pair represents a Weak PBE.

Proof outline: Follows readily from optimality conditions of a minimax pair.

Theorem (6)

Let (σ^*, μ^*) be a minimax pair. In any Weak PBE the payoffs to Alice and Bob are $\pi_A(\sigma^*, \mu^*)$ and $\pi_B(\sigma^*, \mu^*)$, respectively.

Main results (3)

Theorem (5)

There exists a minimax pair such that Bob's beliefs are consistent (Lemma 4). Such a pair represents a Weak PBE.

Proof outline: Follows readily from optimality conditions of a minimax pair.

Theorem (6)

Let (σ^*, μ^*) be a minimax pair. In any Weak PBE the payoffs to Alice and Bob are $\pi_A(\sigma^*, \mu^*)$ and $\pi_B(\sigma^*, \mu^*)$, respectively.

Proof outline: Easy to see that any Weak PBE strategy pair must be minimax, so we can use Lemma 3.

Main results (4)

- Take-home message: The payoffs in any Weak PBE can be computed by finding the minimax value of Alice's payoff function π_A.

Main results (4)

- Take-home message: The payoffs in any Weak PBE can be computed by finding the minimax value of Alice's payoff function π_A.
- Section 5.2 illustrates how one may compute the minimax strategies in practice. The problem is nonlinear.

Noisy channel (1)

$$\begin{aligned}
\mathrm{Eff}_U(\sigma) &= \sum_{x\in\mathcal{X}} \Pr(X=x) \sum_{r'\in\mathcal{R}} \sigma_x(r')\pi_U(f_x, f_0 \to r') \\
&= \lambda_U \sum_{x\in\mathcal{X}} \sum_{r'\in\mathcal{R}} \sum_{u\in\{0,1\}} \Big\{ \Pr(X=x)\Pr(r=r' \mid X=x) \\
&\qquad \times \Pr(U=u \mid X=x) \log \frac{\Pr(U=u \mid r=r')}{\Pr(U=u)} \Big\} \\
&= \lambda_U \sum_{r'} \sum_{u} \Pr(r=r', U=u) \log \frac{\Pr(U=u \mid r=r')}{\Pr(U=u)}
\end{aligned}$$

- Note: $\Pr(r' \mid x)\Pr(u \mid x) = \Pr(r', u \mid x)$.

Noisy channel (2)

- Due to LMSR, Eff_U can directly be thought of as the mutual information between U and the garbled report r.
- Also: Consider a modified game where Alice's signal structure is such that she receives signal $\sigma(x)$ instead. Since (σ, μ) was optimal in the original game, it must now be optimal for Alice to report truthfully in U market.

Limiting ratio (1)

- The greater the ratio λ_V/λ_U, the more does Alice want to conceal her true signal in market U.

Limiting ratio (1)

- The greater the ratio λ_V/λ_U, the more does Alice want to conceal her true signal in market U.
- Fixing all other primitives than the payoff weights, define the limiting ratio to be the maximum ratio λ_V/λ_U such that Alice has no profitable deviation from a truthtelling equilibrium.
- Note: The limiting ratio can be $+\infty$.

Limiting ratio (1)

- The greater the ratio λ_V/λ_U, the more does Alice want to conceal her true signal in market U.
- Fixing all other primitives than the payoff weights, define the limiting ratio to be the maximum ratio λ_V/λ_U such that Alice has no profitable deviation from a truthtelling equilibrium.
- Note: The limiting ratio can be $+\infty$.
- Fixing λ_U, if truthtelling is optimal for λ_V^* it is also optimal for all $\lambda_V \leq \lambda_V^*$, since such a reduction leaves the honest payoff constant while reducing dishonest gains.

Limiting ratio (2)

- The limiting ratio can be interpreted as a measure of alignedness of incentives for honesty.
- Theorem 7 characterizes the limiting ratio of combinations of general forecast values with discrete support in $[0, 1]$.

Discussion

- Is this about two markets?
- Real-life examples of such information tradeoffs?

Chen et al. (2009)

- The problem here is that market participants want to hold on to their information.
- We can think about the game as one market with the following rules:
  - Alice moves market probability from p_0 to p_a.
  - Bob sees Alice's move and moves the market probability from p_a to p_b.

Chen et al. (2009)

- The problem here is that market participants want to hold on to their information.
- We can think about the game as one market with the following rules:
  - Alice moves market probability from p_0 to p_a.
  - Bob sees Alice's move and moves the market probability from p_a to p_b.
- Timing: Would I rather be Bob or Alice in this game?

Chen et al. (2)

This depends on the signal structure. Suppose we have states of the world Ω and type space $T_a \times T_b$, symmetric distinguishable signal distributions and we use a LMSR to elicit information.

Chen et al. (2)

This depends on the signal structure. Suppose we have states of the world Ω and type space $T_a \times T_b$, symmetric distinguishable signal distributions and we use a LMSR to elicit information.

Definition

Say that the signal structure is ***conditionally independent*** *if* $p(t_a, t_b \mid \omega) = p(t_a \mid \omega)p(t_b \mid \omega)$.

Examples: Most information structures we've studied so far.

Chen et al. (2)

This depends on the signal structure. Suppose we have states of the world Ω and type space $T_a \times T_b$, symmetric distinguishable signal distributions and we use a LMSR to elicit information.

Definition

Say that the signal structure is ***conditionally independent*** *if* $p(t_a, t_b \mid \omega) = p(t_a \mid \omega)p(t_b \mid \omega)$.

Examples: Most information structures we've studied so far.

Definition

Say the signal structure is ***unconditionally independent*** *if* $p(t_a, t_b) = p(t_a)p(t_b)$.

Examples: Elections, my type is D or R and outcome is determined based on majority vote.

Chen et al. (3)

Theorem

If signals are conditionally (unconditionally) independent then $\pi^*_{Alice} \geq (\leq) \pi^*_{Bob}$.

Chen et al. (3)

Theorem

If signals are conditionally (unconditionally) independent then $\pi^*_{Alice} \geq (\leq) \pi^*_{Bob}$.

Intuition: All possible information is revealed in the end, so the question is how much of the information pie each person gets.

- If signals are conditionally independent, I want to go **first**: My signal is information substitute for part of your signal.
- If signals are unconditionally independent, I want to go **second**: You learn nothing about my signal from your signal, but I learn something about the state of the world from your report.

Chen et al. (3)

Theorem

If signals are conditionally (unconditionally) independent then $\pi^*_{Alice} \geq (\leq) \pi^*_{Bob}$.

Intuition: All possible information is revealed in the end, so the question is how much of the information pie each person gets.

- If signals are conditionally independent, I want to go **first**: My signal is information substitute for part of your signal.
- If signals are unconditionally independent, I want to go **second**: You learn nothing about my signal from your signal, but I learn something about the state of the world from your report.

This means in infinitely repeated LMSR games we might not get information revelation in finite time.

What's next?

- What are the restrictive assumptions?

What's next?

- What are the restrictive assumptions?
  - Common prior
  - Exogenous timing
  - Exogenous entry

What's next?

- What are the restrictive assumptions?
  - Common prior
  - Exogenous timing
  - Exogenous entry
- Future research?

What's next?

- What are the restrictive assumptions?
  - Common prior
  - Exogenous timing
  - Exogenous entry
- Future research?
  - Test qualitative predictions/assumptions
  - How to get closer to real world?