INCENTIVE-COMPATIBLE EXPERIMENTAL DESIGN

Panos Toulis

June 18, 2015

Harvard University, Department of Statistics
Joint work with David C. Parkes (SEAS), Elery Pfeffer (SEAS), James Zou (Microsoft)

EXPERIMENTAL TREATMENTS AS AGENTS

Marketing Firm 2

Marketing Firm 1

OUTCOMES WITHOUT INCENTIVES

Marketing Firm 1

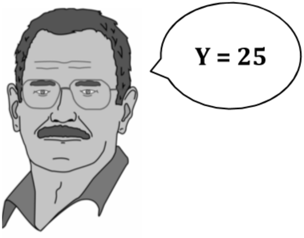

Marketing Firm 2

AMERICAN EXPRESS

INTRODUCING INCENTIVES..

Marketing Firm 1

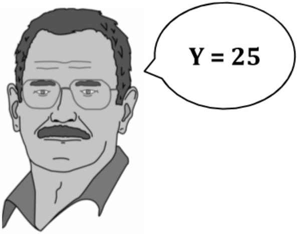

Marketing Firm 2

INCENTIVES...CAUSE DIFFERENT EXPERIMENTAL OUTCOMES

Marketing Firm 1

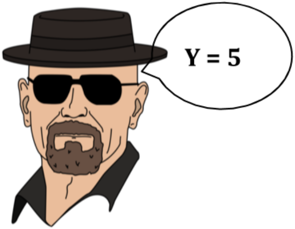

Marketing Firm 2

CHALLENGES

The high-level statistical inference challenge:

- Goal is to estimate performance **without** competition.
- Experiment has data about performance **with** competition.

The specific challenges:

- **Risk vs Return**: A weaker agent may increase its probability of winning by deviating to a“higher risk action”.
- **Strategic Interference**: Action of an agent may affect outcome on units assigned to another agent.

- Formalize *incentive-compatible experimental design*.
- Method to design IC experiments under no interference using *variance stabilizing transformations*.
- Method to design IC experiments under interference using *augmented designs*.

NO INTERFERENCE

Experiment design $\mathcal{D} = (\psi, \phi)$:

(1) Each **experimental unit** assigned to one agent according to ψ.

(2) Agent i picks **treatment version** α_i (action) applied to its units.

(3) **Outcome** Y_{ui} for unit u assigned to agent i.

- Assume, $Y_{ui} \sim [\mu(\alpha_i), \sigma^2(\alpha_i)]$; $Y = \{Y_{ui}\}$.

(4) **Score** of agent i, $\phi_i(Y) \in \mathbb{R}$.

- Assume, $\phi_i(Y) \sim [\mu_\phi(\alpha_i), \sigma_\phi^2(\alpha_i)]$.

(5) **Winner** $\hat{\tau} = \arg\max_i\{\phi_i(Y)\}$. (winner-take-all game.)

INCENTIVE-COMPATIBLE EXPERIMENTAL DESIGN: SETUP

Experiment design $\mathcal{D} = (\psi, \phi)$:

(1) Each **experimental unit** assigned to one agent according to ψ.

(2) Agent i picks **treatment version** α_i (action) applied to its units.

(3) **Outcome** Y_{ui} for unit u assigned to agent i.

- Assume, $Y_{ui} \sim [\mu(\alpha_i), \sigma^2(\alpha_i)]$; $Y = \{Y_{ui}\}$.

(4) **Score** of agent i, $\phi_i(Y) \in \mathbb{R}$.

- Assume, $\phi_i(Y) \sim [\mu_\phi(\alpha_i), \sigma^2_\phi(\alpha_i)]$.

(5) **Winner** $\hat{\tau} = \arg\max_i \{\phi_i(Y)\}$. (winner-take-all game.)

The experimental goal (what the experimenter wants to learn.)

(1) $\alpha_i^\star = \arg\max_{\alpha_i} \mu(\alpha_i)$. (natural action, $\mu(\alpha_i^\star)$ = agent quality.)

(2) $\tau = \arg\max_i \{\mu(\alpha_i^\star)\}$. (highest-quality agent.)

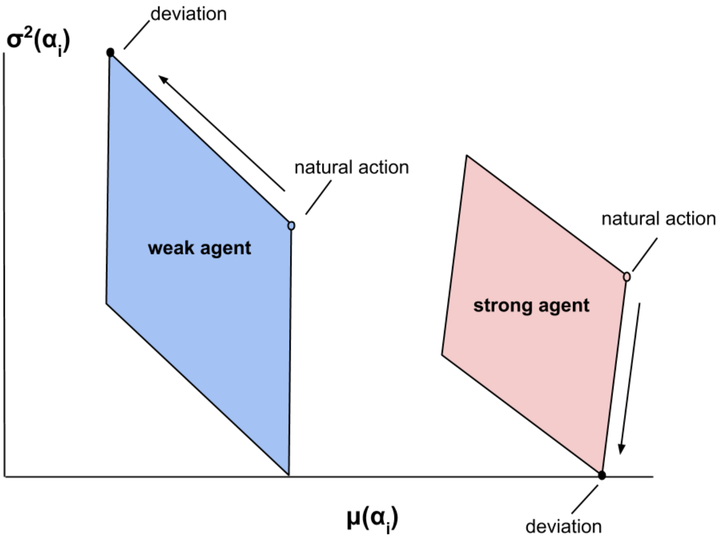
deviation
σ²(αᵢ)
natural action
weak agent
natural action
strong agent
μ(αᵢ)
deviation

EXPERIMENT DESIGN $D = (\psi, \phi)$

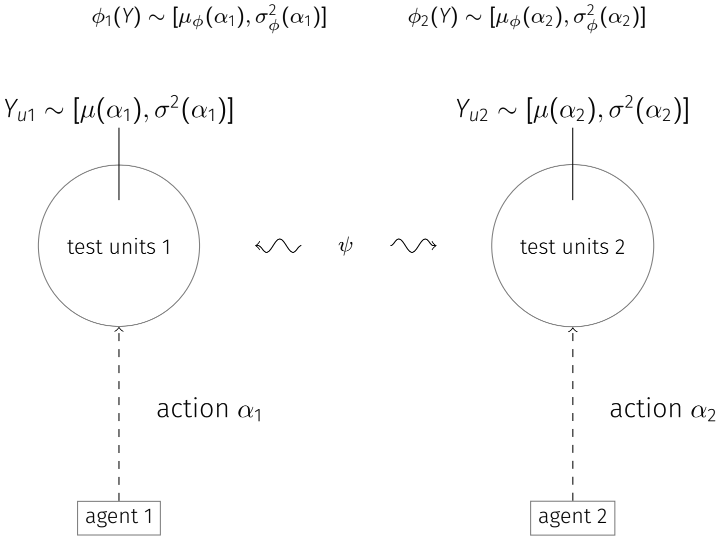

An experiment design $\mathcal{D}$ is *incentive-compatible* if the natural action $\alpha_i^\star$ is dominant strategy for each agent *i*, i.e.,

$$\arg\max_{\alpha_i} \{P_{\text{win},i}(\alpha_i, \boldsymbol{\alpha}_{-i} | \mathcal{D})\} = \alpha_i^\star,$$

for every agent *i*, actions $\boldsymbol{\alpha}_{-i}$.

(1) Start with score $\phi_i(Y) = \overline{Y_{.i}}$. (other options possible, e.g., median.)

(2) Derive **probability of win** for agent i:

$$P_i(\boldsymbol{\alpha}|\mathcal{D}) \stackrel{\text{def}}{=} \Pr(i = \arg\max_j \phi_j(Y)) = \Phi\left(\sqrt{k}\frac{\mu_\phi(\alpha_1) - \mu_\phi(\alpha_2)}{\sqrt{\sigma^2_\phi(\alpha_1) + \sigma^2_\phi(\alpha_2)}}\right).$$

(3) Check if P_i and $\mu(\alpha_i)$ are **both** maximized at natural action $\alpha_i^\star$.

(4) If not, think of a different scoring function ϕ. Goto 2.

High-level **ideas and tools**:

- For (2): Use central limit theorem; pivotal statistics, etc.
- For (3): Calculus.
- For (4): Variance-stabilizing ϕ, such that $\sigma^2_\phi(\alpha_i) \propto \text{const}$.

EXAMPLE (NEGATIVE)

- Assume score $\phi_i(Y) = \overline{Y_{.i}}$. Then, $\mu_\phi(\alpha_i) = \mu(\alpha_i), \sigma^2_\phi(\alpha_i) = \sigma^2(\alpha_i)/k$.
- Probability of win (agent 1):

$$P_1(\boldsymbol{\alpha}|\mathcal{D}) = \Phi(\sqrt{k}\frac{\mu(\alpha_1) - \mu(\alpha_2)}{\sqrt{\sigma^2(\alpha_1) + \sigma^2(\alpha_2)}}).$$

EXAMPLE (NEGATIVE)

- Assume score $\phi_i(Y) = \overline{Y_{.i}}$. Then, $\mu_\phi(\alpha_i) = \mu(\alpha_i), \sigma^2_\phi(\alpha_i) = \sigma^2(\alpha_i)/k$.
- Probability of win (agent 1):

$$P_1(\boldsymbol{\alpha}|\mathcal{D}) = \Phi(\sqrt{k}\frac{\mu(\alpha_1) - \mu(\alpha_2)}{\sqrt{\sigma^2(\alpha_1) + \sigma^2(\alpha_2)}}).$$

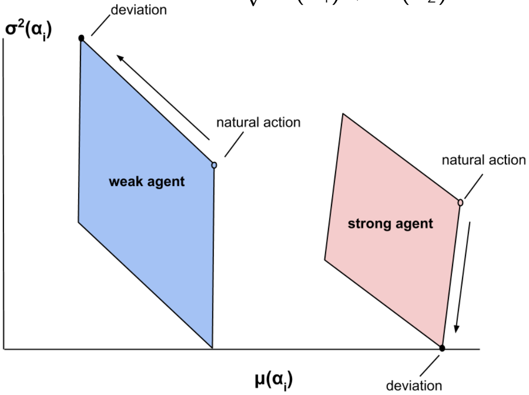

EXAMPLE (POSITIVE): CHARACTERIZATION

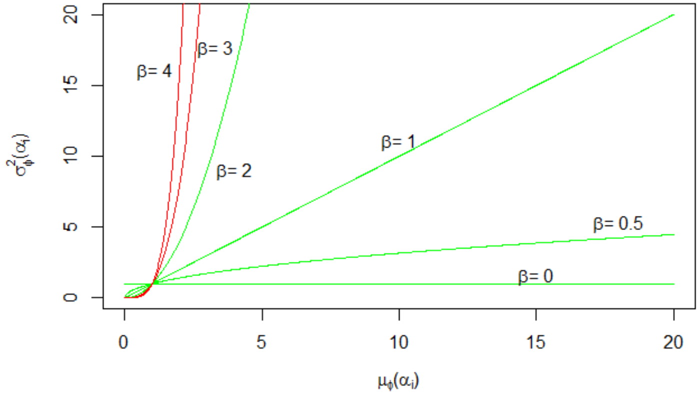

Figure 2: Family of models with $\sigma^2_\phi(\alpha_i) = (\mu_\phi(\alpha_i))^\beta$. When $\beta \leq 2$ designs are incentive-compatible with sample-mean score. When $\beta > 2$ a **transformation** of data is necessary to produce incentive-compatible design.

Figure 3: Assume $\sigma^2_\phi(\alpha_i) = (\mu_\phi(\alpha_i))^4$; fix $\mu(\alpha_2) = 2.0$. **solid line**: prob. 1 wins using $\phi =$sample mean. **dotted line**: prob. 1 wins using $\phi = -1/$sample mean; in this case, $\sigma^2_\phi(\alpha_i) = \text{const.}!$

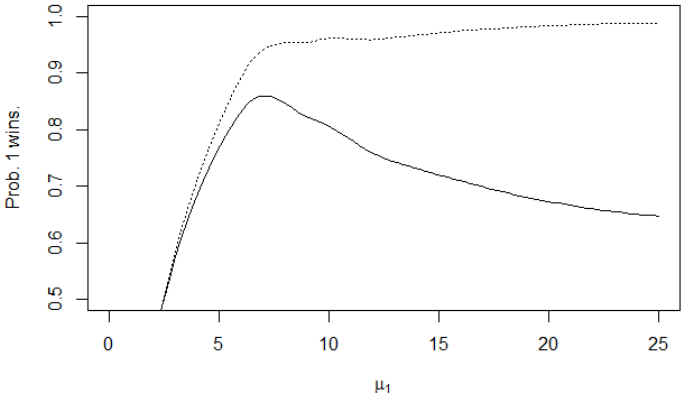

- Variance stabilizing transformation very useful in statistical hypothesis testing. (variance of transformed statistic does not depend on unknown parameter.)
- Theorem 3.1 (in paper) shows how such transformations can be leveraged to remove return-risk tradeoffs (Challenge #1) and produce IC designs.

INTERFERENCE

STRATEGIC INTERFERENCE

$$\phi_1(Y) \sim [\mu_\phi(\alpha_1 + \gamma\alpha'_2), \sigma^2_\phi(\alpha_1 + \gamma\alpha'_2)] \qquad \phi_2(Y) \sim [\mu_\phi(\alpha_2 + \gamma\alpha'_1), \sigma^2_\phi(\alpha_2 + \gamma\alpha'_1)]$$

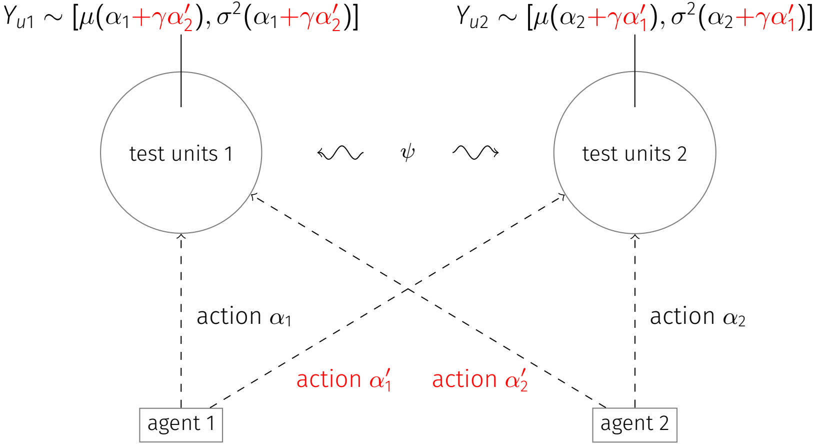

(1) **Augmentation**: extend design to gather *more relevant data.*
(2) **Estimation**: use interference model to estimate agent actions, $\hat{\alpha}_i$.
(3) **Reduction**: use $\hat{\alpha}_i$ and no-interference theory to design score ϕ.

(1) **Augmentation**: extend design to gather *more relevant data*.
(2) **Estimation**: use interference model to estimate agent actions, $\hat{\alpha}_i$.
(3) **Reduction**: use $\hat{\alpha}_i$ and no-interference theory to design score ϕ.

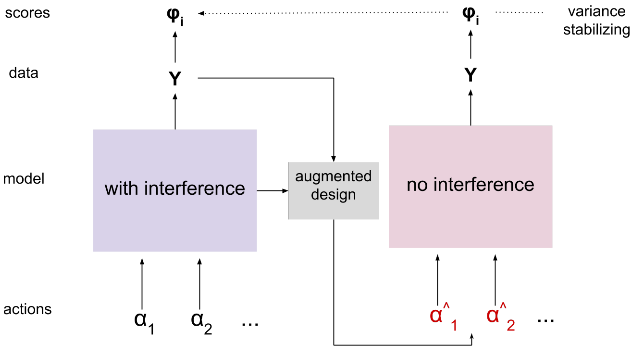

APPROACH TO IC DESIGN WITH INTERFERENCE (EXAMPLE)

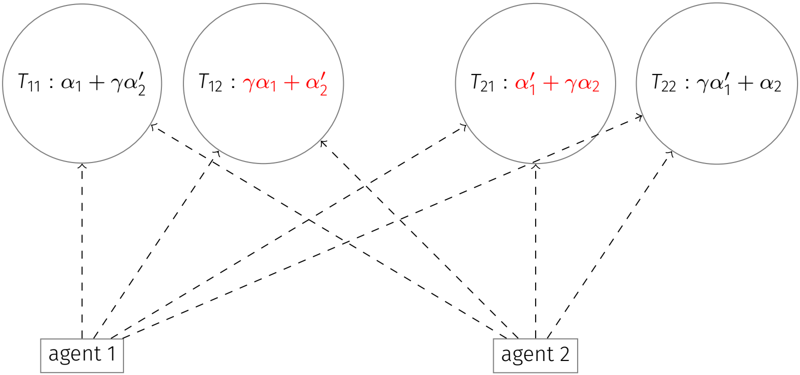

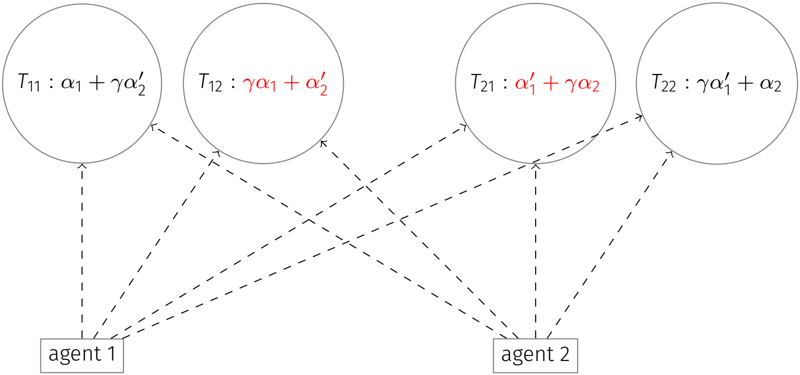

Agent actions α_i, α_i' can now be identified if $\gamma \neq 1$.

CONCLUSION

- Game theory with experimental design; treatments are strategic agents and can interfere with each other.
- Goal of **incentive-compatible experimental design** is to make agents act in the experiment as if there is no competition.

Contributions:

- Variance-stabilizing transformations can alleviate risk-return tradeoffs. (assumes a parametric model on outcomes.)
- Augmented designs can alleviate strategic interference. (assumes a model of interference.)

- Game theory with experimental design; treatments are strategic agents and can interfere with each other.
- Goal of **incentive-compatible experimental design** is to make agents act in the experiment as if there is no competition.

Contributions:

- Variance-stabilizing transformations can alleviate risk-return tradeoffs. (assumes a parametric model on outcomes.)
- Augmented designs can alleviate strategic interference. (assumes a model of interference.)

Future work:

- Incentives of more elaborate designs (e.g., latin squares, factorial).
- (Many) more applications (e.g., industrial settings, A/B testing).