A/B Contracts

George Georgiadis Michael Powell

Northwestern Kellogg

In a Nutshell

- Firms understand that there is a lot to learn from experimentation
    - *e.g.*, firms use it for product design, pricing, advertising, etc...
- A crucial area in managing a firm is designing compensation structures and how people should be rewarded for outcomes
- This question has largely evaded trends in data-driven decision-making
- We show that under mild assumptions about the way people respond to incentives and value rewards, simple experimentation coupled with theoretical insights can lead a long way towards optimal contracting.

In a Nutshell

- Firms understand that there is a lot to learn from experimentation
  - *e.g.*, firms use it for product design, pricing, advertising, etc...
- A crucial area in managing a firm is designing compensation structures and how people should be rewarded for outcomes
- This question has largely evaded trends in data-driven decision-making
- We show that under mild assumptions about the way people respond to incentives and value rewards, simple experimentation coupled with theoretical insights can lead a long way towards optimal contracting.

How to Improve upon an Incentive Contract?

- Imagine that you run a company that sells kitchen knife sets, you hire teenagers every summer to sell door-to-door, and you pay them a simple piece rate
- Your expected profits

$$\Pi = (m - \alpha)a$$

where m is profit margin, α is your piece rate, and a are mean sales

- You want to know whether and how to change your piece rate
- If you marginally increase α, then your profits change by

$$\frac{d\Pi}{d\alpha} = (m - \alpha)\frac{da}{d\alpha} - a\,.$$

So you want to know whether $d\Pi/d\alpha \gtrless 0$

How to Improve upon an Incentive Contract?

- Imagine that you run a company that sells kitchen knife sets,
  you hire teenagers every summer to sell door-to-door, and
  you pay them a simple piece rate
- Your expected profits
  $$\Pi = (m - \alpha)a$$
  where m is profit margin, α is your piece rate, and a are mean sales
- You want to know whether and how to change your piece rate
- If you marginally increase α, then your profits change by
  $$\frac{d\Pi}{d\alpha} = (m - \alpha)\frac{da}{d\alpha} - a\,.$$
  So you want to know whether $d\Pi/d\alpha \gtrless 0$

Experiments Reveal Behavioral Responses

- To answer this question, must know the behavioral response, $da/d\alpha$
- Given just observational data, need to know *production environment*

  i.e., employee's effort costs, mapping from effort to sales, etc...
- You do not need this knowledge if you run an A/B test:
  - Split your salespeople into a treatment and control group,
  - Perturb the piece rate for treatment group, and collect sales data
- You can use this data to estimate $da/d\alpha$, and determine whether you should increase or decrease your piece rate.

Experiments Reveal Behavioral Responses

- To answer this question, must know the behavioral response, $da/d\alpha$
- Given just observational data, need to know *production environment* *i.e.*, employee's effort costs, mapping from effort to sales, etc...
- You do not need this knowledge if you run an A/B test:
  - Split your salespeople into a treatment and control group,
  - Perturb the piece rate for treatment group, and collect sales data
- You can use this data to estimate $da/d\alpha$, and determine whether you should increase or decrease your piece rate.

Two Main Theoretical Issues to Address

This paper: How to improve upon a given contract?

- What does the principal need to know & how to use information

1. It restricted attention to linear contracts
   - A nonlinear contract can be modified in a continuum of ways
   - Need to know how productivity responds to every possible modification
   - **Key lemma:** A single A/B test together with an assumption about the agent's preferences for money provides all the needed information.
2. It asked a *local* question
   - In practice, one is interested in non-local changes to the contract
   - We provide conditions so that a single A/B test suffices to extrapolate, and determine how to optimally adjust a given contract (non-locally).

Two Main Theoretical Issues to Address

This paper: How to improve upon a given contract?

- What does the principal need to know & how to use information

1. It restricted attention to linear contracts
   - A nonlinear contract can be modified in a continuum of ways
   - Need to know how productivity responds to every possible modification
   - **Key lemma:** A single A/B test together with an assumption about the agent's preferences for money provides all the needed information.
2. It asked a *local* question
   - In practice, one is interested in non-local changes to the contract
   - We provide conditions so that a single A/B test suffices to extrapolate, and determine how to optimally adjust a given contract (non-locally).

Two Main Theoretical Issues to Address

This paper: How to improve upon a given contract?

- What does the principal need to know & how to use information

1. It restricted attention to linear contracts
   - A nonlinear contract can be modified in a continuum of ways
   - Need to know how productivity responds to every possible modification
   - **Key lemma:** A single A/B test together with an assumption about the agent's preferences for money provides all the needed information.
2. It asked a *local* question
   - In practice, one is interested in non-local changes to the contract
   - We provide conditions so that a single A/B test suffices to extrapolate, and determine how to optimally adjust a given contract (non-locally).

Two Empirical Exercises

- Evaluate methodology using dataset from DellaVigna & Pope (2017)
  - Real-effort experiment with several different incentive treatments

I. Test our model's ability to predict out-of-sample performance

- For each pair of treatments, take this pair to be our A/B test, and use the model to predict performance in the remaining treatments
- Correlation between predicted and actual performance > 0.9, and mean APE $< 2\%$ (performance varies 18% across treatments)

II. Assess performance of adjusted contract generated by our procedure

- Use all treatments to construct a benchmark
- Use data from each A/B test to compute *test-optimal* contract
- On average, this contract attains $> 2/3$ of the profit gap between the status quo and the *benchmark-optimal* contract.

Two Empirical Exercises

- Evaluate methodology using dataset from DellaVigna & Pope (2017)
  - Real-effort experiment with several different incentive treatments

I. Test our model's ability to predict out-of-sample performance
  - For each pair of treatments, take this pair to be our A/B test, and use the model to predict performance in the remaining treatments
  - Correlation between predicted and actual performance > 0.9, and mean APE $< 2\%$ (performance varies 18% across treatments)

II. Assess performance of adjusted contract generated by our procedure
  - Use all treatments to construct a benchmark
  - Use data from each A/B test to compute *test-optimal* contract
  - On average, this contract attains $> 2/3$ of the profit gap between the status quo and the *benchmark-optimal* contract.

Two Empirical Exercises

- Evaluate methodology using dataset from DellaVigna & Pope (2017)
  - Real-effort experiment with several different incentive treatments

I. Test our model's ability to predict out-of-sample performance
  - For each pair of treatments, take this pair to be our A/B test, and use the model to predict performance in the remaining treatments
  - Correlation between predicted and actual performance > 0.9, and mean APE $< 2\%$ (performance varies 18% across treatments)

II. Assess performance of adjusted contract generated by our procedure
  - Use all treatments to construct a benchmark
  - Use data from each A/B test to compute *test-optimal* contract
  - On average, this contract attains $> 2/3$ of the profit gap between the status quo and the *benchmark-optimal* contract.

Related Literature

- Agency problems — Theory*:*
  - Mirrlees (1976), Holmström (1979), ...
  - Carroll (2015), Gottlieb & Moreira (2017), Chade & Swinkels (2019), ...
- Agency problems — Empirics:
  - Lazear (2000), Shearer (2004), Fehr & Goette (2007), Guiteras & Jack (2018), Balbuzanov et al. (2017), Hong et al. (2018), ...
  - Prendergast (2014), d'Haultfoeuille & Fevrier (2020), ...
- Sufficient statistics:
  - Monopoly pricing: Lerner (1934), Wilson (1993), ...
  - Optimal taxation: Saez (2001), ...
  - Welfare analysis: Chetty (2009), ...

Model

- Principal-agent model with the following timing:
  1. Principal offers a contract $w(\cdot)$.
  2. Agent observes $w(\cdot)$ and chooses effort a.
  3. Output $x \sim f(\cdot|a)$ and payoffs are realized. (Normalize $\mathbb{E}[x|a] = a$.)
- *Preferences:*
  - Agent's utility: $\int v(w(x))f(x|a)dx - c(a)$
  - Principal's profit: $\pi(w) := ma(w) - \int w(x)f(x|a(w))dx$.
- *Information:* $P = (f, c)$ is the **production environment**
  - The agent knows P
  - Principal knows v and has access to outcome data from a *status quo contract* w^A and a *test contract* w^B; *i.e.*, $f(x|a(w^A))$ and $f(x|a(w^B))$

Model

- Principal-agent model with the following timing:
    1. Principal offers a contract $w(\cdot)$.
    2. Agent observes $w(\cdot)$ and chooses effort a.
    3. Output $x \sim f(\cdot|a)$ and payoffs are realized. (Normalize $\mathbb{E}[x|a] = a$.)
- *Preferences:*
    - Agent's utility: $\int v(w(x))f(x|a)dx - c(a)$
    - Principal's profit: $\pi(w) := ma(w) - \int w(x)f(x|a(w))dx$.
- *Information:* $P = (f, c)$ is the **production environment**
    - The agent knows P
    - Principal knows v and has access to outcome data from a *status quo contract* w^A and a *test contract* w^B; *i.e.*, $f(x|a(w^A))$ and $f(x|a(w^B))$

Model

- Principal-agent model with the following timing:
  1. Principal offers a contract $w(\cdot)$.
  2. Agent observes $w(\cdot)$ and chooses effort a.
  3. Output $x \sim f(\cdot|a)$ and payoffs are realized. (Normalize $\mathbb{E}[x|a] = a$.)
- *Preferences:*
  - Agent's utility: $\int v(w(x))f(x|a)dx - c(a)$
  - Principal's profit: $\pi(w) := ma(w) - \int w(x)f(x|a(w))dx$.
- *Information:* $P = (f, c)$ is the **production environment**
  - The agent knows P
  - Principal knows v and has access to outcome data from a *status quo contract* w^A and a *test contract* w^B; *i.e.*, $f(x|a(w^A))$ and $f(x|a(w^B))$

Agent's Problem

- Given contract w, the agent's expected utility

$$u(w) = \max_a \int v(w(x))f(x|a)dx - c(a)$$

- Define the agent's *marginal incentives as*

$$I(w,a) = \int v(w(x))f_a(x|a)dx$$

- Assume optimal effort $a(w)$ is implicitly defined by

$$I(w,a) = c'(a)$$

i.e., optimal effort equates marginal benefit to marginal cost

Principal's Problem

- Principal's objective is to choose a profit-maximizing contract that gives the agent at least as much expected utility as w^A:

$$\max_{w(x),\, a} \; ma - \int w(x) f(x|a) dx$$

$$a \in \arg\max_{\tilde{a}} \int v(w(x)) f(x|\tilde{a}) dx - c(\tilde{a})$$

$$\text{s.t.} \int v(w(x)) f(x|a) dx - c(a) \geq u(w^A)$$

Outline

- **Theoretical results:**
  1. **Local adjustments**
     - Suppose the principal focuses on w such that $\left\| w - w^A \right\|$ is small
     - How to find the optimal adjustment?
     - Will show how a local A/B test provides the needed information
  2. Non-local Adjustments
     - Consider the full set of contracts
     - Provide conditions such that we can extrapolate logic
     - Will show that an A/B test suffices to find optimal adjustment
- **Empirical exercises**

Definition: Gateaux differential

- We will want to evaluate how profits are affected if the status quo contract is perturbed in some arbitrary direction.

  i.e., if the contract $w(x)$ is replaced by $w(x) + \theta t(x)$ for some small θ.

- Given contract w and function $q(w)$, **define** the *Gateaux differential* in the direction t:

$$\mathcal{D}q(w,t) = \lim_{\theta \to 0} \frac{q(w+\theta t) - q(w)}{\theta}$$

- Intuitively $\mathcal{D}q(w,t)$ measures how $q(w+\theta t)$ changes with θ for $\theta \simeq 0$

Definition: Gateaux differential

- We will want to evaluate how profits are affected if the status quo contract is perturbed in some arbitrary direction.

  i.e., if the contract $w(x)$ is replaced by $w(x) + \theta t(x)$ for some small θ.

- Given contract w and function $q(w)$, **define** the *Gateaux differential* in the direction t:

$$\mathcal{D}q(w, t) = \lim_{\theta \to 0} \frac{q(w + \theta t) - q(w)}{\theta}$$

- Intuitively $\mathcal{D}q(w, t)$ measures how $q(w + \theta t)$ changes with θ for $\theta \simeq 0$

Agent's Responses

- How does $a(w)$ and $u(w)$ change if w is replaced by $w + \theta t$ for $\theta \simeq 0$?

Lemma 1.

Locally adjusting contract w in direction t changes the agent's effort by

$$\mathcal{D}a\left(w,t\right) = \frac{\mathcal{D}I\left(w,t\right)}{c'' - \int v\left(w\right) f_{aa}\, dx}$$

where $\mathcal{D}I\left(w,t\right) := \int tv'\left(w\right) f_a\, dx$, and his expected utility by

$$\mathcal{D}u\left(w,t\right) = \int tv'\left(w\right) f\, dx$$

- Observations:
  1. The ratio $\mathcal{D}a\left(w,t\right)/\mathcal{D}I\left(w,t\right)$ does not depend on t
  2. Change in agent's utility does not directly depend on his cost function

Agent's Responses

- How does $a(w)$ and $u(w)$ change if w is replaced by $w + \theta t$ for $\theta \simeq 0$?

Lemma 1.

Locally adjusting contract w in direction t changes the agent's effort by

$$\mathcal{D}a\left(w,t\right)=\frac{\mathcal{D}I\left(w,t\right)}{c''-\int v\left(w\right)f_{aa}\,dx}$$

where $\mathcal{D}I\left(w,t\right) := \int tv'\left(w\right)f_a\,dx$, and his expected utility by

$$\mathcal{D}u\left(w,t\right)=\int tv'\left(w\right)f\,dx$$

- Observations:
  1. The ratio $\mathcal{D}a\left(w,t\right)/\mathcal{D}I\left(w,t\right)$ does not depend on t
  2. Change in agent's utility does not directly depend on his cost function

Principal's Problem

- The principal's expected profit under contract w is

$$\pi(w) = ma(w) - \int w(x) f(x|a(w)) dx$$

- If w is adjusted in direction t, profits change according to

$$\mathcal{D}\pi(w,t) = \left[m - \int w(x) f_a(x|a(w)) dx \right] \mathcal{D}a(w,t) - \int t(x) f(x|a(w)) dx$$

- The principal's goal is to

$$\max_{t : \|t\| \leq 1} \ \mathcal{D}\pi(w^A, t)$$
$$\text{s.t. } \mathcal{D}u(w^A, t) \geq 0$$

i.e., seeks direction t in which profits increase at fastest rate subject to giving the agent at least as much utility as the status quo contract.

Principal's Problem

- The principal's expected profit under contract w is

$$\pi(w) = ma(w) - \int w(x) f(x|a(w)) dx$$

- If w is adjusted in direction t, profits change according to

$$\mathcal{D}\pi(w,t) = \left[m - \int w(x) f_a(x|a(w)) dx\right] \mathcal{D}a(w,t) - \int t(x) f(x|a(w)) dx$$

- The principal's goal is to

$$\max_{t : \|t\| \leq 1} \mathcal{D}\pi(w^A, t)$$
$$\text{s.t. } \mathcal{D}u(w^A, t) \geq 0$$

i.e., seeks direction t in which profits increase at fastest rate subject to giving the agent at least as much utility as the status quo contract.

Principal's Problem

- The principal's expected profit under contract w is

$$\pi(w) = ma(w) - \int w(x)f(x|a(w))dx$$

- If w is adjusted in direction t, profits change according to

$$\mathcal{D}\pi(w,t) = \left[m - \int w(x)f_a(x|a(w))dx\right]\mathcal{D}a(w,t) - \int t(x)f(x|a(w))dx$$

- The principal's goal is to

$$\max_{t:\|t\|\leq 1} \mathcal{D}\pi(w^A,t)$$
$$\text{s.t. } \mathcal{D}u(w^A,t) \geq 0$$

i.e., seeks direction t in which profits increase at fastest rate subject to giving the agent at least as much utility as the status quo contract.

Informational Requirements

- What does the principal need to know to solve this problem?
- Principal's problem can be written in terms of primitives as

$$\max_{t:\|t\|\leq 1} \left(m - \int wf_a dx\right) \mathcal{D}a(w^A, t) - \int tf\, dx \qquad (\text{Adj}_{loc})$$

$$\text{s.t.} \int tv'\left(w^A\right) f\, dx \geq 0$$

Lemma 2 shows that:

The relevant aspects of production environment for solving (Adj_{loc}) are:

- $f(x|a(w^A))$ and $f_a(x|a(w^A))$, and
- $\mathcal{D}a(w^A, t)$ for all t

- We will argue that knowing $f_a(x|a(w^A))$ and $\mathcal{D}a(w^A, t)$ for *some* t suffices to evaluate $\mathcal{D}a(w^A, t')$ for every other t'.

Informational Requirements

- What does the principal need to know to solve this problem?
- Principal's problem can be written in terms of primitives as

$$\max_{t:\|t\|\leq 1} \left(m - \int w f_a dx\right) \mathcal{D}a(w^A, t) - \int tf\, dx \qquad (\text{Adj}_{loc})$$

$$\text{s.t.} \int tv'\left(w^A\right) f\, dx \geq 0$$

Lemma 2 shows that:

The relevant aspects of production environment for solving (Adj_{loc}) are:

- $f(x|a(w^A))$ and $f_a(x|a(w^A))$, and
- $\mathcal{D}a(w^A, t)$ for all t

- We will argue that knowing $f_a(x|a(w^A))$ and $\mathcal{D}a(w^A, t)$ for *some* t suffices to evaluate $\mathcal{D}a(w^A, t')$ for every other t'.

Informational Requirements

- What does the principal need to know to solve this problem?
- Principal's problem can be written in terms of primitives as

$$\max_{t:\|t\|\leq 1} \left(m - \int w f_a dx\right) \mathcal{D}a(w^A, t) - \int tf\, dx \qquad (\text{Adj}_{loc})$$

$$\text{s.t.} \int tv'\left(w^A\right) f\, dx \geq 0$$

Lemma 2 shows that:

The relevant aspects of production environment for solving (Adj_{loc}) are:

- $f(x|a(w^A))$ and $f_a(x|a(w^A))$, and
- $\mathcal{D}a(w^A, t)$ for all t

- We will argue that knowing $f_a(x|a(w^A))$ and $\mathcal{D}a(w^A, t)$ for *some* t suffices to evaluate $\mathcal{D}a(w^A, t')$ for every other t'.

Definitions: A/B Test and Local A/B Test

- An *A/B test* for contracts w^A and w^B is a pair

$$AB(w^A, w^B) = \left(f^A, f^B\right)$$

- A *local A/B test* is a triple

$$LAB(w^A, w^B) = \left(f^A, f_a^A, \mathcal{D}a(w^A, w^B)\right)$$

  - *Interpretation:* Test comprises data for w^A and $w^A + \theta w^B$ as $\theta \to 0$

- Assume that $\mathcal{D}a(w^A, w^B) \neq 0$; *i.e.*, local A/B test is informative

Definitions: A/B Test and Local A/B Test

- An *A/B test* for contracts w^A and w^B is a pair

$$AB(w^A, w^B) = \left(f^A, f^B\right)$$

- A *local A/B test* is a triple

$$LAB(w^A, w^B) = \left(f^A,\ f_a^A,\ \mathcal{D}a(w^A, w^B)\right)$$

  - *Interpretation:* Test comprises data for w^A and $w^A + \theta w^B$ as $\theta \to 0$

- Assume that $\mathcal{D}a(w^A, w^B) \neq 0$; *i.e.*, local A/B test is informative

Sufficient Statistic Result

- *Recall:* The principal seeks to solve

$$\max_{t:\|t\|\leq 1} \left(m - \int w f_a^A dx\right) \mathcal{D}a(w^A, t) - \int t f^A \, dx \qquad (\text{Adj}_{loc})$$

$$\text{s.t.} \int t v'\left(w^A\right) f^A \, dx \geq 0$$

Proposition 1 shows that:

The information provided by a local A/B test suffices to solve problem.

- Knowing f_a^A, the principal can evaluate for every t,

$$\mathcal{D}I\left(w,t\right) := \int t v'\left(w\right) f_a^A \, dx$$

- Knowing f_a^A and $\mathcal{D}a(w^A, w^B)$, she can evaluate for every t,

$$\mathcal{D}a\left(w^A, t\right) = \frac{\mathcal{D}a\left(w^A, w^B\right)}{\mathcal{D}I\left(w^A, w^B\right)} \mathcal{D}I\left(w^A, t\right)$$

Sufficient Statistic Result

- *Recall:* The principal seeks to solve

$$\max_{t:\|t\|\leq 1} \left(m - \int w f_a^A dx\right) \mathcal{D}a(w^A, t) - \int t f^A \, dx \qquad (\text{Adj}_{loc})$$

$$\text{s.t.} \int t v'\left(w^A\right) f^A \, dx \geq 0$$

Proposition 1 shows that:

The information provided by a local A/B test suffices to solve problem.

- Knowing f_a^A, the principal can evaluate for every t,

$$\mathcal{D}I\left(w, t\right) := \int t v'\left(w\right) f_a^A \, dx$$

- Knowing f_a^A and $\mathcal{D}a(w^A, w^B)$, she can evaluate for every t,

$$\mathcal{D}a\left(w^A, t\right) = \frac{\mathcal{D}a\left(w^A, w^B\right)}{\mathcal{D}I\left(w^A, w^B\right)} \mathcal{D}I\left(w^A, t\right)$$

Sufficient Statistic Result

- *Recall:* The principal seeks to solve

$$\max_{t:\|t\|\leq 1} \left(m - \int w f_a^A dx\right) \mathcal{D}a(w^A, t) - \int t f^A \, dx \qquad (\text{Adj}_{loc})$$

$$\text{s.t.} \int t v'\left(w^A\right) f^A \, dx \geq 0$$

Proposition 1 shows that:

The information provided by a local A/B test suffices to solve problem.

- Knowing f_a^A, the principal can evaluate for every t,

$$\mathcal{D}I\left(w, t\right) := \int t v'\left(w\right) f_a^A \, dx$$

- Knowing f_a^A and $\mathcal{D}a(w^A, w^B)$, she can evaluate for every t,

$$\mathcal{D}a\left(w^A, t\right) = \frac{\mathcal{D}a\left(w^A, w^B\right)}{\mathcal{D}I\left(w^A, w^B\right)} \mathcal{D}I\left(w^A, t\right)$$

Sufficient Statistic Result

- *Recall:* The principal seeks to solve

$$\max_{t:\|t\|\leq 1} \left(m - \int wf_a^A dx\right)\mathcal{D}a(w^A, t) - \int tf^A\, dx \qquad (\text{Adj}_{loc})$$

$$\text{s.t.} \quad \int tv'\left(w^A\right) f^A\, dx \geq 0$$

Proposition 1 shows that:

The information provided by a local A/B test suffices to solve problem.

- Knowing f^A, she can evaluate
  - the effect of t on her compensation costs; *i.e.*, $\int tf^A dx$
  - the (participation) constraint

Optimal Local Adjustment

- (Adj_{loc}) is a standard convex optimization program

Proposition 2:

- Characterizes the optimal local adjustment; and
- Gives a condition for w^A to be locally optimal; *i.e.*, $t^* \equiv 0$

Outline

- **Theoretical results:**
  1. Local adjustments
     - Suppose the principal focuses on w such that $\left\| w - w^A \right\|$ is small
     - How to find the optimal adjustment?
     - Will show how a local A/B test provides the needed information
  2. **Non-local Adjustments**
     - Consider the full set of contracts
     - Provide conditions such that we can extrapolate logic
     - Will show that an A/B test suffices to find optimal adjustment
- **Empirical exercises**

Non-Local Adjustments

- In practice,
  - A/B tests are *not* local, and
  - firms are interested in non-local adjustments
- To find the optimal non-local adjustment, in general, one must know the entire production environment $P = (f, c)$.
- We provide two conditions allowing us to extrapolate the ideas from previous part to assess such adjustments with only an A/B test.

Non-Local Adjustments

- In practice,
  - A/B tests are *not* local, and
  - firms are interested in non-local adjustments
- To find the optimal non-local adjustment, in general, one must know the entire production environment $P = (f, c)$.
- We provide two conditions allowing us to extrapolate the ideas from previous part to assess such adjustments with only an A/B test.

Condition 1: Output distribution is affine in effort

Condition 1.

The output distribution $f(x|a)$ is affine in a; *i.e.*, $f(x|a) = g(x) + ah(x)$ for some functions $g(x)$ and $h(x)$

1. Given an A/B test (f^A, f^B), one can determine $f(\cdot|a)$ for every a
2. Because $f_a(x|a) \equiv h(x)$ for all a, the marginal incentives

$$I(w) = \int v(w(x))h(x)dx$$

do not depend on a.

Condition 1: Output distribution is affine in effort

Condition 1.

The output distribution $f(x|a)$ is affine in a; *i.e.*, $f(x|a) = g(x) + ah(x)$ for some functions $g(x)$ and $h(x)$

1. Given an A/B test (f^A, f^B), one can determine $f(\cdot|a)$ for every a
2. Because $f_a(x|a) \equiv h(x)$ for all a, the marginal incentives

$$I(w) = \int v(w(x))h(x)dx$$

do not depend on a.

Condition 1: Output distribution is affine in effort

Condition 1.

The output distribution $f(x|a)$ is affine in a; *i.e.*, $f(x|a) = g(x) + ah(x)$ for some functions $g(x)$ and $h(x)$

1. Given an A/B test (f^A, f^B), one can determine $f(\cdot|a)$ for every a
2. Because $f_a(x|a) \equiv h(x)$ for all a, the marginal incentives

$$I(w) = \int v(w(x))h(x)dx$$

do not depend on a.

Condition 2: Isoelastic effort costs

Condition 2.

The agent has isoelastic effort costs:

$$c'(a) = e^{-\beta/\epsilon} a^{1/\epsilon}$$

- This condition implies that for any contract w, we have

$$\ln a(w) = \beta + \epsilon \ln I(w)$$

- Given an A/B test,
  - One can evaluate $I(w)$ for any w, and
  - $a(w^A)$ and $a(w^B)$ since $a(w)$ is the expected output given w.
- Thus, principal can pin down β and ϵ, and predict $a(w)$ for any w

Condition 2: Isoelastic effort costs

Condition 2.

The agent has isoelastic effort costs:

$$c'(a) = e^{-\beta/\epsilon} a^{1/\epsilon}$$

- This condition implies that for any contract w, we have

$$\ln a(w) = \beta + \epsilon \ln I(w)$$

- Given an A/B test,
  - One can evaluate $I(w)$ for any w, and
  - $a(w^A)$ and $a(w^B)$ since $a(w)$ is the expected output given w.
- Thus, principal can pin down β and ϵ, and predict $a(w)$ for any w

Principal's Problem

- The principal's profit if she offers contract w is

$$\pi(w) = ma(w) - \int w(x)\left[g(x) + a(w)h(x)\right]dx$$

- Given status quo contract w^A, she solves

$$\max_{w} \ \pi(w)$$
$$\text{s.t.} \ u(w) \geq u(w^A)$$

Proposition 3 shows that:

The information provided by an A/B test suffices to solve this problem.

- This problem can be solved in two stages a-la Grossman and Hart:
  i. Fix an a and find cost-minimizing contract that implements this effort
  ii. Find profit-maximizing $a(w)$

Principal's Problem

- The principal's profit if she offers contract w is

$$\pi(w) = ma(w) - \int w(x)\left[g(x) + a(w)h(x)\right]dx$$

- Given status quo contract w^A, she solves

$$\max_{w} \ \pi(w) \quad \text{s.t. } u(w) \geq u(w^A)$$

Proposition 3 shows that:

The information provided by an A/B test suffices to solve this problem.

- This problem can be solved in two stages a-la Grossman and Hart:
  - i. Fix an a and find cost-minimizing contract that implements this effort
  - ii. Find profit-maximizing $a(w)$

Dataset

- *Goal:* Assess performance of our model
- Dataset from DellaVigna and Pope (2017)
- Real-effort experiment on M-Turk: Subjects press a-b keys for 10 min
- 7 treatments with different monetary incentives:

	Contract (in ¢)	Avg. #points (x)	N
No incentives	$w_1(x) = 100$	1521	540
Piece-rate	$w_2(x) = 100 + 0.001x$	1883	538
	$w_3(x) = 100 + 0.01x$	2029	558
	$w_4(x) = 100 + 0.04x$	2132	566
	$w_5(x) = 100 + 0.10x$	2175	538
Bonus	$w_6(x) = 100 + 40\,\mathbb{I}_{\{x \geq 2000\}}$	2136	545
	$w_7(x) = 100 + 80\,\mathbb{I}_{\{x \geq 2000\}}$	2187	532

- Each subject participated in a single treatment, once.

Two Exercises

I. Assess our model's ability to predict performance out-of-sample
   - Use data from each pair of treatments to predict mean performance in the remaining treatments.
   - We then compare our predictions to observed performance

II. Assess the performance of optimal adjustments
   - Use all treatments to construct production environment (f, c)
   - Using (f, c), compute the *benchmark-optimal* contract
   - For each pair of treatments, take this pair to constitute our A/B test, and use its data to compute the *test-optimal* contract
   - Compare its profit to that of w^A and benchmark-optimal contract

Two Exercises

I. Assess our model's ability to predict performance out-of-sample
   - Use data from each pair of treatments to predict mean performance in the remaining treatments.
   - We then compare our predictions to observed performance

II. Assess the performance of optimal adjustments
   - Use all treatments to construct production environment (f, c)
   - Using (f, c), compute the *benchmark-optimal* contract
   - For each pair of treatments, take this pair to constitute our A/B test, and use its data to compute the *test-optimal* contract
   - Compare its profit to that of w^A and benchmark-optimal contract

Prediction Exercise: Procedure

1. Assume each subject has CRRA utility: $v'(\omega) = \omega^{-\rho}$ with $\rho = 0.3$
2. Normalize $a(w_i) = (\text{Avg. \#points})_i$.
3. Take an arbitrary pair of treatments, labeled w^A and w^B
   i. Using a kernel estimator, construct the pdfs $\hat{f}^A$ and $\hat{f}^B$
   ii. For every treatment C, compute the marginal incentives

$$\hat{I}_C^{AB} = \int v\left(w^C(x)\right)\hat{h}^{AB}(x)\,dx \text{ , where } \hat{h}^{AB}(x) = \frac{\hat{f}^A(x) - \hat{f}^B(x)}{a^A - a^B}$$

   iii. Estimate the cost parameters $\hat{\varepsilon}^{AB}$ and $\hat{\beta}^{AB}$
   iv. Predict performance for every treatment $C \notin \{A, B\}$ using

$$\ln \hat{a}_C^{AB} = \hat{\beta}^{AB} + \hat{\varepsilon}^{AB} \ln \hat{I}_C^{AB}$$

- Focus on A/B tests where w^A and w^B belong to same class.

Prediction Exercise: Procedure

1. Assume each subject has CRRA utility: $v'(\omega) = \omega^{-\rho}$ with $\rho = 0.3$
2. Normalize $a(w_i) = (\text{Avg. } \#\text{points})_i$.
3. Take an arbitrary pair of treatments, labeled w^A and w^B
   - i. Using a kernel estimator, construct the pdfs $\hat{f}^A$ and $\hat{f}^B$
   - ii. For every treatment C, compute the marginal incentives

$$\hat{I}_C^{AB} = \int v\left(w^C(x)\right) \hat{h}^{AB}(x)\, dx \text{ , where } \hat{h}^{AB}(x) = \frac{\hat{f}^A(x) - \hat{f}^B(x)}{a^A - a^B}$$

   - iii. Estimate the cost parameters $\hat{\epsilon}^{AB}$ and $\hat{\beta}^{AB}$
   - iv. Predict performance for every treatment $C \notin \{A, B\}$ using

$$\ln \hat{a}_C^{AB} = \hat{\beta}^{AB} + \hat{\varepsilon}^{AB} \ln \hat{I}_C^{AB}$$

- Focus on A/B tests where w^A and w^B belong to same class.

Prediction Exercise: Procedure

1. Assume each subject has CRRA utility: $v'(\omega) = \omega^{-\rho}$ with $\rho = 0.3$
2. Normalize $a(w_i) = (\text{Avg. } \#\text{points})_i$.
3. Take an arbitrary pair of treatments, labeled w^A and w^B
   i. Using a kernel estimator, construct the pdfs $\hat{f}^A$ and $\hat{f}^B$
   ii. For every treatment C, compute the marginal incentives

$$\hat{I}_C^{AB} = \int v\left(w^C(x)\right)\hat{h}^{AB}(x)\,dx \text{ , where } \hat{h}^{AB}(x) = \frac{\hat{f}^A(x) - \hat{f}^B(x)}{a^A - a^B}$$

   iii. Estimate the cost parameters $\hat{\epsilon}^{AB}$ and $\hat{\beta}^{AB}$
   iv. Predict performance for every treatment $C \notin \{A, B\}$ using

$$\ln \hat{a}_C^{AB} = \hat{\beta}^{AB} + \hat{\varepsilon}^{AB} \ln \hat{I}_C^{AB}$$

- Focus on A/B tests where w^A and w^B belong to same class.

Prediction Exercise Illustrated

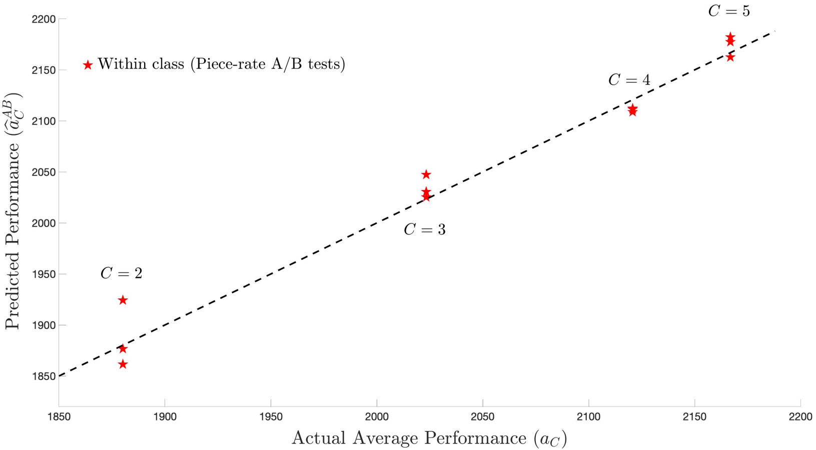

i. Predicted performance is close to actual performance

ii. Under-predicts performance in (bonus) treatments 6 and 7

iii. Predictions are similar no matter which A/B test is used

Prediction Exercise Illustrated

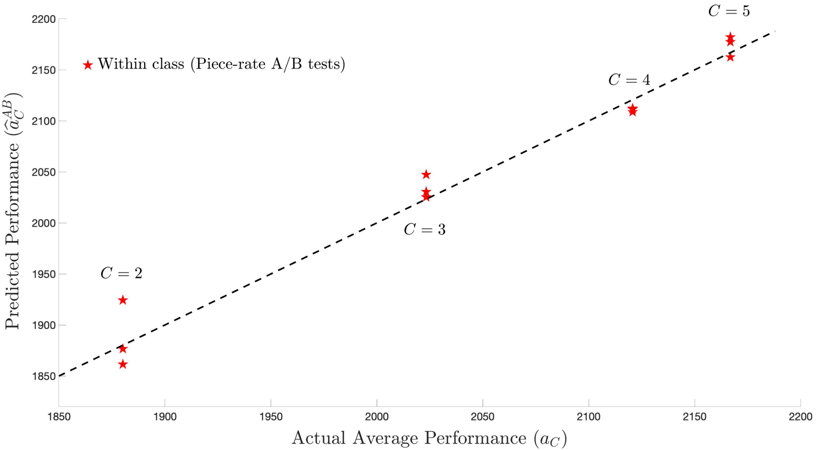

i. Predicted performance is close to actual performance

ii. Under-predicts performance in (bonus) treatments 6 and 7

iii. Predictions are similar no matter which A/B test is used

Prediction Exercise Illustrated

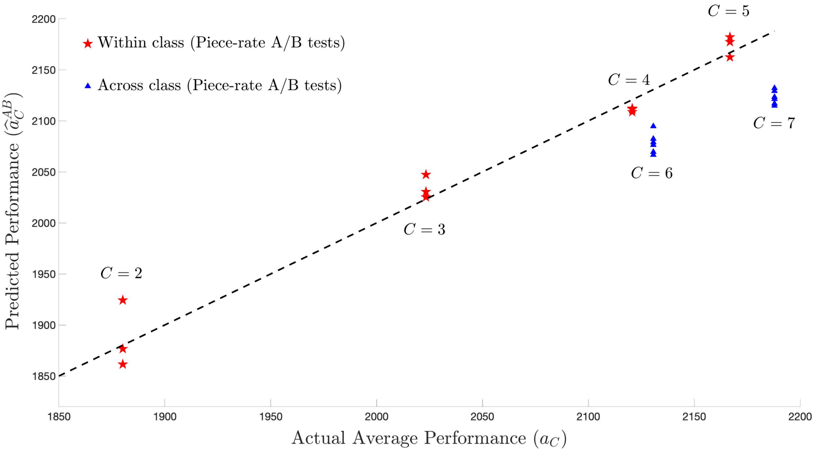

i. Predicted performance is close to actual performance

ii. Under-predicts performance in (bonus) treatments 6 and 7

iii. Predictions are similar no matter which A/B test is used

Prediction Exercise Illustrated

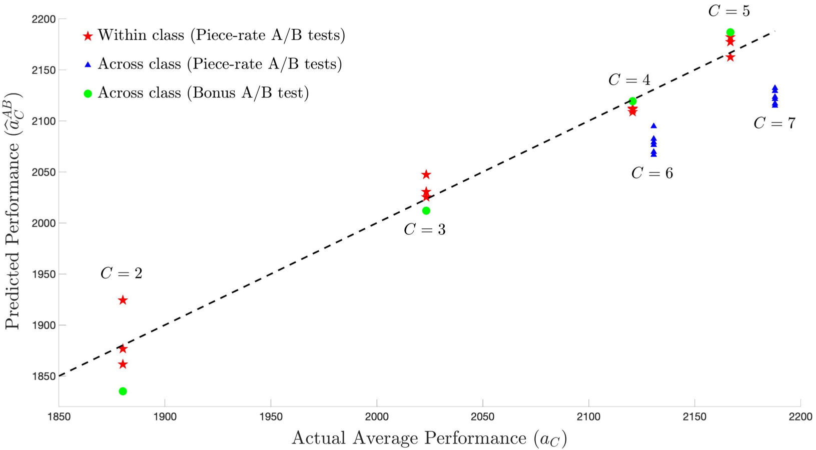

i. Predicted performance is close to actual performance

ii. Under-predicts performance in (bonus) treatments 6 and 7

iii. Predictions are similar no matter which A/B test is used

Prediction Exercise Illustrated

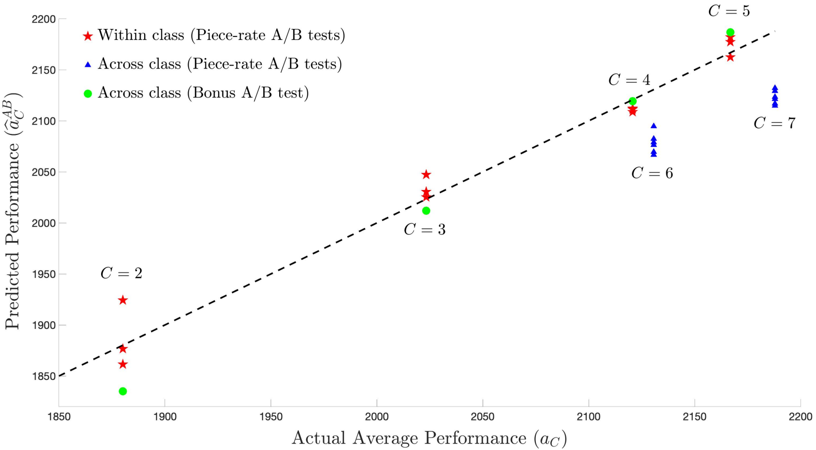

i. Predicted performance is close to actual performance
ii. Under-predicts performance in (bonus) treatments 6 and 7
iii. Predictions are similar no matter which A/B test is used

Prediction Exercise: Sensitivity

Coefficient of relative risk aversion (ρ)	0.3

Homogeneous A/B Tests (w^A and w^B belong to same class)	
Corr $\left(\hat{a}_C^{AB}, a_C\right)$	0.94
Mean Absolute Percentage Error (APE)	1.59
Worst-case APE	3.34

i. Hybrid A/B tests sometimes generate poor predictions

ii. Prediction accuracy is insensitive to the coefficient of risk aversion

iii. Similar accuracy if we assume quadratic utility function

Prediction Exercise: Sensitivity

Coefficient of relative risk aversion (ρ)	0.3
Homogeneous A/B Tests (w^A and w^B belong to same class)	
Corr$\left(\hat{a}_C^{AB}, a_C\right)$	0.94
Mean Absolute Percentage Error (APE)	1.59
Worst-case APE	3.34
Hybrid A/B Tests (w^A and w^B belong to different classes)	
Corr$\left(\hat{a}_C^{AB}, a_C\right)$	0.84
Mean APE	2.16
Worst-case APE	10.70

i. Hybrid A/B tests sometimes generate poor predictions

ii. Prediction accuracy is insensitive to the coefficient of risk aversion

iii. Similar accuracy if we assume quadratic utility function

Prediction Exercise: Sensitivity

Coefficient of relative risk aversion (ρ)	0	0.3	0.5	1
Homogeneous A/B Tests (w^A and w^B belong to same class)				
$\text{Corr}\left(\hat{a}_C^{AB}, a_C\right)$	0.92	0.94	0.96	0.97
Mean Absolute Percentage Error (APE)	1.76	1.59	1.54	1.64
Worst-case APE	3.65	3.34	3.08	4.30
Hybrid A/B Tests (w^A and w^B belong to different classes)				
$\text{Corr}\left(\hat{a}_C^{AB}, a_C\right)$	0.86	0.84	0.83	0.78
Mean APE	2.19	2.16	2.14	2.18
Worst-case APE	10.61	10.70	11.40	12.69

i. Hybrid A/B tests sometimes generate poor predictions

ii. Prediction accuracy is insensitive to the coefficient of risk aversion

iii. Similar accuracy if we assume quadratic utility function

Exercise 2: Constructing a benchmark

- Use all 7 treatments to construct production environment (f, c)
- Constructing f:
  - For each $a = a(w_i)$, we use a kernel estimator to construct pdf $f(x|a)$
  - For each $a \neq a(w_i)$, we use a spline interpolation to construct $f(x|a)$
- Constructing c:
  - Assume $c'(a) = c_0 a^p - I_0$ for some parameters c_0, p, and I_0 TBD
  - Assume $u'(\omega) = \omega^{-\rho}$, $\rho = 0.3$, and fit parameters with NLS estimation
- Assume that each unit of x is worth $m = 0.2$¢ to the principal
- For each treatment C, compute the profit-maximizing contract that gives the agent at least as much utility as w^C.
  - Denote profit of the benchmark-optimal contract by $\pi^*(w^C)$

Exercise 2: Constructing a benchmark

- Use all 7 treatments to construct production environment (f, c)
- Constructing f:
  - For each $a = a(w_i)$, we use a kernel estimator to construct pdf $f(x|a)$
  - For each $a \neq a(w_i)$, we use a spline interpolation to construct $f(x|a)$
- Constructing c:
  - Assume $c'(a) = c_0 a^p - I_0$ for some parameters c_0, p, and I_0 TBD
  - Assume $u'(\omega) = \omega^{-\rho}$, $\rho = 0.3$, and fit parameters with NLS estimation
- Assume that each unit of x is worth $m = 0.2$¢ to the principal
- For each treatment C, compute the profit-maximizing contract that gives the agent at least as much utility as w^C.
  - Denote profit of the benchmark-optimal contract by $\pi^*(w^C)$

Exercise 2: Constructing a benchmark

- Use all 7 treatments to construct production environment (f, c)
- Constructing f:
  - For each $a = a(w_i)$, we use a kernel estimator to construct pdf $f(x|a)$
  - For each $a \neq a(w_i)$, we use a spline interpolation to construct $f(x|a)$
- Constructing c:
  - Assume $c'(a) = c_0 a^p - I_0$ for some parameters c_0, p, and I_0 TBD
  - Assume $u'(\omega) = \omega^{-\rho}$, $\rho = 0.3$, and fit parameters with NLS estimation
- Assume that each unit of x is worth $m = 0.2$¢ to the principal
- For each treatment C, compute the profit-maximizing contract that gives the agent at least as much utility as w^C.
  - Denote profit of the benchmark-optimal contract by $\pi^*(w^C)$

Exercise 2: Constructing a benchmark

- Use all 7 treatments to construct production environment (f, c)
- Constructing f:
  - For each $a = a(w_i)$, we use a kernel estimator to construct pdf $f(x|a)$
  - For each $a \neq a(w_i)$, we use a spline interpolation to construct $f(x|a)$
- Constructing c:
  - Assume $c'(a) = c_0 a^p - l_0$ for some parameters c_0, p, and l_0 TBD
  - Assume $u'(\omega) = \omega^{-\rho}$, $\rho = 0.3$, and fit parameters with NLS estimation
- Assume that each unit of x is worth $m = 0.2$¢ to the principal
- For each treatment C, compute the profit-maximizing contract that gives the agent at least as much utility as w^C.
  - Denote profit of the benchmark-optimal contract by $\pi^*(w^C)$

Exercise 2: Constructing a benchmark

- Use all 7 treatments to construct production environment (f, c)
- Constructing f:
  - For each $a = a(w_i)$, we use a kernel estimator to construct pdf $f(x|a)$
  - For each $a \neq a(w_i)$, we use a spline interpolation to construct $f(x|a)$
- Constructing c:
  - Assume $c'(a) = c_0 a^p - I_0$ for some parameters c_0, p, and I_0 TBD
  - Assume $u'(\omega) = \omega^{-\rho}$, $\rho = 0.3$, and fit parameters with NLS estimation
- Assume that each unit of x is worth $m = 0.2$¢ to the principal
- For each treatment C, compute the profit-maximizing contract that gives the agent at least as much utility as w^C.
  - Denote profit of the benchmark-optimal contract by $\pi^*(w^C)$

Exercise 2: Optimal Adjustments

- Take an arbitrary pair of treatments, labeled w^A and w^B
- Using the same procedure as in the prediction exercise, construct the pdfs $\hat{f}^A$ and $\hat{f}^B$, and the parameters $\hat{\epsilon}^{AB}$ and $\hat{\beta}^{AB}$
- For every treatment C, compute the test-optimal contract that gives the agent at least as much utility as w^C
- Using the constructed production environment, evaluate the profit of each test-optimal contract, which we denote $\pi^{AB}(w^C)$

Exercise 2: Optimal Adjustments

- Take an arbitrary pair of treatments, labeled w^A and w^B
- Using the same procedure as in the prediction exercise, construct the pdfs $\hat{f}^A$ and $\hat{f}^B$, and the parameters $\hat{\epsilon}^{AB}$ and $\hat{\beta}^{AB}$
- For every treatment C, compute the test-optimal contract that gives the agent at least as much utility as w^C
- Using the constructed production environment, evaluate the profit of each test-optimal contract, which we denote $\pi^{AB}(w^C)$

Exercise 2: Optimal Adjustments

- Take an arbitrary pair of treatments, labeled w^A and w^B
- Using the same procedure as in the prediction exercise, construct the pdfs $\hat{f}^A$ and $\hat{f}^B$, and the parameters $\hat{\epsilon}^{AB}$ and $\hat{\beta}^{AB}$
- For every treatment C, compute the test-optimal contract that gives the agent at least as much utility as w^C
- Using the constructed production environment, evaluate the profit of each test-optimal contract, which we denote $\pi^{AB}(w^C)$

Evaluating the Performance of Optimal Adjustments

- **Maximum available gains** for treatment C:

$$\text{MaxGains}^{C} = \pi^{*}\left(w^{C}\right) - \pi\left(w^{C}\right)$$

- **Average Realized gains** for treatment C:

$$\text{AvgGains}^{C} = \frac{1}{|Hom.\ Tests|} \sum_{A,B \in Hom} \pi^{AB}\left(w^{C}\right) - \pi\left(w^{C}\right)$$

i.e., we average the realized gains across all homogeneous A/B tests.

- Averaging across treatments $C \in \{2, \ldots, 7\}$, the average realized gains are 68% of the maximum available gains.
- Ratio is insensitive to different choices of m and coefficient of RRA

Evaluating the Performance of Optimal Adjustments

- **Maximum available gains** for treatment C:

$$\text{MaxGains}^C = \pi^* \left(w^C\right) - \pi \left(w^C\right)$$

- **Average Realized gains** for treatment C:

$$\text{AvgGains}^C = \frac{1}{|Hom.\ Tests|} \sum_{A,B \in Hom} \pi^{AB} \left(w^C\right) - \pi \left(w^C\right)$$

i.e., we average the realized gains across all homogeneous A/B tests.

- Averaging across treatments $C \in \{2, \dots, 7\}$, the average realized gains are 68% of the maximum available gains.

- Ratio is insensitive to different choices of m and coefficient of RRA

Evaluating the Performance of Optimal Adjustments

- **Maximum available gains** for treatment C:

$$\text{MaxGains}^C = \pi^* \left(w^C\right) - \pi \left(w^C\right)$$

- **Average Realized gains** for treatment C:

$$\text{AvgGains}^C = \frac{1}{|Hom.\ Tests|} \sum_{A,B \in Hom} \pi^{AB} \left(w^C\right) - \pi \left(w^C\right)$$

i.e., we average the realized gains across all homogeneous A/B tests.

- Averaging across treatments $C \in \{2, \ldots, 7\}$, the average realized gains are 68% of the maximum available gains.
- Ratio is insensitive to different choices of m and coefficient of RRA

Performance of Optimal Adjustments: Illustrated

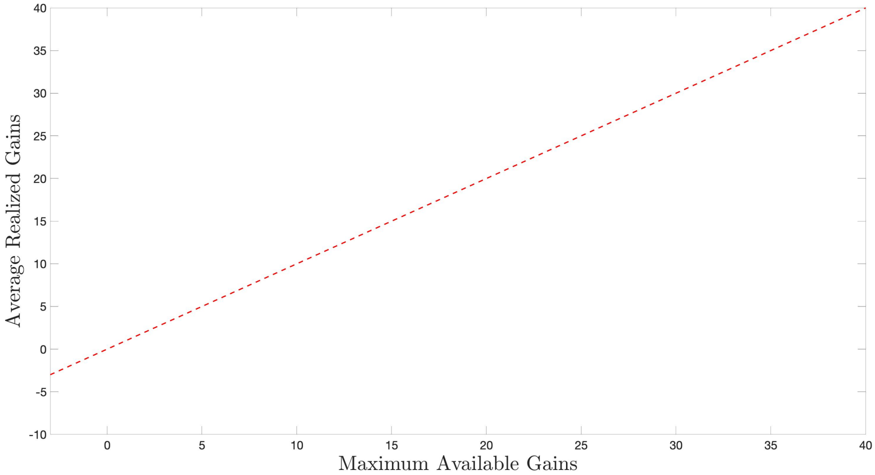

- The $AvgGains_C$ are approximately $3 smaller than $MaxGains^C$

Performance of Optimal Adjustments: Illustrated

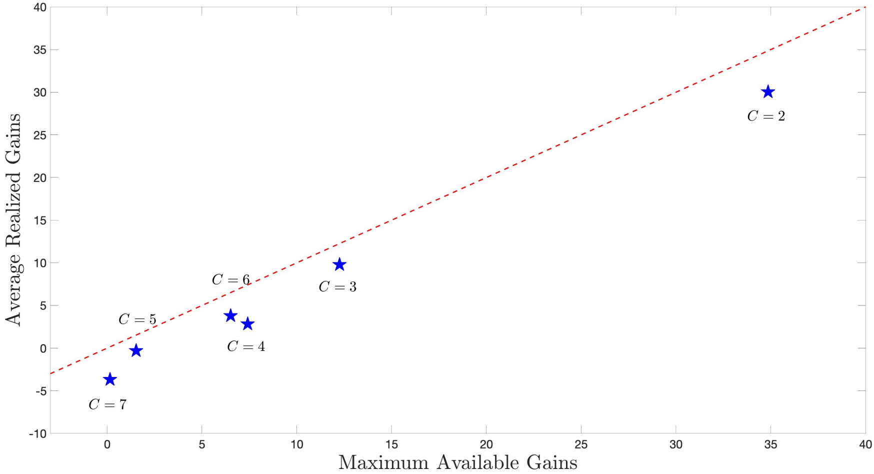

- The AvgGains_C are approximately \$3 smaller than MaxGains^C

Performance of Optimal Adjustments: Illustrated

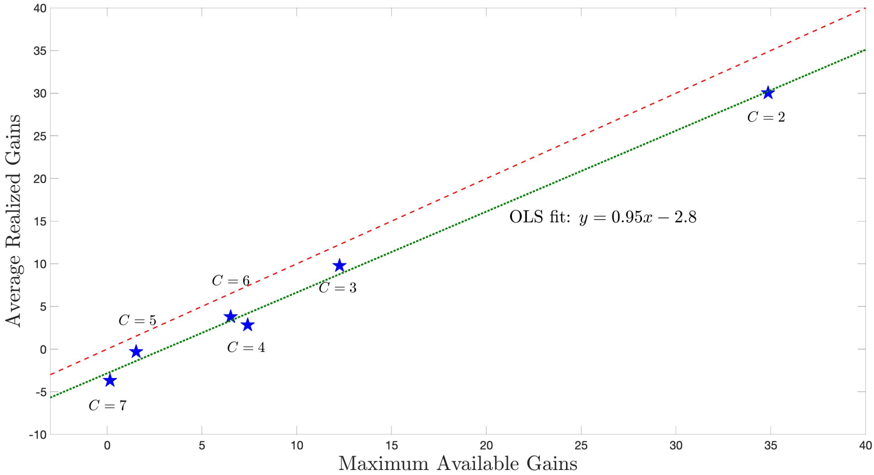

- The AvgGains_C are approximately \$3 smaller than MaxGains^C

Decomposing the \$3 Gap

- Two reasons why AvgGains is smaller than MaxGains:
  a. Principal overpays to implement a given effort; or
  b. Implements an effort that is not profit-maximizing
- For each treatment C, we compare
  (A) Wage bill of the test-optimal contract to wage bill of the cost-minimizing contract that implements the same effort.
  (B) Effort implemented by test-optimal contract to the optimal effort.
- On average:
  (A') The test-optimal contract overpays by $\sim$ \$1.8
  (B') Implements an effort ~ 7 units too low, *losing* $7 \times (m = 0.2) = \$1.4$

Decomposing the $3 Gap

- Two reasons why AvgGains is smaller than MaxGains:
  a. Principal overpays to implement a given effort; or
  b. Implements an effort that is not profit-maximizing
- For each treatment C, we compare
  (A) Wage bill of the test-optimal contract to wage bill of the cost-minimizing contract that implements the same effort.
  (B) Effort implemented by test-optimal contract to the optimal effort.
- On average:
  (A') The test-optimal contract overpays by ~ $1.8
  (B') Implements an effort ~ 7 units too low, *losing* $7 \times (m = 0.2) = \$1.4$

Decomposing the $3 Gap

- Two reasons why AvgGains is smaller than MaxGains:
  a. Principal overpays to implement a given effort; or
  b. Implements an effort that is not profit-maximizing
- For each treatment C, we compare
  (A) Wage bill of the test-optimal contract to wage bill of the cost-minimizing contract that implements the same effort.
  (B) Effort implemented by test-optimal contract to the optimal effort.
- On average:
  (A') The test-optimal contract overpays by $\sim \$1.8$
  (B') Implements an effort ~ 7 units too low, *losing* $7 \times (m = 0.2) = \$1.4$

Decomposing the \$3 Gap

- Two reasons why AvgGains is smaller than MaxGains:
  a. Principal overpays to implement a given effort; or
  b. Implements an effort that is not profit-maximizing
- For each treatment C, we compare
  (A) Wage bill of the test-optimal contract to wage bill of the cost-minimizing contract that implements the same effort.
  (B) Effort implemented by test-optimal contract to the optimal effort.
- On average:
  (A') The test-optimal contract overpays by $\sim$ \$1.8
  (B') Implements an effort ~ 7 units too low, *losing* $7 \times (m = 0.2) = \$1.4$

Decomposing the \$3 Gap

- Two reasons why AvgGains is smaller than MaxGains:
  - a. Principal overpays to implement a given effort; or
  - b. Implements an effort that is not profit-maximizing
- For each treatment C, we compare
  - (A) Wage bill of the test-optimal contract to wage bill of the cost-minimizing contract that implements the same effort.
  - (B) Effort implemented by test-optimal contract to the optimal effort.
- On average:
  - (A') The test-optimal contract overpays by ~ \$1.8
  - (B') Implements an effort ~ 7 units too low, *losing* $7 \times (m = 0.2) = \$1.4$

Agent's Utility

- For ea. A/B test and treatment C, we sought a profit-maximizing contract that gives the agent at least as much utility as w^C.
- This figure illustrates the CDF of the ratio of the agent's utility under the test-optimal contract to the target utility.
- Ratio ranges from 0.97 to 1.07, and it is greater than 1 in ~ 75% of cases.

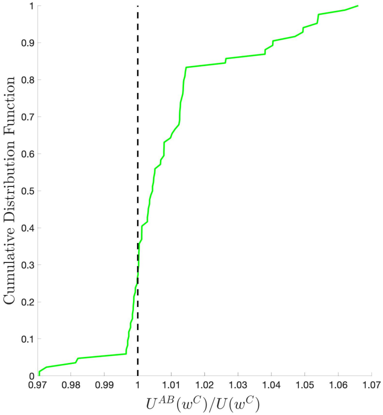

Agent's Utility

- For ea. A/B test and treatment C, we sought a profit-maximizing contract that gives the agent at least as much utility as w^C.
- This figure illustrates the CDF of the ratio of the agent's utility under the test-optimal contract to the target utility.
- Ratio ranges from 0.97 to 1.07, and it is greater than 1 in ~ 75% of cases.

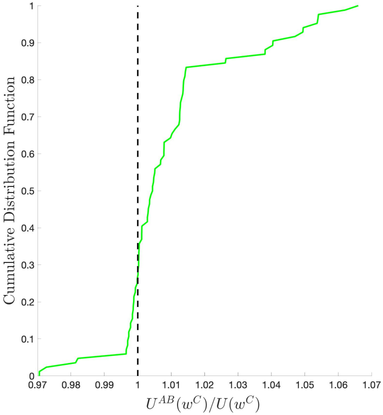

Beyond the Classic Model

I. *Multitasking.* Effort $\mathbf{a} \in \mathbb{R}^M$ and output $\mathbf{x} \in \mathbb{R}^M$

- *e.g.*, effort towards quantity & quality, or selling different products.
- Need $\lceil (M+1)/2 \rceil$ linearly independently test contracts

II. *Parametric contract classes.* Restrict attention to contracts of the form w_α, where α is a vector of parameters.

- *e.g.*, linear, piecewise linear, or bonus contracts
- Similar logic and same informational requirements

III. *Heterogeneous workers.* Principal offers a common contract to agents with heterogeneous effort costs.

- Straightforward application
- Can induce *selection* by imposing participation for subset of types

Summary & Future Work

- What does a firm need to know to improve an existing contract?
  - We showed how an A/B test can provide this information
  - Provided a proof of concept
- Many open questions, lots to do:
  - Statistical and approximation error?
  - How to design an A/B test optimally?
  - How to account for strategic manipulation; *e.g.*, ratchet effects
  - Intertemporal (dynamic) incentives?
  - Incentive design for teams of workers?

Summary & Future Work

- What does a firm need to know to improve an existing contract?
  - We showed how an A/B test can provide this information
  - Provided a proof of concept
- Many open questions, lots to do:
  - Statistical and approximation error?
  - How to design an A/B test optimally?
  - How to account for strategic manipulation; *e.g.*, ratchet effects
  - Intertemporal (dynamic) incentives?
  - Incentive design for teams of workers?