Improving Lifecycle Regret with Tie-Breaker Designs

Apoorva Lal [1] and Kyle Woodward [1]

[1]Amazon

Motivation and key findings

The Randomized Controlled Trial (RCT) is the "gold standard" for causal inference [1]. By randomly assigning users to treatment or control, RCTs provide unbiased impact estimates. But this purity comes at a cost: RCTs ignore prior knowledge, and units which are likely to benefit from treatment have the same probability of treatment as low-benefit units. In practice the experimenter must balance

- **Getting it now**, and obtaining short-term value from likely-high-return units
- **Getting it right**, and making statistically correct long-term decisions

We study **tie-breaker designs** [3] when the experimenter seeks to measure a heterogeneous effect but faces uncertainty [2] regarding the effect's true value and has preferences over both in-experiment and post-experiment payoffs. We show that RCTs are generically nonoptimal. The optimal experiment tends toward an RCT as the experimenter increasingly prefers getting it right, or as the experiment's outcome becomes uncertain, and tends toward a Regression Discontinuity Design (RDD) as the experimenter increasingly prefers getting it now, or as the experiment's outcome becomes certain.

Tie-breaker designs

A **tie-breaker design** with respect to mean-zero characteristic X_i is defined by $\Delta \in [0,1]$ and $\chi \in \{-1,1\}$,

$$\Pr(W_i = \chi | X_i) = \begin{cases} 0 & \text{if} \quad X_i < -\Delta, \\ \frac{1}{2} & \text{if} \quad -\Delta \le X_i \le \Delta, \\ 1 & \text{if} \quad \Delta < X_i \end{cases}$$

Tie-breaker designs nest classical experimental designs:

- When $\Delta = 0$, the experiment is a **regression discontinuity design** (RDD)
- When $\Delta = 1$, the experiment is a **randomized controlled trial** (RCT)

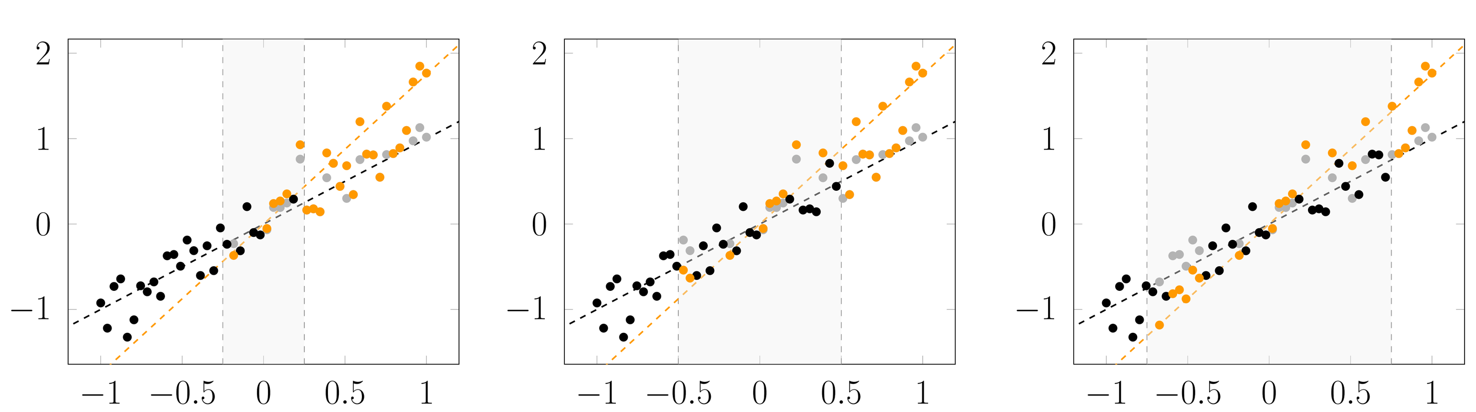

Figure. Forward tie-breaker outcomes for varying tie-breaker widths.

References

[1] Carlos Cinelli, Avi Feller, Guido Imbens, Edward Kennedy, Sara Magliacane, and Jose Zubizarreta. Challenges in statistics: A dozen challenges in causality and causal inference, 2025.

[2] Frank Hyneman Knight. *Risk, uncertainty and profit*, volume 31. Houghton Mifflin, 1921.

[3] Art B Owen and Hal Varian. Optimizing the tie-breaker regression discontinuity design. 2020.

Statistical model

- There are $2n+1$ units, defined by $X_i \in \{-\frac{n}{n}, -\frac{n-1}{n}, \ldots, \frac{n-1}{n}, \frac{n}{n}\}$
- There is a **baseline effect** β and a heterogeneous **treatment effect** γ
- $W_i \in \{-1, 1\}$ determines whether a unit is treated
- The experimenter's payoff from unit i, defined by (X_i, W_i), is
  $$Y_i = \beta_0 + \beta_1 X_i + \gamma X_i W_i + \varepsilon_i$$
  We assume ε_i is independent of X_i and W_i, and $\mathbb{E}[\varepsilon_i] = 0$ and $\text{Var}(\varepsilon_i) = \sigma^2$

Experimental model

- The experimenter runs a **tie-breaker design** defined by (Δ, χ)
- The experiment's **ship probability** is $P(\Delta, \chi; \gamma)$
- The experimenter has discount rate δ
- Conditional on γ, the experimenter's **lifecycle utility** is
  $$\begin{aligned} V(\Delta, \chi; \gamma) &= (1-\delta) \times \text{In-XP} + \delta \times \text{Post-XP} \\ \text{In-XP} &= \Pi(\Delta, \chi; \gamma) \\ \text{Post-XP} &= P(\Delta, \chi; \gamma)\, \Pi(\text{ship}; \gamma) + (1 - P(\Delta, \chi; \gamma))\, \Pi(\neg\text{ship}) \end{aligned}$$

Decision model

- Ship probability is determined by a normal hypothesis test
- Experimenter knows only that $\gamma \in \Gamma = [\underline{\gamma}, \overline{\gamma}]$
- Experimenter **minimizes maximum lifecycle regret**,
  $$\min_{\Delta,\chi} \max_{\Delta',\chi',\gamma'} V(\Delta', \chi'; \gamma') - V(\Delta, \chi; \gamma')$$

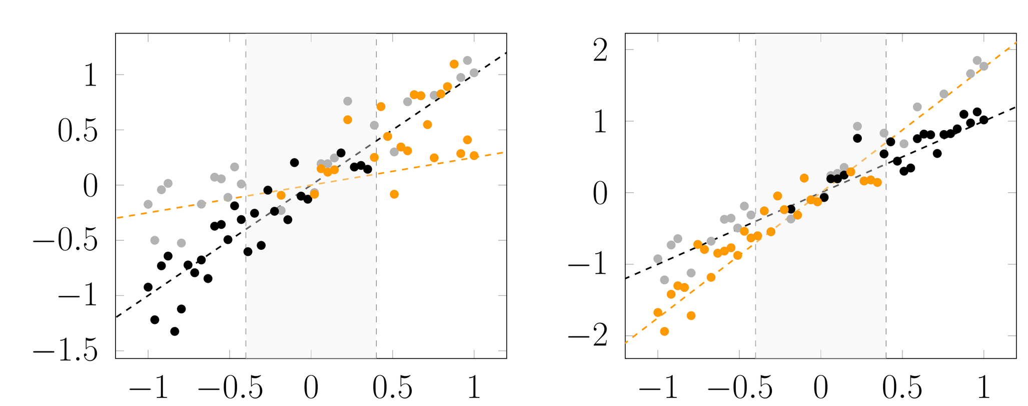

Figure. Regret-maximizing γ for forward (left) and reverse experiments (right).

Note. The paper model is more general. This model is sufficient for key insights.

Illustrative result: Impatient experimenter

When the experimenter is impatient, the lifecycle regret minimization problem is

$$\Delta^\star, \chi^\star \in \arg\min_{\Delta,\chi} \left[\max_{\Delta',\chi',\gamma'} \Pi(\Delta', \chi'; \gamma') - \Pi(\Delta, \chi; \gamma') \right]$$

Worst-case regret depends on the experimental design:

- In a forward experiment ($\chi = 1$), the worst case is $\gamma = \underline{\gamma}$
- In a backward experiment ($\chi = -1$), the worst case is $\gamma = \overline{\gamma}$

Proposition. When $\delta = 0$, the optimal tie-breaker design is:

- An RDD in the "direction" of γ when $\underline{\gamma} \overset{\text{sgn}}{=} \overline{\gamma}$
- An RCT when $-\underline{\gamma} = \overline{\gamma}$
- An interior tie-breaker design otherwise.

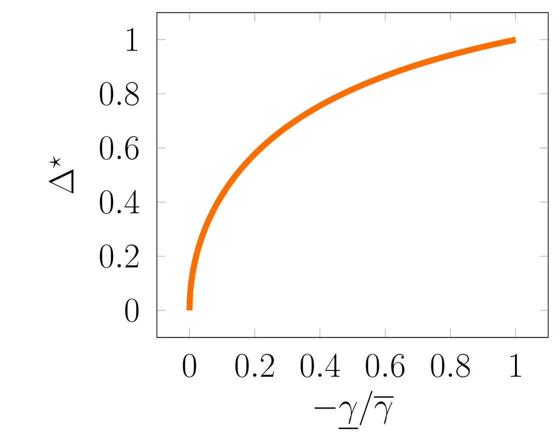

Illustrative result: Known treatment effect

When the treatment effect is known, regret minimization is utility maximization.

Proposition. When $\Gamma = \{\gamma\}$, the width of the optimal tie-breaker design:

- Increases with future preference δ
- Eventually decreases with statistical significance
- Eventually decreases with $n\gamma^2/\sigma^2$

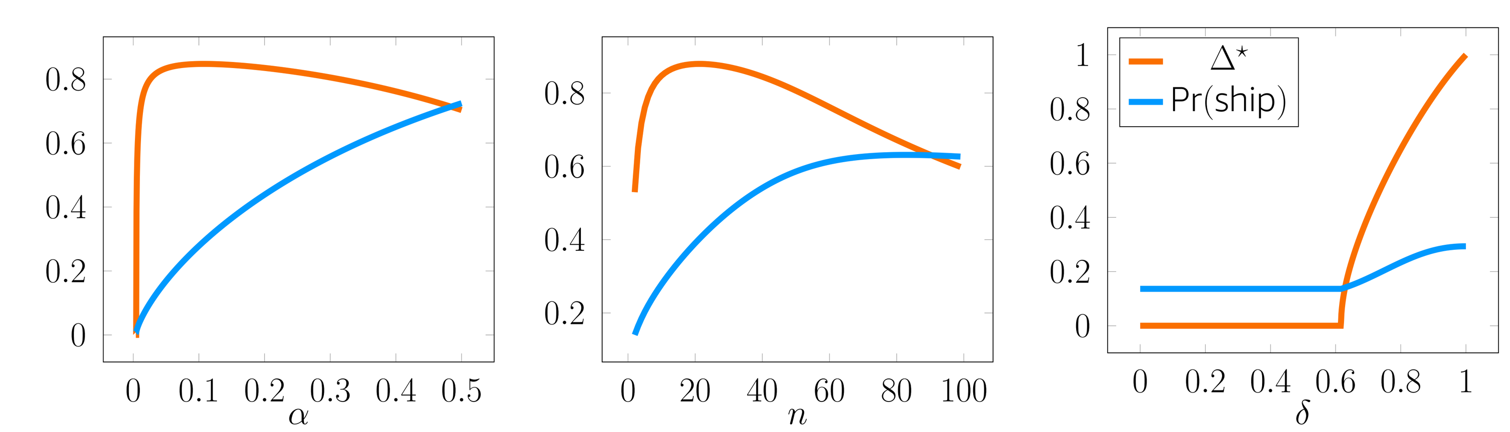

☆ General result ☆

Proposition. The breadth of the optimal tiebreaker design $\Delta^\star$ increases with patience, and decreases with statistical significance (when significance is sufficiently large). An RDD may be optimal, but **an RCT is never optimal unless $\delta = 0$.**

aprvlal@amazon.com and kwdwrd@amazon.com

apoorvalal.github.io and kylewoodward.com