Data-Driven Mechanism Design: Jointly Eliciting Preferences and Information

Dirk Bergemann[1] Marek Bojko[1] Paul Dütting[2]
Renato Paes Leme[2] Haifeng Xu[3,2] Song Zuo[2]

[1]Yale

[2]Google Research

[3]University of Chicago

Introduction

- In many economic environments (particularly online and digital), agents have:
  - Private preferences
  - Private information about a common payoff-relevant state
- Efficient implementation is difficult: tension of separating information from preferences
- We build a framework for mechanism design that utilizes often abundant additional data
- Our mechanisms condition transfers on this additional information about the state
- This *data-driven* approach enables implementation of efficient allocations

Motivating Examples

Sponsored search auctions

- Advertisers with private info on values (preferences) and click-through rates (state)
- Additional data revealed through clicks on ads, outcomes more generally

Motivating Examples

Sponsored search auctions

- Advertisers with private info on values (preferences) and click-through rates (state)
- Additional data revealed through clicks on ads, outcomes more generally

Emerging settings with LLMs

- Content providers privately know desired outcomes (preferences) and user tastes (state)
- Additional data revealed through user engagement and platform feedback

Motivating Examples

Sponsored search auctions

- Advertisers with private info on values (preferences) and click-through rates (state)
- Additional data revealed through clicks on ads, outcomes more generally

Emerging settings with LLMs

- Content providers privately know desired outcomes (preferences) and user tastes (state)
- Additional data revealed through user engagement and platform feedback

Supply Chain Coordination

- Vendors have information about private costs (preferences) and aggregate demand (state)
- Additional data revealed through sales by retailer

Preferences

- $N \equiv \{1, \ldots, n\}$ agents
- Utility
$$U_i(x, \omega, \theta_i, t_i) \equiv u_i(x, \omega, \theta_i) + t_i$$
with continuous u_i
- Outcome $x \in X$ (compact)
- State $\omega \in \Omega$ (compact)
- Preference type $\theta_i \in \Theta_i$

Information

- Signal $s_i \in \mathcal{S}_i$
- θ_i independent of s_j and ω, $\forall i, j$
- (θ_i, s_i) jointly referred to as i's type
- Commonly known *information structure* $\mu \in \Delta(\Omega \times \mathcal{S})$ (relaxed in the full paper)
- Expected payoff

$$v_i(x, \theta_i, s) \equiv \mathbb{E}[u_i(x, \omega, \theta_i)|s]$$

Efficiency

Definition (Efficient allocation)

$x : \Theta \times \mathcal{S} \to X$ is efficient if, $\forall \theta, s$,

$$x(\theta, s) \in \arg\max_{x \in X} \sum_{i \in N} v_i(x, \theta_i, s)$$

Denote an efficient allocation by x^*.

- Posterior notion: after revelation of types but allowing for residual uncertainty about ω
- The goal is to implement this object

Mechanisms and Implementation

Message-Driven Mechanisms

Definition (Message-driven direct mechanism)

(x, t): allocation

$$x : \Theta \times \mathcal{S} \to X$$

and transfers

$$t : \Theta \times \mathcal{S} \to \mathbb{R}^n$$

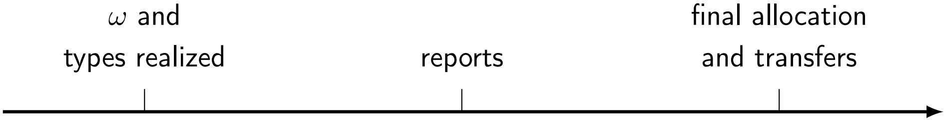

Ex-Post Implementation

Definition (Ex-post implementation)

(x, t) permits implementation in ex-post equilibrium if $\forall i, s, \theta$:

$$v_i(x(\theta, s), \theta_i, s) + t_i(\theta, s) \geq$$
$$v_i(x(\theta'_i, \theta_{-i}, s'_i, s_{-i}), \theta_i, s) + t_i(\theta'_i, \theta_{-i}, s'_i, s_{-i})$$

$\forall s'_i, \theta'_i$.

- After the uncertainty about types is resolved, no loss from having reported truthfully
- Ex-post with respect to agent types but allowing for residual uncertainty about ω
- Robust to belief misspecification

Impossibility

- Implementation is impossibile with interdependent values if signals are multi-dimensional (Maskin, 1992; Dasgupta and Maskin, 2000; Jehiel and Moldovanu, 2001; Jehiel et al., 2006)
- Even with one-dimensional signals, here types remain multi-dimensional
- Implementation is impossible under conditions sufficient for implementation in common value settings

Data-Driven Mechanisms

Estimator of ω

Assumption (Estimator)

There is an estimator $\widehat{\omega}$ of ω such that:

1. *The distribution of $\widehat{\omega}$ conditional on ω is commonly known.*
2. *$\widehat{\omega}$ is independent of agents' signals conditional on ω.*

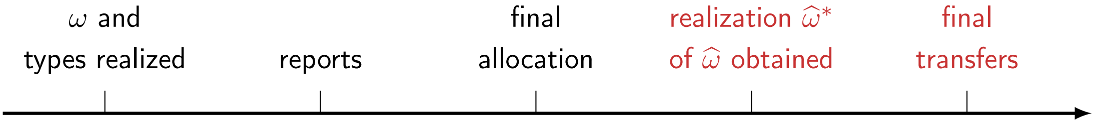

Data-Driven Mechanisms

Definition (Data-driven direct mechanism)

(x, t): allocation

$$x : \Theta \times \mathcal{S} \to X$$

and transfers

$$t : \Theta \times \mathcal{S} \times \Omega \to \mathbb{R}^n$$

- Ex-post implementation defined based on expected transfers

$$\bar{t}_i(\theta', s', s) \equiv \mathbb{E}[t_i(\theta', s', \widehat{\omega})|s]$$

Data-Driven VCG

Definition (Data-driven VCG)

(x^*, t) is a *data-driven VCG mechanism* if x^* is efficient and $\forall i$,

$$t_i(\theta, s, \widehat{\omega}^*) = h_i(\theta_{-i}, s_{-i}, \widehat{\omega}^*) + \sum_{j \neq i} u_j(x^*(\theta, s), \widehat{\omega}^*, \theta_j),$$

for an integrable function h_i of others' reports and a realization $\widehat{\omega}^*$ of $\widehat{\omega}$.

- Transfers do not depend on reported signals other than through the allocation
- Conditional on an allocation, does not "fix" beliefs

Implementation

Theorem (Full revelation of the state)

If $\widehat{\omega} = \omega$, every data-driven VCG mechanism permits implementation in ex-post equilibrium.

Implementation

Theorem (Full revelation of the state)

If $\widehat{\omega} = \omega$, every data-driven VCG mechanism permits implementation in ex-post equilibrium.

- Suppose agents $-i$ report truthfully; i evaluates the transfer in expectation:

$$h_i(\theta_{-i}, s_{-i}) + \sum_{j \neq i} v_j(x^*(\theta_i', \theta_{-i}, s_i', s_{-i}), \theta_j, s)$$

- i's net utility becomes

$$h_i(\theta_{-i}, s_{-i}) + \sum_{j \in N} v_j(x^*(\theta_i', \theta_{-i}, s_i', s_{-i}), \theta_j, s)$$

Maximized at $(\theta_i', s_i') = (\theta_i, s_i)$ by efficiency.

Noisy data

- Designer likely has only a noisy estimate
- We consider estimators with two desirable properties:
  - Unbiased $\widehat{\omega} \rightarrow$ exact implementation when utilities are affine in the state
  - Consistent $\widehat{\omega} \rightarrow$ approximate implementation (exact in the limit)

ϵ-Ex-Post Implementation

Definition (ϵ-ex-post implementation)

Fix $\epsilon \geq 0$. (x, t) permits implementation in ϵ-ex-post equilibrium if $\forall i, s, \theta$:

$$v_i(x(\theta, s), \theta_i, s) + \bar{t}_i(\theta, s, s) + \epsilon \geq$$
$$v_i(x(\theta_i', \theta_{-i}, s_i', s_{-i}), \theta_i, s) + \bar{t}_i(\theta_i', \theta_{-i}, s_i', s_{-i}, s)$$

$\forall s_i', \theta_i'$.

Implementation Continuity

Sequence of estimators $\{\widehat{\omega}_m\}_m$ is *uniformly consistent* if

$$\forall \epsilon > 0 : \quad \lim_{m \to \infty} \sup_{\omega \in \Omega} \mathbb{P}(d(\widehat{\omega}_m, \omega) > \epsilon \mid \omega) = 0.$$

Theorem (Consistent estimators)

Suppose u_i is Lipschitz for each agent i. Fix a uniformly consistent sequence of estimators $\{\widehat{\omega}_m\}_m$. Then there is a non-negative sequence $\{\epsilon_m\}_m$, with $\epsilon_m \to 0$ as $m \to \infty$, such that every data-driven VCG mechanism for $\widehat{\omega}_m$ permits implementation in ϵ_m-ex-post equilibrium for every $m \in \mathbb{N}$.

Additional Results and Related Approaches

- Applications to our motivating examples:
  - Data-driven VCG mechanisms encompass per-click payment rules in ad auctions
  - In the LLM application, we introduce regularized data-driven VCG mechanisms that maximize total agent rewards while penalizing deviations from a reference distribution

Additional Results and Related Approaches

- Applications to our motivating examples:
  - Data-driven VCG mechanisms encompass per-click payment rules in ad auctions
  - In the LLM application, we introduce regularized data-driven VCG mechanisms that maximize total agent rewards while penalizing deviations from a reference distribution
- They illustrate the practical appeal of data-driven VCG over related approaches
- Mezzetti (2004): agents observe ex-post payoffs and report them in a 2^{nd} stage, with transfer for agent i given by $\sum_{j \neq i} u_j + h_i(\theta_{-i}, s_{-i}, u_{-i})$
  - Data-driven VCG require one round of communication and less information
  - Data-driven VCG are also less fragile and can be "pivotal"
- Riordan and Sappington (1988): use the estimator in Crémer&McLean-style transfers
  - Data-driven VCG transfers are bounded uniformly across environments

Conclusion

- We offered an approach to mechanism design that harnesses the natural flow of information in digital environments
- By conditioning transfers on post-allocation data, we showed how to achieve implementation even in challenging multidimensional settings
- Our framework provides a foundation for designing practical mechanisms in modern applications where rich feedback data is readily available

Thank you!