Robust Pricing in Dynamic Mechanism Design

July, 2020 @ ICML

Yuan Deng, Duke University => Google Research

Sébastien Lahaie, Google Research

Vahab Mirrokni, Google Research

Online Advertising

- The popularity of selling online advertising opportunities via ***repeated auctions***
    - *the set of advertisers is the same*
    - *the ad slots are different*
        - *users / ad locations / timing*

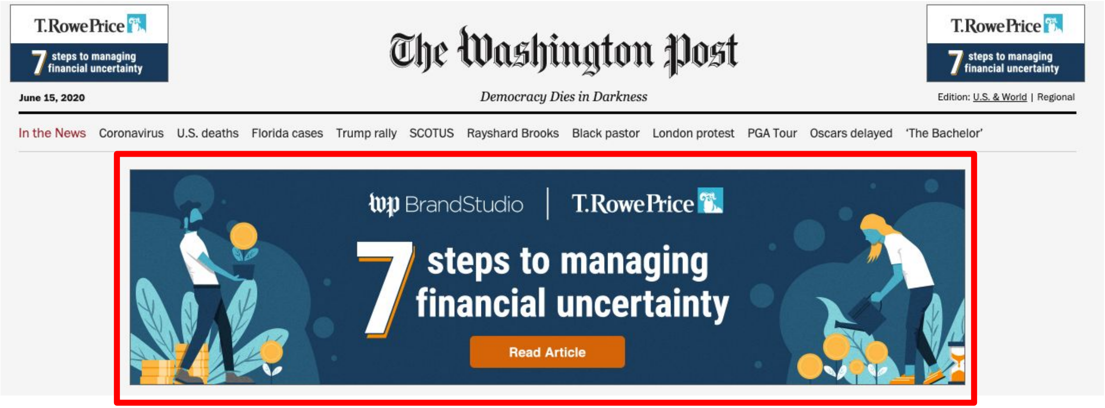

- A standard approach to ***monetize*** online web services;
    - generate *hundreds of billions of dollars of **revenue** annually*.

Dynamic Mechanism Design

- Selling online advertisements via repeated auctions inspires the research on ***dynamic mechanism design*** in the past decade [ADH 16, MPTZ 18]:

Dynamic Mechanism	Static Mechanism
• Mechanism ***depends*** on the history	• Mechanism ***ignores*** the history
For example,	For example,
• Dynamic reserve pricing	• Repeated second-price auctions

- Dynamic auctions open up the possibility of evolving the auctions across time to **boost *revenue***.
    - The ***revenue*** gap between ***dynamic*** and ***static*** mechanism can be ***arbitrarily large*** [PPPR 16]

Dynamic Mechanism Design

- Dynamic auctions open up the possibility of evolving the auctions across time to **boost *revenue***.
  - The revenue gap between ***dynamic*** and ***static*** mechanism can be ***arbitrarily large*** [PPPR 16]

However

- Dynamic mechanism ***complicates*** the buyer's ***long-term incentive***
  - the buyers' ***current*** bids may change the ***future*** mechanism
  - e.g., shading the bids in past may lower the reserve in the future

To align the buyer's incentives, ***perfect distributional knowledge*** is usually required

- Such a reliance limits the application of dynamic mechanism design in practice
  - The seller may only have access to ***estimated*** distributions
  - The seller may need to ***learn*** the distributions

Our Contribution

To align the buyer's incentives, ***perfect distributional knowledge*** is usually required

- We develop a framework for robust **dynamic mechanism design**
  - its **revenue performance** is robust against
    - **estimation error** on the valuation distributions and the buyer's **strategic behavior**
    - **i.e., the revenue loss can be bounded by the estimation error**

- We apply our framework to **contextual auctions**
  - where the seller needs to learn the valuation distributions
  - obtain **the first**, to the best of our knowledge, **no-regret dynamic pricing policy against revenue-optimal dynamic mechanism that has perfect distributional knowledge**

Bayesian Dynamic Environment

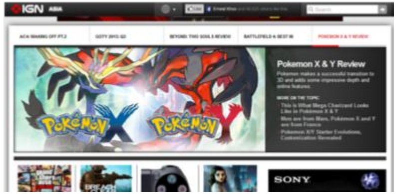

$v_1 \sim F_1$ $v_2 \sim F_2$ $v_3 \sim F_3$

1. One item arrives at stage t
2. The buyer observes private v_t drawn *independently* from F_t
3. The buyer submits bid b_t to the seller
4. The seller only knows an estimated distribution F'_t, and he will determine:
   - *Allocation probability* $x_t(b_{(1,t)}, F'_{(1,T)})$ and *Payment* $p_t(b_{(1,t)}, F'_{(1,T)})$

- The buyer's utility is $u_t(b_{(1,t)}, F'_{(1,T)}) = v_t \cdot x_t(b_{(1,t)}, F'_{(1,T)}) - p_t(b_{(1,t)}, F'_{(1,T)})$
  - **additive across items**

Impatient Buyer & Imperfect Distributional Knowledge

- We assume the buyer is **impatient**
  - she discounts her future utility at a factor γ
  - it is impossible to obtain a no-regret policy for a patient buyer [ARS 13]

- **Imperfect distributional knowledge (estimation error)**
  - The estimation error is Δ if there exists a coupling between a random draw $\boldsymbol{v_t}$ drawn *independently* from $\boldsymbol{F_t}$ and $\boldsymbol{v'_t}$ drawn *independently* from $\boldsymbol{F'_t}$ such that

$$v_t = v'_t + \epsilon_t \text{ with } \epsilon_t \in [-\Delta, \Delta]$$

  - Intuitively, samples from the estimated distribution have **a bounded bias**
  - This measurement is **consistent** with the model of contextual auctions

approximate Dynamic Incentive Compatibility

exact dynamic-IC notion [MPTZ 18] (for long-term utility maximizers):

- For every stage, reporting truthfully is an optimal strategy
  - assuming the buyer plays ***optimally (to maximize her cumulative utility)*** in the future

---

- Impossible to achieve exact dynamic-IC without perfect distributional knowledge
  - with a non-trivial dynamic mechanism

approximate dynamic-IC notion:

- For every stage, reporting **a bid close to her true valuation** is an optimal strategy
  - assuming the buyer plays ***optimally (to maximize her cumulative utility)*** in the future

Challenges

- Impossible to achieve exact dynamic-IC
  - Attempt to achieve approximate dynamic-IC
    - <u>How to bound the magnitude of the misreport for dynamic mechanisms?</u>

- Revenue performance
  - Future mechanism depends on the buyer’s reports in the past
    - A misreport could change the structure of future mechanisms and their revenues
    - <u>How to bound the revenue loss due to misreport for dynamic mechanisms?</u>

- We propose a **framework** to **robustify** dynamic mechanism so that
  - the magnitude of misreport can be bounded **by the estimation errors**
  - the revenue loss due to misreport can be bounded **by the magnitude of misreport**

=> the revenue loss against strategic buyers can be bounded **by the estimation errors**

Bound the Misreport

Our framework is based on the **bank account mechanism** [MPTZ 18]

- it is without loss of generality to consider bank account mechanism: any dynamic mechanism can be reduced to a bank account mechanism without loss of any **revenue** or **welfare**

- Bank account mechanism enjoys a property called **utility independence**
    - the buyer's ***expected utility*** (under truthful bidding) at a stage is ***independent of the history***
    - i.e., the buyer's ***historical bids*** have ***no impact*** on her ***future expected utility***

    - **Remark**: although the expected utility is the same, the mechanism can be different

Utility Independence (Example) [PPPR, SODA'16]

Stage 1

$$\Pr\left[v_1 = 2^i\right] = \frac{1}{2^i}, \quad \text{for } i \in \{1, \cdots, n\}$$

- Run the **first-price** auction
  - bid $\mathbf{b_1}$; get the item and pay $\mathbf{b_1}$
- Buyer's utility under valuation $\mathbf{v_1}$

$$v_1 - b_1$$

Stage 2

$$\Pr\left[v_2 = 2^j\right] = \frac{1}{2^j}, \quad \text{for } j \in \{1, \cdots, 2^n\}$$

- Give the item **for free** with prob. $\mathbf{b_1/2^n}$
  - no matter what $\mathbf{b_2}$ is
- Buyer's expected utility

$$E_{v_2}\left[v_2 \cdot \frac{b_1}{2^n}\right] = E_{v_2}[v_2] \cdot \frac{b_1}{2^n} = b_1$$

Dynamic-IC and **Revenue is n**

- (discrete) *equal revenue distributions* for both stages
  - Selling separately using the ***optimal static*** mechanism gives **revenue 2** per stage

Utility Independence (Example) [PPPR, SODA'16]

Stage 1

$$\Pr\left[v_1 = 2^i\right] = \frac{1}{2^i}, \quad \text{for } i \in \{1, \cdots, n\}$$

- Run the **first-price** auction
  - bid $\mathbf{b_1}$; get the item and pay $\mathbf{b_1}$
- Buyer's utility under valuation $\mathbf{v_1}$

$$v_1 - b_1$$

Stage 2

$$\Pr\left[v_2 = 2^j\right] = \frac{1}{2^j}, \quad \text{for } j \in \{1, \cdots, 2^n\}$$

- Give the item **for free** with prob. $\mathbf{b_1/2^n}$
  - no matter what $\mathbf{b_2}$ is
- Buyer's expected utility

$$E_{v_2}\left[v_2 \cdot \frac{b_1}{2^n}\right] = E_{v_2}[v_2] \cdot \frac{b_1}{2^n} = b_1$$

depend on Stage 2

~~**Dynamic IC**~~ and **Revenue is n**

- (discrete) *equal revenue distributions* for both stages
  - Selling separately using the ***optimal static*** mechanism gives **revenue 2** per stage

Payment Realignment

Stage 1

$$\Pr\left[v_1 = 2^i\right] = \frac{1}{2^i}, \quad \text{for } i \in \{1, \cdots, n\}$$

- Run the **first-price** auction
  - bid $\mathbf{b_1}$; get the item and pay $\mathbf{b_1}$
- Buyer's utility under valuation $\mathbf{v_1}$

$$v_1 - b_1$$

Stage 2

$$\Pr\left[v_2 = 2^j\right] = \frac{1}{2^j}, \quad \text{for } j \in \{1, \cdots, 2^n\}$$

- Give the item **for free** with prob. $\mathbf{b_1/2^n}$
  - no matter what $\mathbf{b_2}$ is
- Buyer's expected utility

$$E_{v_2}\left[v_2 \cdot \frac{b_1}{2^n}\right] = E_{v_2}[v_2] \cdot \frac{b_1}{2^n} = b_1$$

Dynamic-IC and **Revenue is n**

- (discrete) *equal revenue distributions* for both stages
  - Selling separately using the ***optimal static*** mechanism gives **revenue 2** per stage

Payment Realignment

Stage 1

$$\Pr\left[v_1 = 2^i\right] = \frac{1}{2^i}, \quad \text{for } i \in \{1, \cdots, n\}$$

- Run the ~~[first-price]~~ **[give-for-free]** auction
  - bid $\mathbf{b_1}$; get the item ~~and pay $\mathbf{b_1}$~~
- Buyer's utility under valuation $\mathbf{v_1}$

$$v_1 \cancel{- b_1}$$

Stage 2

$$\Pr\left[v_2 = 2^j\right] = \frac{1}{2^j}, \quad \text{for } j \in \{1, \cdots, 2^n\}$$

- Give the item **for free** with prob. $\mathbf{b_1/2^n}$
  - no matter what $\mathbf{b_2}$ is
- Buyer's expected utility

$$E_{v_2}\left[v_2 \cdot \frac{b_1}{2^n}\right] = E_{v_2}[v_2] \cdot \frac{b_1}{2^n} = b_1$$

$\boxed{\textbf{Dynamic-IC} \quad \text{and} \quad \textbf{Revenue is n}}$

- (discrete) *equal revenue distributions* for both stages
  - Selling separately using the ***optimal static*** mechanism gives **revenue 2** per stage

Payment Realignment

Stage 1

$$\Pr\left[v_1 = 2^i\right] = \frac{1}{2^i}, \quad \text{for } i \in \{1, \cdots, n\}$$

- Run the ~~[first-price]~~ **[give-for-free]** auction
  - bid $\mathbf{b_1}$; get the item ~~and pay $\mathbf{b_1}$~~
- Buyer's utility under valuation $\mathbf{v_1}$

$$v_1 \cancel{- b_1}$$

Stage 2

$$\Pr\left[v_2 = 2^j\right] = \frac{1}{2^j}, \quad \text{for } j \in \{1, \cdots, 2^n\}$$

- Charge $\mathbf{b_1}$
- Give the item ~~[for free]~~ with prob. $\mathbf{b_1/2^n}$
  - no matter what $\mathbf{b_2}$ is

Buyer's expected utility

$$\boxed{E_{v_2}\left[v_2 \cdot \frac{b_1}{2^n}\right] = E_{v_2}[v_2] \cdot \frac{b_1}{2^n} = b_1 - b_1 = 0}$$

independent *of Stage 2 :)*

Dynamic-IC and **Revenue is n**

History UI

- (discrete) *equal revenue distributions* for both stages
  - Selling separately using the ***optimal static*** mechanism gives **revenue 2** per stage

Utility Independence

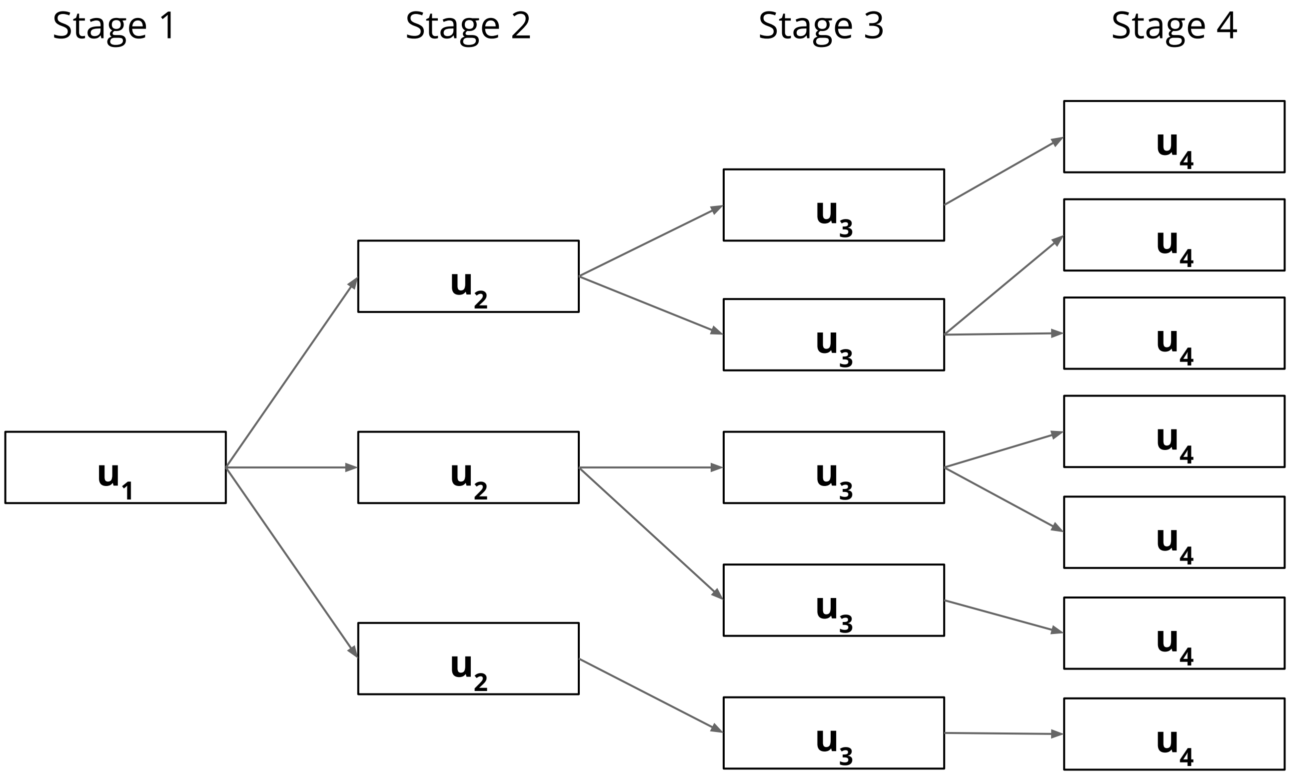

Bound the Misreport

- Bank account mechanism enjoys a property called **utility independence**
  - the buyer's ***expected utility*** at a stage is ***independent of the history***
  - i.e., the buyer's ***historical bids*** have ***no impact*** on her ***future expected utility***
  - **(under perfect distributional knowledge)**

Under imperfect distributional knowledge

- the buyer's ***expected utility*** at a stage is within a range related to **the estimation error**

approximate Utility Independence

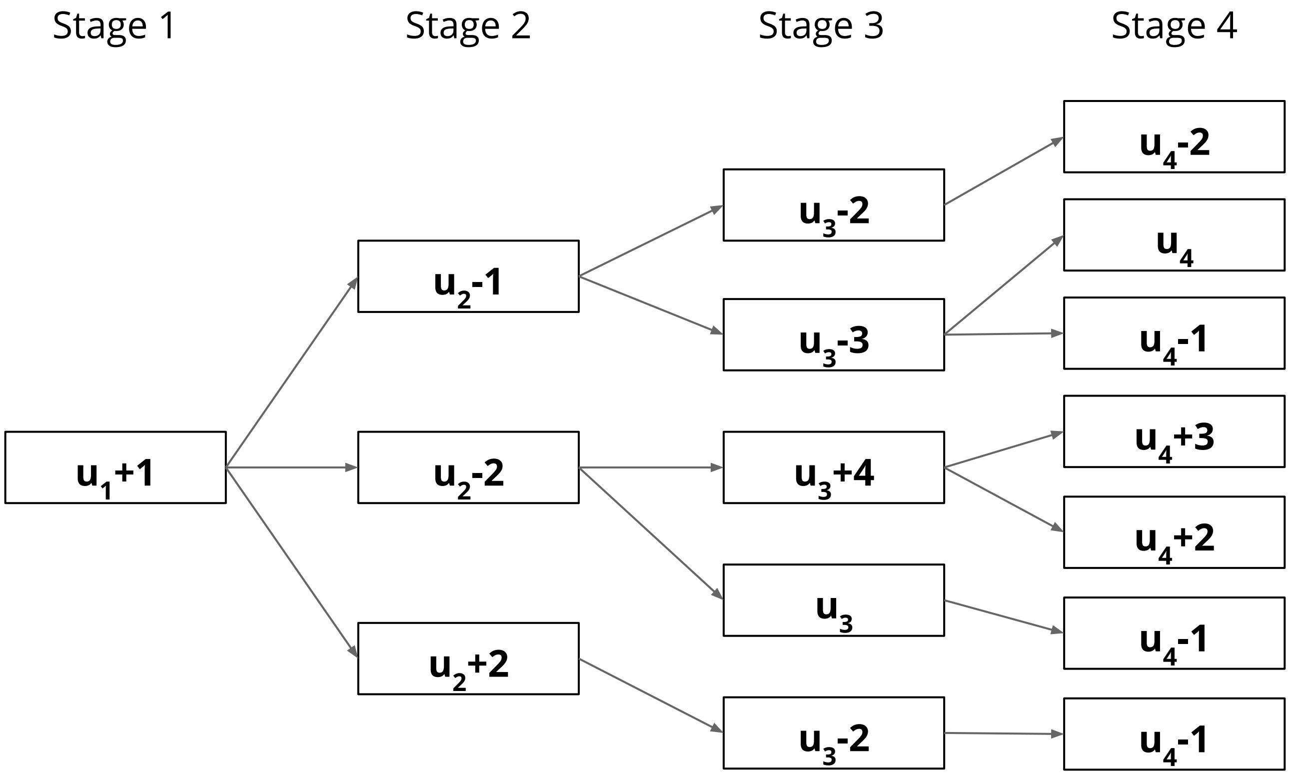

approximate Utility Independence

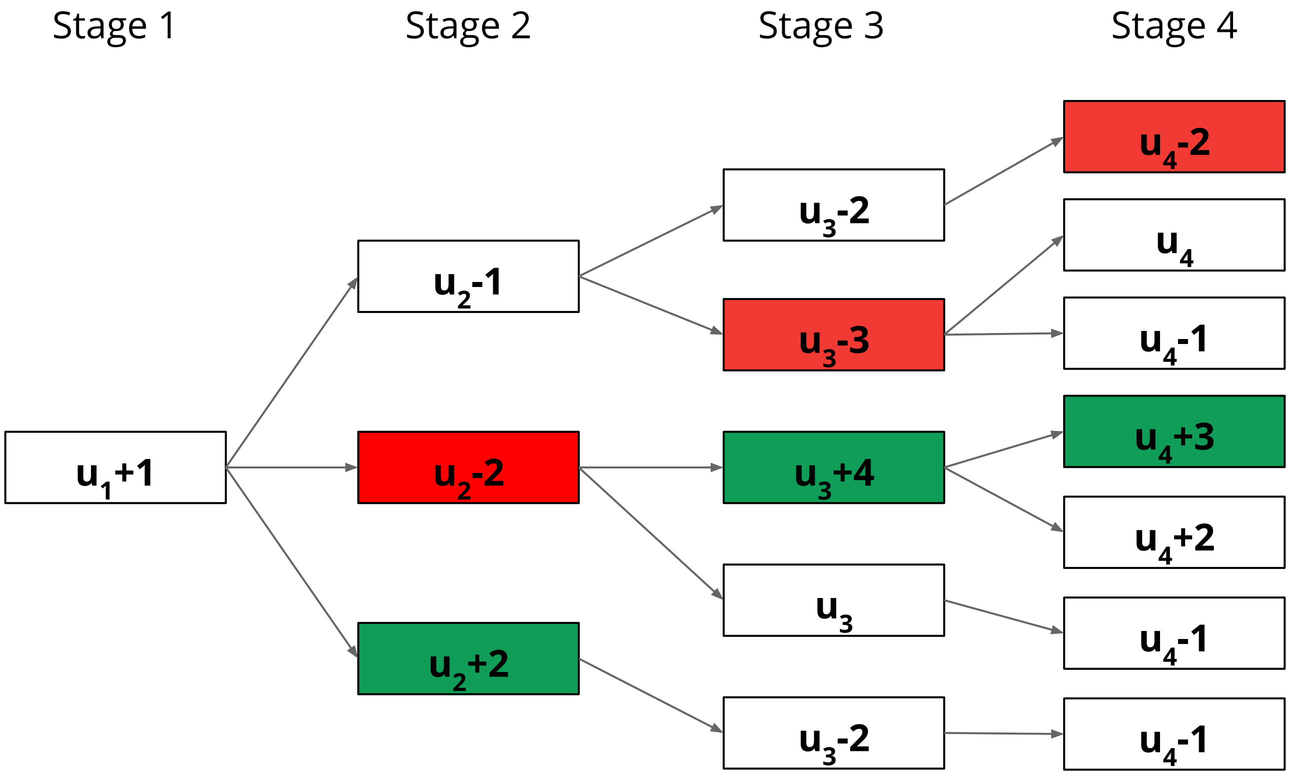

Bound the Misreport

- Bank account mechanism enjoys a property called **utility independence**
  - the buyer's ***expected utility*** at a stage is ***independent of the history***
  - i.e., the buyer's ***historical bids*** have ***no impact*** on her ***future expected utility***

Under imperfect distributional knowledge

- the buyer's ***expected utility*** at a stage is within a range related to **the estimation error**
- **so that** the buyer's utility gain at this stage from misreporting in the past is at most the range

High-level idea [GJM19]**:** create punishment for misreporting

- Mix the dynamic mechanism with **a random posted-price auction**
  - where a take-it-or-leave-it price is randomly drawn
  - **Property**: the larger the misreport is, the larger the utility loss would be

Bound the Revenue Loss

Extensively exploit the structure of bank account mechanisms

- Develop **new tools** for analyzing bank account mechanisms:
  - new ways to **edit** and **concatenate** bank account mechanisms for robustification
    - change the dynamics of the mechanism
    - while preserve the bank account structure
  - a program to **compute the revenue performance** with strategic buyers even when the distributional information is not perfect
    - leads to bounds on revenue loss due to misreport

- With tools at hand
  - Develop bank account mechanisms whose **revenue is robust against misreport**
  - i.e., the revenue loss can be bounded by the magnitude of the misreport

Challenges

- Impossible to achieve exact dynamic-IC
  - Attempt to achieve approximate dynamic-IC
    - How to bound the magnitude of the misreport for dynamic mechanisms?

- Revenue performance
  - Future mechanism depends on the buyer's reports in the past
    - A misreport could change the structure of future mechanisms and their revenues
    - How to bound the revenue loss due to misreport for dynamic mechanisms?

- We propose a **framework** to **robustify** dynamic mechanism so that
  - the magnitude of the misreport can be bounded
    - **mix in random posted-price auctions**
  - the revenue loss due to misreport can be bounded
    - **revenue-robust dynamic mechanism**

Conclusion & Future Work

Summary:

- We develop a framework for robust dynamic mechanism design
  - revenue robust against estimation error on distribution and strategic behavior
- As an application, we obtain a no-regret dynamic pricing policy for contextual auctions

Future Work:

- Improve our bounds
  - better revenue loss bound of the framework
  - better no-regret bound for contextual auctions
  - lower bounds?
- Apply our framework to environments more general than contextual auctions

References

[**ARS13**] Kareem Amin, Afshin Rostamizadeh, Umar Syed. *Learning Prices for Repeated Auctions with Strategic Buyers*. **NeurIPS'13**.

[**ADH16**] Itai Ashlagi, Constantinos Daskalakis, and Nima Haghpanah. *Sequential mechanisms with expost participation guarantees*. **EC'16**

[**PPPR16**] Christos Papadimitriou, George Pierrakos, Christos-Alexandros Psomas, and Aviad Rubinstein. *On the complexity of dynamic mechanism design*. **SODA'16**

[**MPTZ18**] Vahab Mirrokni, Renato Paes Leme, Pingzhong Tang, and Song Zuo. *Non-clairvoyant dynamic mechanism design*. **EC'18, Econometrica**

[**GJM19**] Negin Golrezaei, Adel Javanmard, and Vahab Mirrokni. *Dynamic Incentive-Aware Learning: Robust Pricing in Contextual Auctions*. **NeurIPS'19**.