Non Clairvoyant Dynamic Mechanism Design

Vahab Mirrokni
(Google)

Renato Paes Leme
(Google)

Pingzhong Tang
(Tsinghua)

Song Zuo
(Tsinghua)

This talk in one slide

This talk in one slide

This talk in one slide

This talk in one slide

This talk in one slide

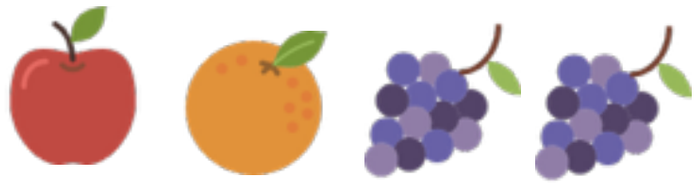

This talk in one slide

This talk in one slide

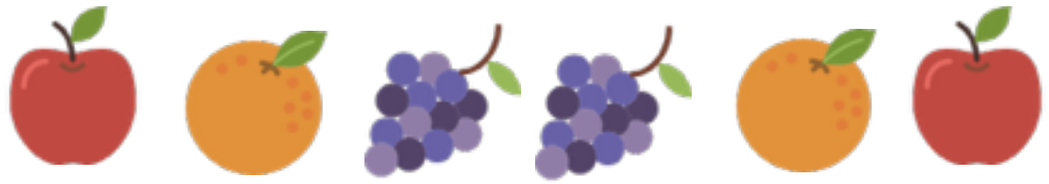

This talk in one slide

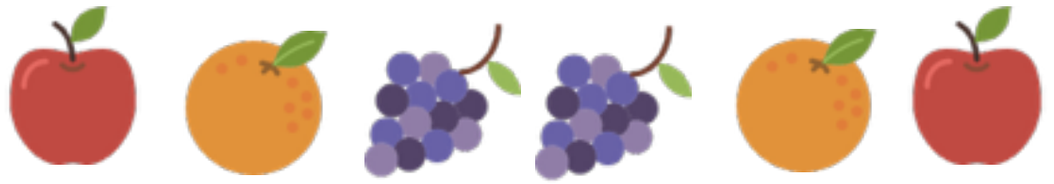

This talk in one slide

Clairvoyant seller

Sees present, past and future.

Non Clairvoyant seller

Remembers the past, but doesn't see the future.

Static seller

Has no memory of the past.

This talk in one slide

Clairvoyant seller

Sees present, past and future.

1 2 3

Non Clairvoyant seller

Remembers the past, but doesn't see the future.

1 2 3

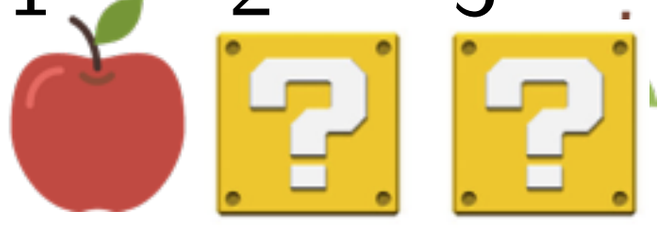

Static seller

Has no memory of the past.

1 2 3

This talk in one slide

Clairvoyant seller

Sees present, past and future.

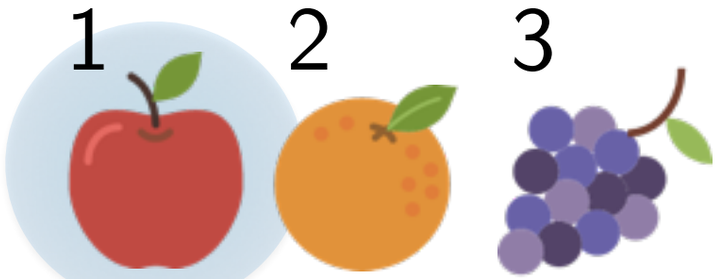

Non Clairvoyant seller

Remembers the past, but doesn't see the future.

Static seller

Has no memory of the past.

This talk in one slide

Clairvoyant seller

Sees present, past and future.

1 2 3

Non Clairvoyant seller

Remembers the past, but doesn't see the future.

1 2 3

Static seller

Has no memory of the past.

1 2 3

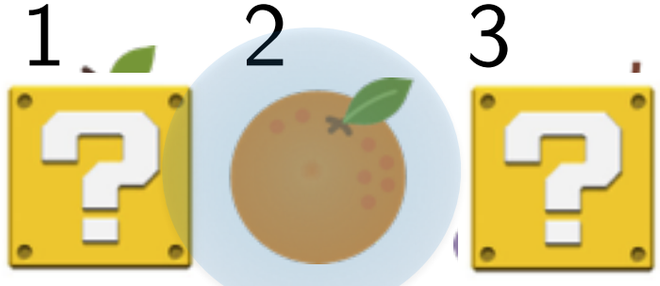

This talk in one slide

Clairvoyant seller

Sees present, past and future.

1 2 3

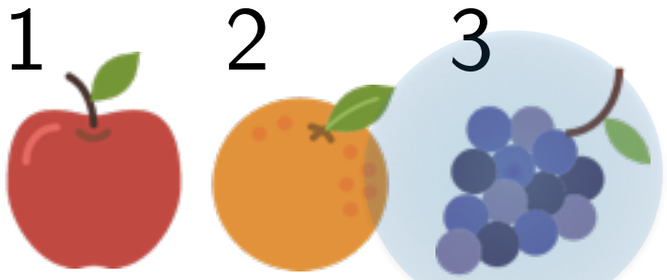

Non Clairvoyant seller

Remembers the past, but doesn't see the future.

1 2 3

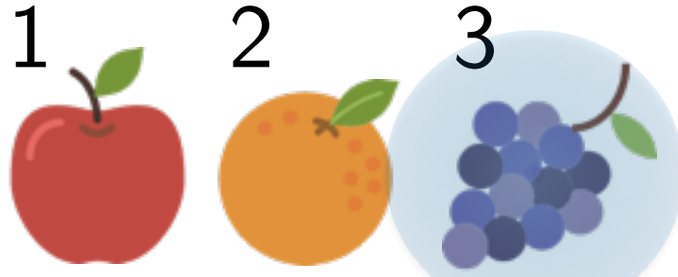

Static seller

Has no memory of the past.

1 2 3

This talk in one slide

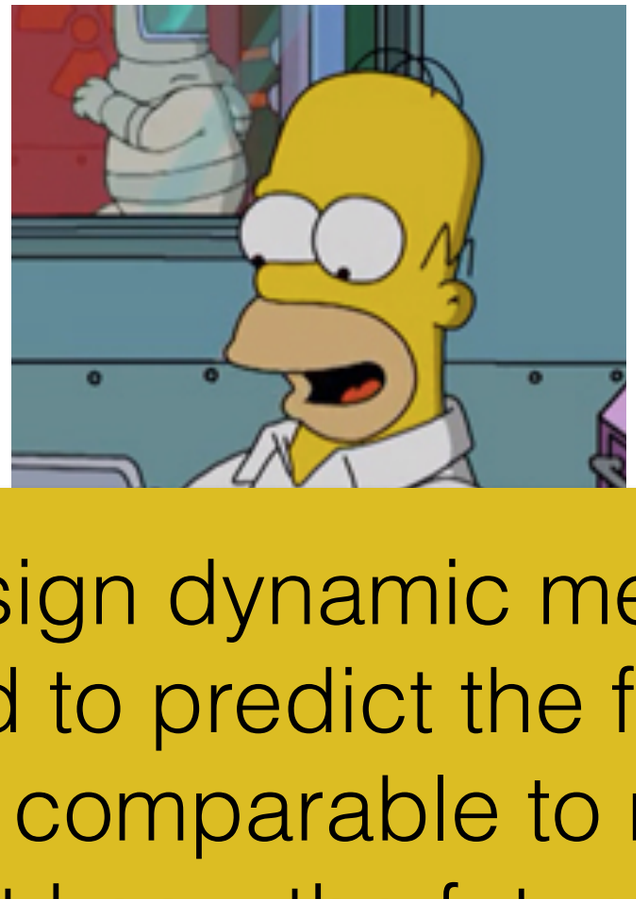

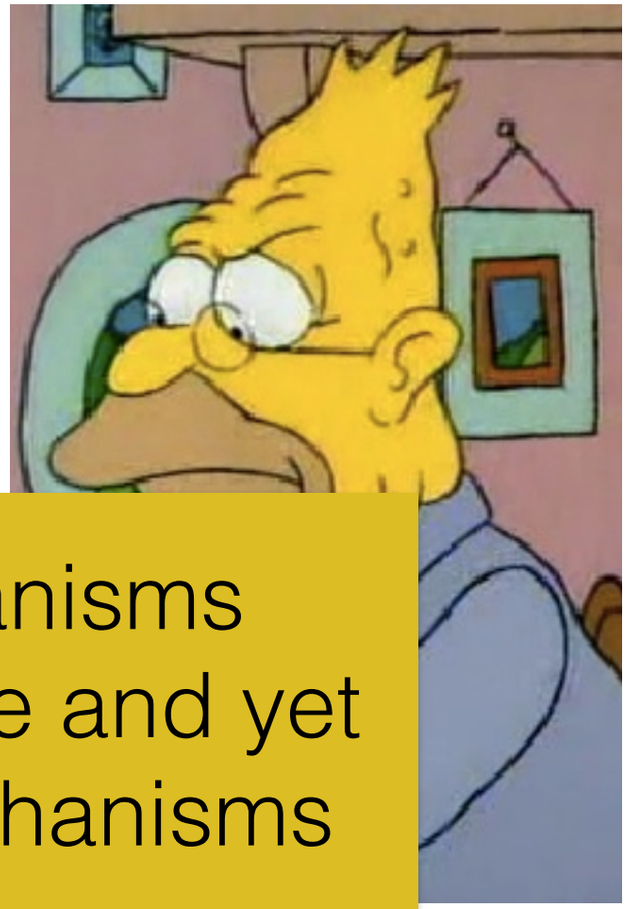

Can we design dynamic mechanisms that don't need to predict the future and yet have revenue comparable to mechanisms that know the future?

Clai **ler**

Sees present, past and future.

1 2 3

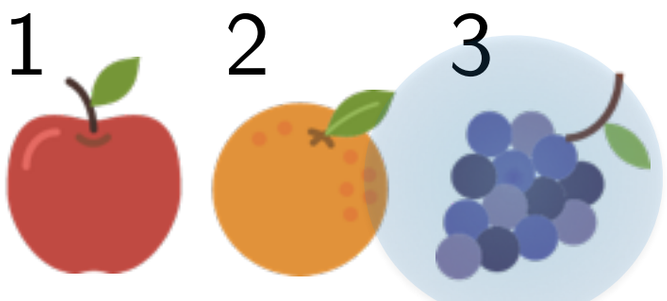

Remembers the past, but doesn't see the future.

1 2 3

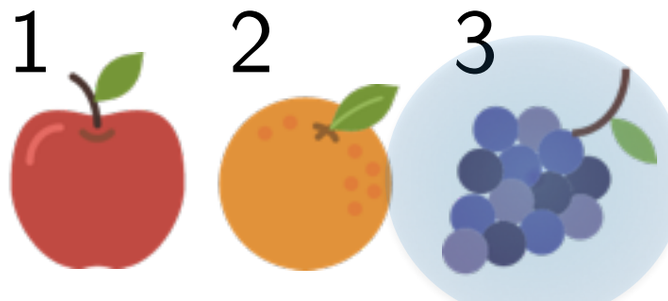

Has no memory of the past.

1 2 3

Problem Setup

- Items arrive in sequence.

Problem Setup

- Items arrive in sequence.
- One seller and many buyers: item sold when it arrives.

Problem Setup

- Items arrive in sequence.
- One seller and ~~many~~ one buyers: item sold when it arrives.

Problem Setup

- Items arrive in sequence.
- One seller and ~~many~~ one buyers: item sold when it arrives.
- Each item type has an distribution, e.g.

 $\sim F_1$  $\sim F_2$  $\sim F_3$

Problem Setup

- Items arrive in sequence.
- One seller and ~~many~~ one buyers: item sold when it arrives.
- Each item type has an distribution, e.g.

 $\sim F_1$  $\sim F_2$  $\sim F_3$

- The value for the t-th item is realized at time t.

Problem Setup

- Items arrive in sequence.
- One seller and ~~many~~ one buyers: item sold when it arrives.
- Each item type has an distribution, e.g.

 $\sim F_1$  $\sim F_2$  $\sim F_3$

- The value for the t-th item is realized at time t.
- Buyer's utility: $U = \sum_t v_t x_t - p_t$
  - allocation $x_t \in [0, 1]$
  - payment $p_t \geq 0$

Static Seller

- Sells one item at a time, without memory of the past or knowledge about the future : each auction is a standard Myersonian problem.
- Revelation principle: focus on mechanism specified as $x(v), p(v)$ and subject to two constraints:
  - Incentive compatibility:
    $$v = \text{argmax}_{\hat{v}} v \cdot x(\hat{v}) - p(\hat{v})$$
  - Individual rationality:
    $$v \cdot x(v) - p(v) \geq 0$$
- Simple recipe to $\max_{v \sim F}[p(v)]$
  e.g. if F = U[0,1], price at 1/2.

Dynamic Seller

- Mechanism is now described as a function of the reports in this and prev rounds: $x_t(v_1, v_2, \ldots, v_t), p_t(v_1, v_2, \ldots, v_t)$
- Linking independent problems together can improve revenue and efficiency [Jackson-Sonnenschein, Manelli-Vincent, Papadimitriou et al].
  - arbitrarily more revenue

Dynamic Seller

- Mechanism is now described as a function of the reports in this and prev rounds: $x_t(v_{1..t}), p_t(v_{1..t})$
- Linking independent problems together can improve revenue and efficiency [Jackson-Sonnenschein, Manelli-Vincent, Papadimitriou et al].
  - arbitrarily more revenue

Dynamic Seller

- Mechanism is now described as a function of the reports in this and prev rounds: $x_t(v_{1..t}), p_t(v_{1..t})$
- Linking independent problems together can improve revenue and efficiency [Jackson-Sonnenschein, Manelli-Vincent, Papadimitriou et al].
  - arbitrarily more revenue
- Incentive constraint: buyer is better of reporting his true type in each round.
- Individual rationality: buyer derives non-negative utility from the mechanism.

$$\sum_t v_t x_t - p_t \geq 0$$

Dynamic Incentive Compatibility

- Incentive constraint: buyer is better off reporting his true type in each round.

- Backwards induction: last round he is better off reporting his value conditioned on history:

$$v_T = \text{argmax}_{\hat{v}} v_T x_T(v_{1..T-1}\hat{v}) - p_T(v_{1..T-1}\hat{v})$$

Before to last period:

$$v_t = \text{argmax}\ u_{T-1}(v_{T-1}; v_{1..T-2}\hat{v}) + \mathbb{E}_{v_T} u_\tau(v_T; v_{1..T-2}\hat{v}v_T)$$

effect of my report in this round

expected effect of my report in next round

where $u_t(w; v_{1..t}) = w \cdot x_t(v_{1..t}) - p(v_{1..t})$

Dynamic Incentive Compatibility

- Incentive constraint: buyer is better off reporting his true type in each round.

- Dynamic Incentive Compatibility:

$$v_t = \text{argmax}\ u_t(v_t; v_{1..t-1}\hat{v}) + \mathbb{E}_{v_{t+1..T}}[\sum_{\tau=t+1}^{T} u_\tau(v_\tau; v_{1..t-1}\hat{v}v_{t+1..\tau})]$$

effect of my report in this round

expected effect of my report in future round

where $u_t(w; v_{1..t}) = w \cdot x_t(v_{1..t}) - p(v_{1..t})$

Clairvoyant Seller

- Revenue maximization $\max \mathbb{E}[\sum_t p_t(v_{1..t})]$ s.t IC and IR.
  - Solving this LP/DP requires knowledge about the future.

- Selling two apples, 🍎 $\sim U[0,1]$
  - Optimal static: price each at 1/2, optimal revenue is 0.5.
  - Improved dynamic:
    - elicit v_1 and sell first item for 1/2
    - charge $f = \min((v_1 - 1/2)^+, 3/8)$ to inspect the item and then post price $1 - \sqrt{2f + 1/4}$.
    - Total revenue = 0.617

Clairvoyant Seller

- Optimal dynamic mechanism via dynamic programming [Papadimitriou et al, Ashlagi et al, Mirrokni et al].

- Optimal auction requires clairvoyance: allocation in the first period depends on distribution F_2.

- In practice, information about the second item might not be available when we are selling the first item.

- Requires buyer to have the same belief about the future as the seller.

Non Clairvoyant Seller

- Seller doesn't know the future.
- Buyer doesn't need to agree with the seller about how the future looks like.
- Mechanism now has the following form:

$$x_t(v_{1..t}, \theta_{1..t}), p_t(v_{1..t}, \theta_{1..t}),$$

where $\theta_t \in \{$ 🍎 🍊 🍇 ... $\}$

- How does it look like ?
  - t=1 : 🍊 use $x_1(v_1, 🍊), p_1(v_1, 🍊)$
  - t=2 : 🍇 use $x_2(v_1, v_2, 🍊, 🍇)$

Power of clairvoyance

World 1 : 🍊 🍇

$x_1(v_1,$ 🍊 $)$

$x_2(v_1, v_2,$ 🍊 , 🍇 $)$

World 2 : 🍊 🍎

$x_1(v_1,$ 🍊 $)$

$x_2(v_1, v_2,$ 🍊 , 🍎 $)$

World 1 : 🍊 🍇

$x_1(v_1,$ 🍊 , 🍇 $)$

$x_2(v_1, v_2,$ 🍊 , 🍇 $)$

World 2 : 🍊 🍎

$x_1(v_1,$ 🍊 , 🍎 $)$

$x_2(v_1, v_2,$ 🍊 , 🍎 $)$

Non Clairvoyant Seller

- Non-Clairvoyant Dynamic Incentive Compatibility: if the auction is dynamic incentive compatible for every sequence of items
- e.g static auction is Non-Clairvoyant DIC
- Can we get revenue comparable to the optimal clairvoyant mechanism ?

Non Clairvoyant Revenue Approx

- Define a non-clairvoyant auction.
- Pick a sequence of items: 🍎🍎🍇🍎
- Evaluate NC auction for this sequence.
- Evaluate optimal clairvoyant auction for this sequence.

α -Revenue approximation: if for every sequence of items:

$$\mathrm{NCRev}(\mathrm{items}) \geq \alpha \cdot \mathrm{CRev}(\mathrm{items})$$

Non Clairvoyant Revenue Approx

Theorem: Every non-clairvoyant policy is at most a 1/2-approximation to the optimal clairvoyant revenue.

Theorem: For multiple buyers there is a non-clairvoyant policy that is at least 1/5-approx to the opt clairvoyant.

Non Clairvoyant Revenue Approx

Theorem: Every non-clairvoyant policy is at most a 1/2-approximation to the optimal clairvoyant revenue.

Theorem: For multiple buyers there is a non-clairvoyant policy that is at least 1/5-approx to the opt clairvoyant.

$$\geq \frac{1}{5}.$$

Non Clairvoyant Revenue Approx

Theorem: Every non-clairvoyant policy is at most a 1/2-approximation to the optimal clairvoyant revenue.

Theorem: For multiple buyers there is a non-clairvoyant policy that is at least 1/5-approx to the opt clairvoyant.

Theorem: Can be improved to 1/2 for two periods and for 1/3 for one buyer and multiple periods.

Non Clairvoyant Revenue Approx

Theorem: Every non-clairvoyant policy is at most a 1/2-approximation to the optimal clairvoyant revenue.

Theorem: For multiple buyers there is a non-clairvoyant policy that is at least 1/5-approx to the opt clairvoyant.

Theorem: Can be improved to 1/2 for two periods and for 1/3 for one buyer and multiple periods.

Technique: Bank Account Mechanisms

Theorem: Every non-clairvoyant policy is "isomorphic" to a bank account mechanism.

- Keeps a state variable b_t (balance) for each buyer
- Chooses a per-period IC mechanism based on balance

$$x_t(v_t, b_t), p_t(v_t, b_t)$$

with the balance-independence property

$$\mathbb{E}[v_t x_t(v_t, b_t) - p_t(v_t, b_t)] = \text{const} \geq 0$$

- Updates balance:

$$0 \leq b_{t+1} \leq b_t + [v_t x_t - p_t]$$

Technique: Bank Account Mechanisms

Theorem: Every non-clairvoyant policy is "isomorphic" to a bank account mechanism.

- Keeps a state variable b_t (balance) for each buyer
- Chooses a per-period IC mechanism based on balance

$$x_t(v_t, b_t), p_t(v_t, b_t)$$

with the balance-independence property

$$\mathbb{E}[v_t x_t(v_t, b_t) - p_t(v_t, b_t)] = \text{const} \geq 0$$

- Updates balance:

$$0 \leq b_{t+1} \leq b_t + [v_t x_t - p_t]$$

Technique: Bank Account Mechanisms

Theorem: Every non-clairvoyant policy is "isomorphic" to a bank account mechanism.

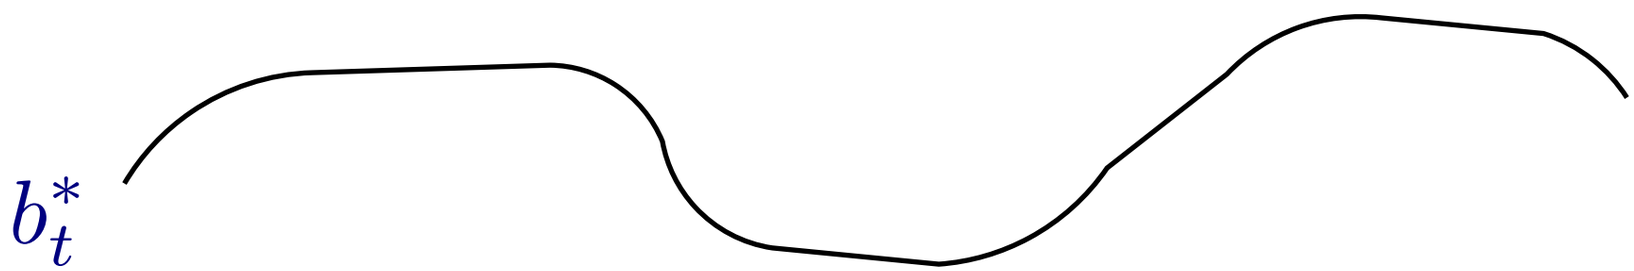

Technique: Bank Account Mechanisms

Theorem: Every non-clairvoyant policy is "isomorphic" to a bank account mechanism.

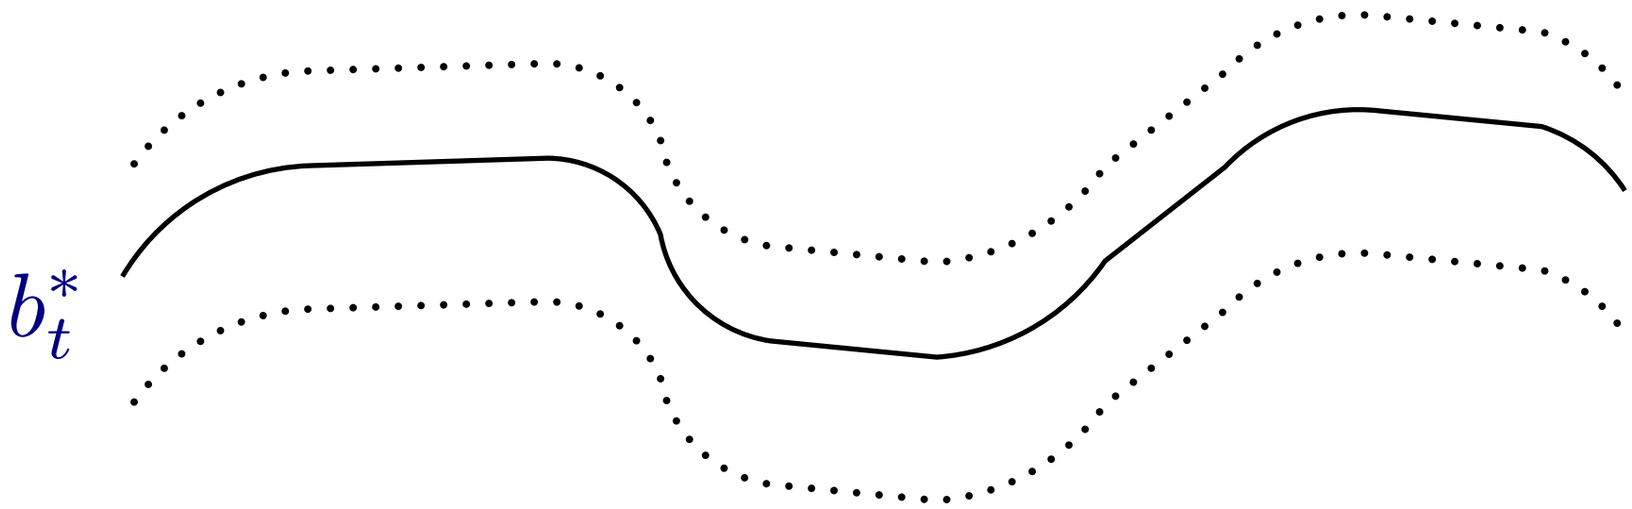

Technique: Bank Account Mechanisms

Theorem: Every non-clairvoyant policy is "isomorphic" to a bank account mechanism.

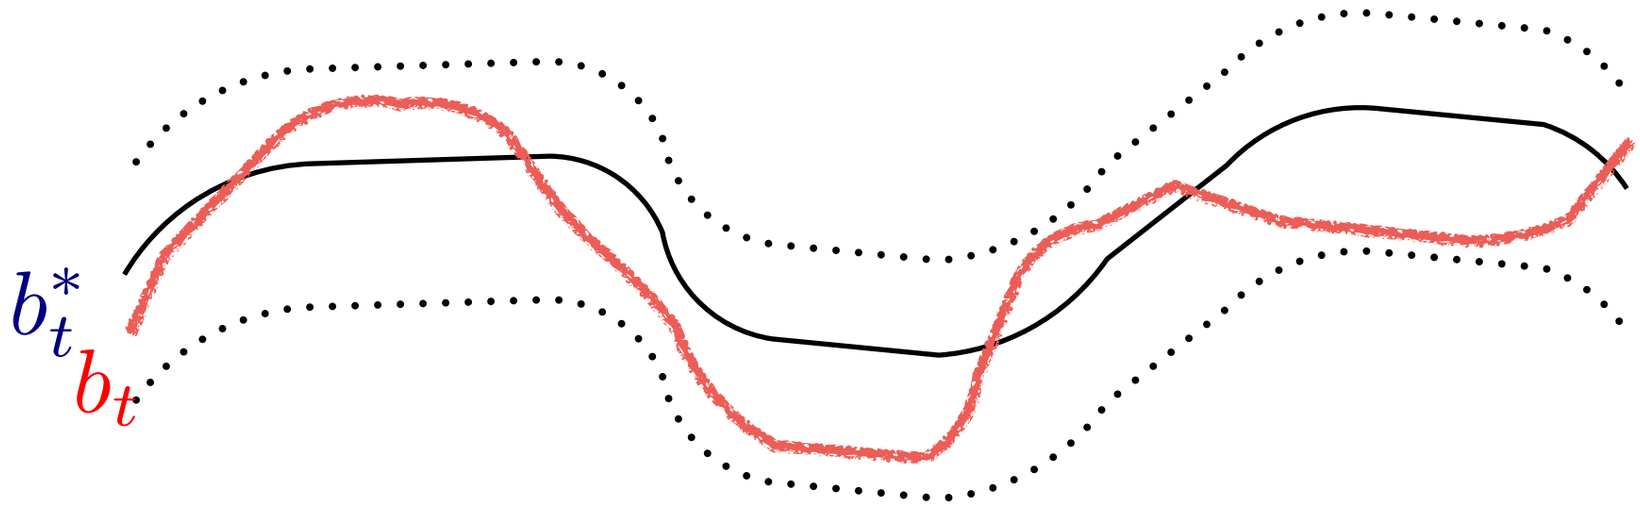

Technique: Bank Account Mechanisms

Theorem: Every non-clairvoyant policy is "isomorphic" to a bank account mechanism.

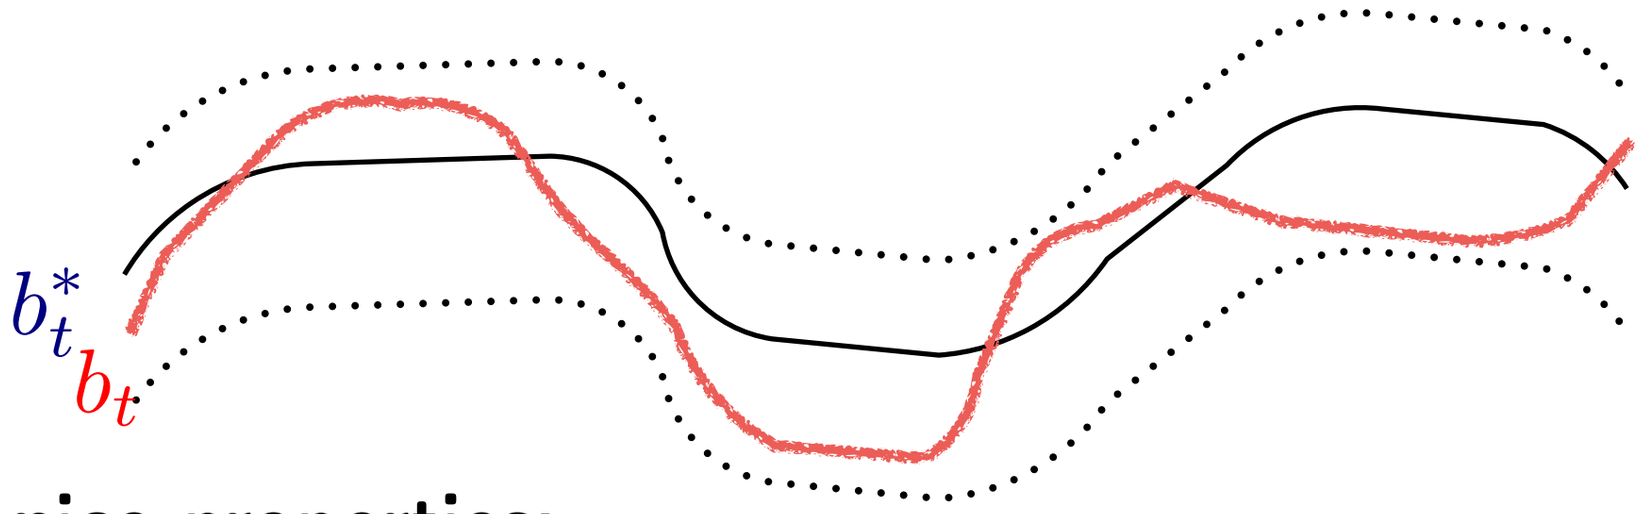

Other nice properties:

- framework to design and prove lower bounds on dynamic mechanisms
- computationally efficient (multi-buyer, multi-item)
- no pre-processing required (LP or DP)

1/3-approximation policy

Keep a variable b called balance initialized to zero.
For every period t, receive an item with distribution F_t
Sell 1/3 of the item with each of the following auctions:

- Myerson's auction for F_t
- Give the item for free and increment balance $b = b + v_t$
- For $f = \min(b, \mathbb{E}_{F_t}[v_t])$
  charge f before the buyer can see the item
  post a price of r such that $\mathbb{E}(v_t - r)^+ = f$
  decrement balance $b = b - f$

1/3-approximation policy

Keep a variable b called balance initialized to zero.
For every period t, receive an item with distribution F_t
Sell 1/3 of the item with each of the following auctions:

- Myerson's auction for F_t
- Give the item for free and increment balance $b = b + v_t$
- For $f = \min(b, \mathbb{E}_{F_t}[v_t])$
  charge f before the buyer can see the item
  post a price of r such that $\mathbb{E}(v_t - r)^+ = f$
  decrement balance $b = b - f$

Balance independence property: E[utility] is balance independent.

Motivation

Dynamic Mechanisms offer a great promise for ad auctions.

- improved revenue, efficiency and match-rate.
- once an ad impression comes, we can estimate distribution from cookies and other metadata
- we can't run expensive DPs
- we can't rely IC on buyers trusting our forecasts.

Larger program

Make dynamic auctions more friendly to industrial auction environments. Some other work:

- Martingale Auctions
  (Balseiro, Mirrokni, PL)
- Dynamic Second Price Auctions with Low Regret
  (Mirrokni, PL, Ren, Zuo)
- Dynamic Revenue Sharing
  (Balseiro, Lin, Mirrokni, PL, Zuo)
- Dynamic Mechanism Design under Positive Commitment
  (Lobel, PL)

Thanks

Non Clairvoyant Mechanism Design
https://ssrn.com/abstract=2873701