Dynamic Allocation of Reusable Resources to Strategic Agents under Long-Term Constraints

(NeurIPS'25; **Winner** of ACM Student Research @ SIGMETRICS'25)

Yan Dai Negin Golrezaei Patrick Jaillet

Massachusetts Institute of Technology

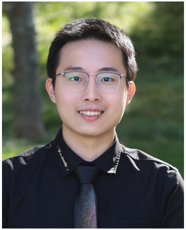

Resource Allocation under Incentives & Constraints

GPU Allocation

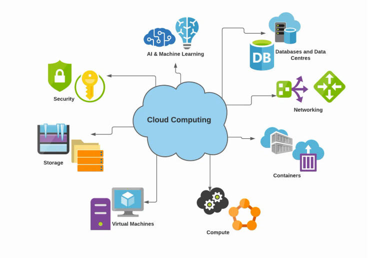

- **Resource:** Reusable GPU
- **Agents:** Research groups
- **Constr:** Energy & budget

Resource Allocation under Incentives & Constraints

GPU Allocation

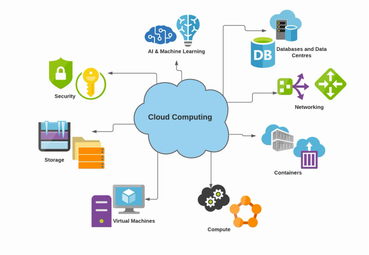

- **Resource:** Reusable GPU
- **Agents:** Research groups
- **Constr:** Energy & budget

Mobile Health Unit

- **Resource:** MHU
- **Agents:** Remote regions
- **Constr:** Staffing & budget

Efficiency-Feasibility-Incentives Trilemma

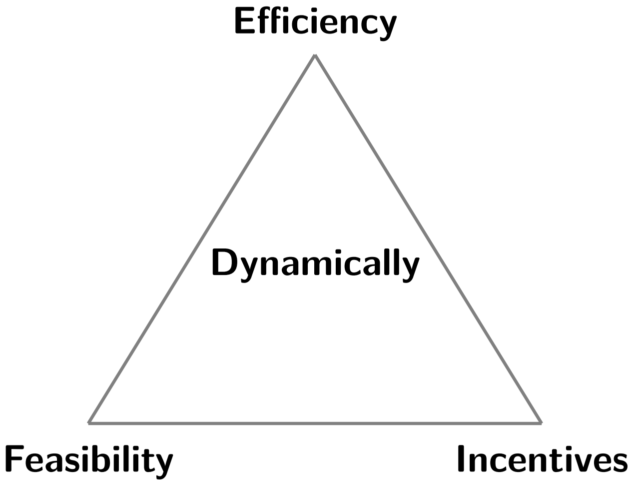

Efficiency-Feasibility-Incentives Trilemma

Efficiency

- T rounds, K agents, value $v_{t,i} \sim$ **unknown** $\mathcal{V}_i$

$$\boxed{\text{Max value: } \sum_t v_{t,i_t}}$$

Efficiency-Feasibility-Incentives Trilemma

Efficiency

- T rounds, K agents, value $v_{t,i} \sim$ **unknown** $\mathcal{V}_i$

$$\boxed{\text{Max value: } \sum_t v_{t,i_t}}$$

- Alloc cost $\boldsymbol{c}_{t,i}$ (d-dim), **iid** $\sim$ **unknown** $\mathcal{C}_i$

$$\boxed{\text{Constr: } \sum_t \boldsymbol{c}_{t,i_t} \le T\boldsymbol{\rho}}$$

Feasibility

Efficiency-Feasibility-Incentives Trilemma

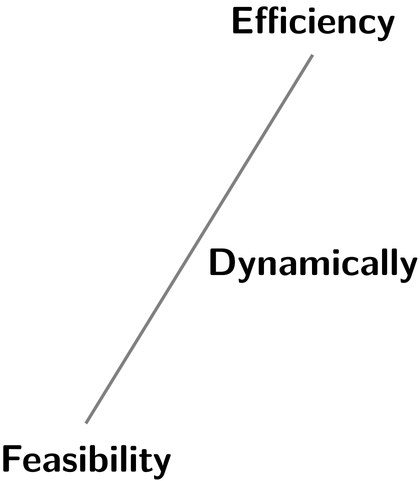

- T rounds, K agents, value $v_{t,i} \sim$ **unknown** $\mathcal{V}_i$

$$\boxed{\text{Max value: } \sum_t v_{t,i_t}}$$

- Alloc cost $\boldsymbol{c}_{t,i}$ (d-dim), **iid** $\sim$ **unknown** $\mathcal{C}_i$

$$\boxed{\text{Constr: } \sum_t \boldsymbol{c}_{t,i_t} \leq T\boldsymbol{\rho}}$$

Efficiency-Feasibility-Incentives Trilemma

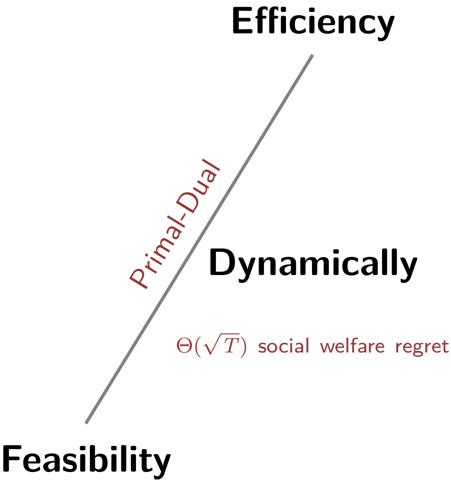

- T rounds, K agents, value $v_{t,i} \sim$ **unknown** $\mathcal{V}_i$

$$\boxed{\text{Max value: } \sum_t v_{t,i_t}}$$

- Alloc cost $\boldsymbol{c}_{t,i}$ (d-dim), **iid** $\sim$ **unknown** $\mathcal{C}_i$

$$\boxed{\text{Constr: } \sum_t \boldsymbol{c}_{t,i_t} \le T\boldsymbol{\rho}}$$

Efficiency-Feasibility-Incentives Trilemma

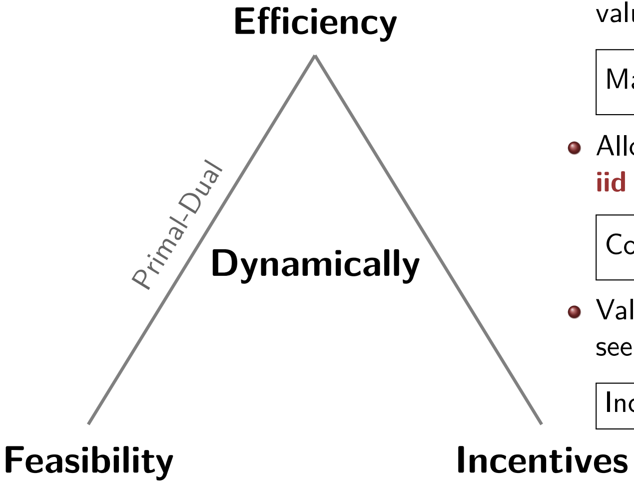

- T rounds, K agents, value $v_{t,i} \sim$ **unknown** $\mathcal{V}_i$

$$\boxed{\text{Max value: } \sum_t v_{t,i_t}}$$

- Alloc cost $\boldsymbol{c}_{t,i}$ (d-dim), **iid** $\sim$ **unknown** $\mathcal{C}_i$

$$\boxed{\text{Constr: } \sum_t \boldsymbol{c}_{t,i_t} \leq T\boldsymbol{\rho}}$$

- Value $v_{t,i}$ private; only see **strategic report** $u_{t,i}$

$$\boxed{\text{Incentivize } u_{t,i} \approx v_{t,i}}$$

Efficiency-Feasibility-Incentives Trilemma

Efficiency

Dynamically

Incentives

- T rounds, K agents, value $v_{t,i} \sim$ **unknown** $\mathcal{V}_i$

$$\boxed{\text{Max value: } \sum_t v_{t,i_t}}$$

- Value $v_{t,i}$ private; only see **strategic report** $u_{t,i}$

$$\boxed{\text{Incentivize } u_{t,i} \approx v_{t,i}}$$

Efficiency-Feasibility-Incentives Trilemma

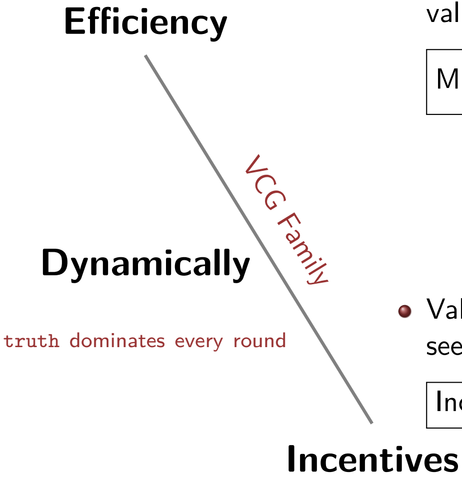

- T rounds, K agents, value $v_{t,i} \sim$ **unknown** $\mathcal{V}_i$

$$\boxed{\text{Max value: } \sum_t v_{t,i_t}}$$

- Value $v_{t,i}$ private; only see **strategic report** $u_{t,i}$

$$\boxed{\text{Incentivize } u_{t,i} \approx v_{t,i}}$$

Efficiency-Feasibility-Incentives Trilemma

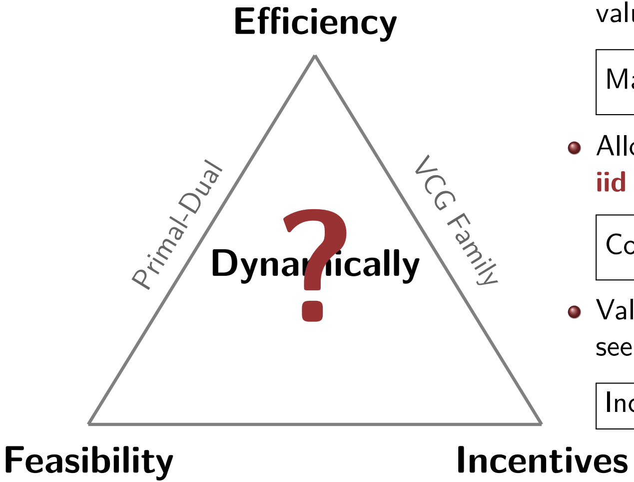

- T rounds, K agents, value $v_{t,i} \sim$ **unknown** $\mathcal{V}_i$

$$\boxed{\text{Max value: } \sum_t v_{t,i_t}}$$

- Alloc cost $\boldsymbol{c}_{t,i}$ (d-dim), **iid** $\sim$ **unknown** $\mathcal{C}_i$

$$\boxed{\text{Constr: } \sum_t \boldsymbol{c}_{t,i_t} \le T\boldsymbol{\rho}}$$

- Value $v_{t,i}$ private; only see **strategic report** $u_{t,i}$

$$\boxed{\text{Incentivize } u_{t,i} \approx v_{t,i}}$$

Efficiency-Feasibility-Incentives Trilemma

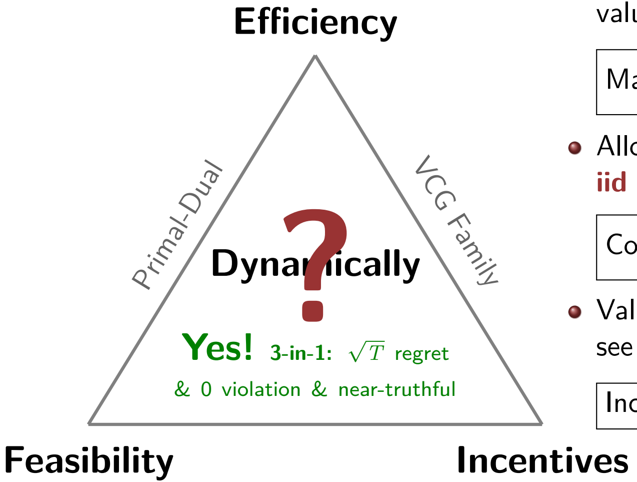

- T rounds, K agents, value $v_{t,i} \sim$ **unknown** $\mathcal{V}_i$

$$\boxed{\text{Max value: } \sum_t v_{t,i_t}}$$

- Alloc cost $\boldsymbol{c}_{t,i}$ (d-dim), **iid** $\sim$ **unknown** $\mathcal{C}_i$

$$\boxed{\text{Constr: } \sum_t \boldsymbol{c}_{t,i_t} \le T\boldsymbol{\rho}}$$

- Value $v_{t,i}$ private; only see **strategic report** $u_{t,i}$

$$\boxed{\text{Incentivize } u_{t,i} \approx v_{t,i}}$$

Classical Primal-Dual Fails with Strategic Agents

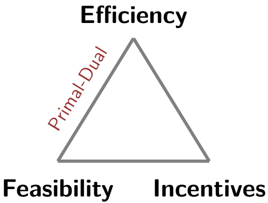

Figure: Classical Primal-Dual in Round t

Classical Primal-Dual Fails with Strategic Agents

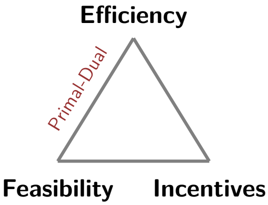

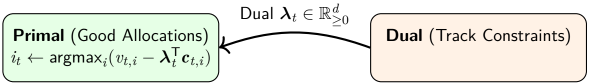

Figure: Classical Primal-Dual in Round t

Classical Primal-Dual Fails with Strategic Agents

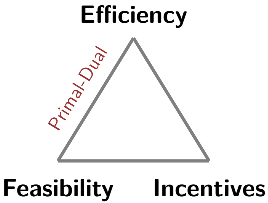

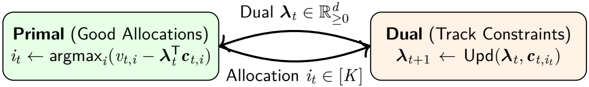

Figure: Classical Primal-Dual in Round t

Classical Primal-Dual Fails with Strategic Agents

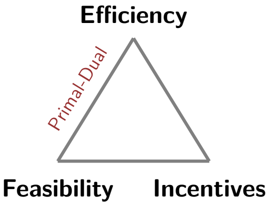

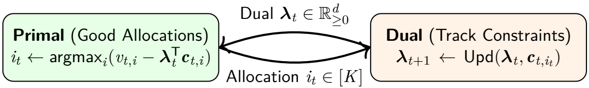

Figure: Classical Primal-Dual in Round t

What Happens With Strategic Agents?

Manipulate $v_{t,i}$

$\downarrow$

Alter i_t

Classical Primal-Dual Fails with Strategic Agents

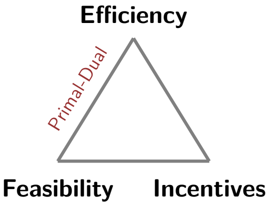

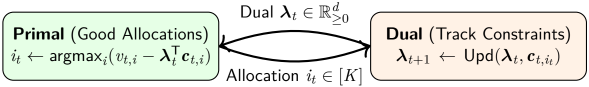

Figure: Classical Primal-Dual in Round t

What Happens With Strategic Agents?

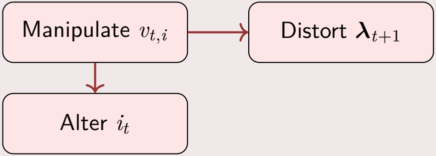

Classical Primal-Dual Fails with Strategic Agents

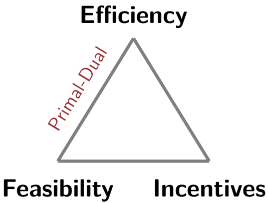

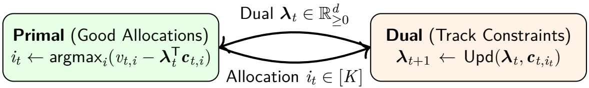

Figure: Classical Primal-Dual in Round t

What Happens With Strategic Agents?

Manipulate $v_{t,i}$ → Distort $\boldsymbol{\lambda}_{t+1}$ → Affect i_{t+1}

Manipulate $v_{t,i}$ → Alter i_t

Classical Primal-Dual Fails with Strategic Agents

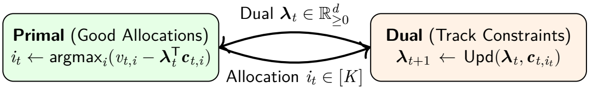

Figure: Classical Primal-Dual in Round t

What Happens With Strategic Agents?

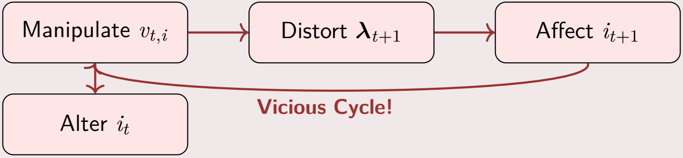

Classical Primal-Dual Fails with Strategic Agents

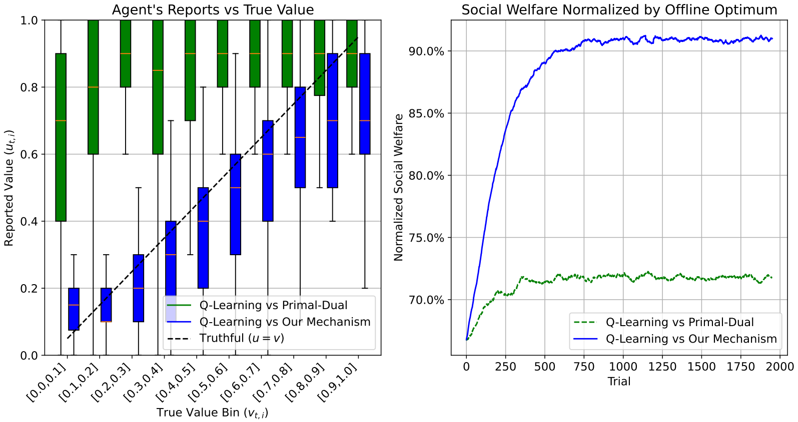

Figure: Vanilla Primal-Dual vs Our Mechanism

Classical Primal-Dual Fails with Strategic Agents

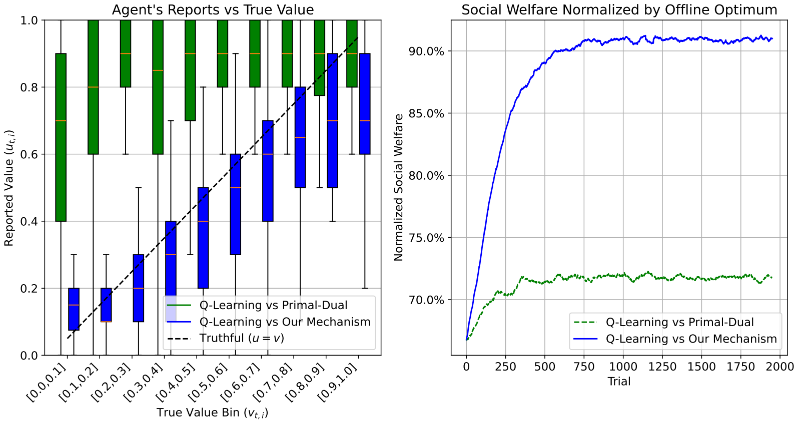

Figure: Vanilla Primal-Dual vs Our Mechanism

misreport vs truthful

Classical Primal-Dual Fails with Strategic Agents

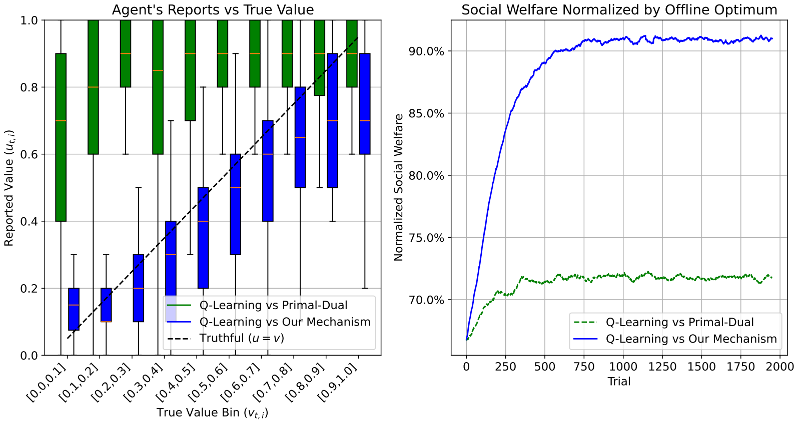

Figure: Vanilla Primal-Dual vs Our Mechanism

misreport vs truthful

low vs high efficiency

Primal Allocations

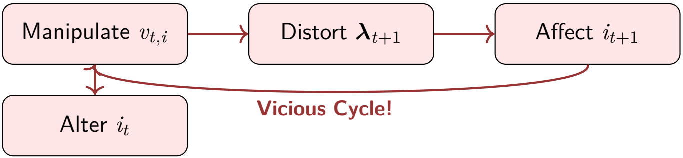

Primal Allocations: Pricing

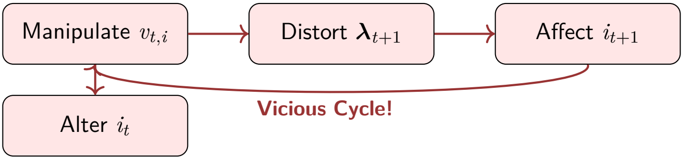

1. **Dual-Adjusted Pricing.** VCG-like rule (adapted for $\boldsymbol{\lambda}$)

Primal Allocations: Pricing

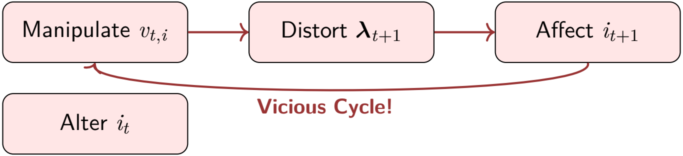

1. **Dual-Adjusted Pricing.** VCG-like rule (adapted for $\boldsymbol{\lambda}$)
   $\implies$ `truth` dominates (for static setups)

Primal Allocations: Pricing + Epoching

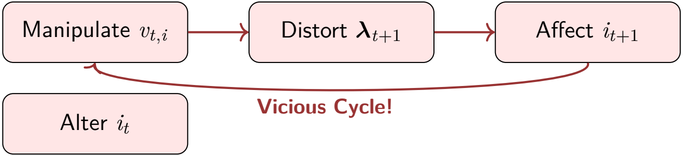

1. **Dual-Adjusted Pricing.** VCG-like rule (adapted for $\boldsymbol{\lambda}$)
   $\implies$ `truth` dominates (for static setups)
2. **Epoch-Based Lazy Updates.** Fix $\boldsymbol{\lambda}$ for $\sqrt{T}$ rounds

Primal Allocations: Pricing + Epoching

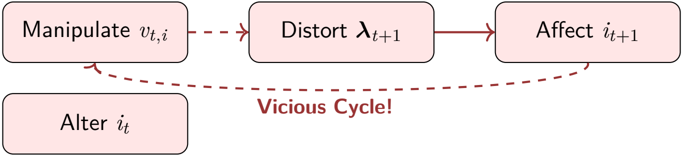

1. **Dual-Adjusted Pricing.** VCG-like rule (adapted for $\boldsymbol{\lambda}$)
   $\Longrightarrow$ `truth` dominates (for static setups)
2. **Epoch-Based Lazy Updates.** Fix $\boldsymbol{\lambda}$ for $\sqrt{T}$ rounds
   $\Longrightarrow$ hard to affect future (but not impossible)

Primal Allocations: Pricing + Epoching + Exploration

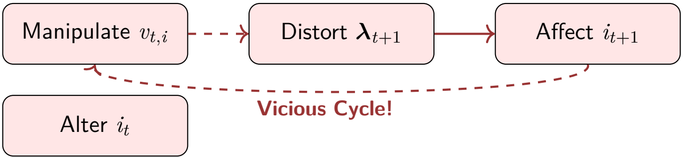

1. **Dual-Adjusted Pricing.** VCG-like rule (adapted for $\boldsymbol{\lambda}$)
   $\Longrightarrow$ `truth` dominates (for static setups)
2. **Epoch-Based Lazy Updates.** Fix $\boldsymbol{\lambda}$ for $\sqrt{T}$ rounds
   $\Longrightarrow$ hard to affect future (but not impossible)
3. **Randomized Exploration.** Misreport means harm

Primal Allocations: Pricing + Epoching + Exploration

Manipulate $v_{t,i}$

Distort $\boldsymbol{\lambda}_{t+1}$ $\rightarrow$ Affect i_{t+1}

Alter i_t

1. **Dual-Adjusted Pricing.** VCG-like rule (adapted for $\boldsymbol{\lambda}$)
   $\Longrightarrow$ `truth` dominates (for static setups)
2. **Epoch-Based Lazy Updates.** Fix $\boldsymbol{\lambda}$ for $\sqrt{T}$ rounds
   $\Longrightarrow$ hard to affect future (but not impossible)
3. **Randomized Exploration.** Misreport means harm
   $\Longrightarrow$ $\exists$ near-truthful equilibrium (if harm $\geq$ gain)

Primal Allocations: Pricing + Epoching + Exploration

Manipulate $v_{t,i}$

Distort $\boldsymbol{\lambda}_{t+1}$ $\longrightarrow$ Affect i_{t+1}

Alter i_t

1. **Dual-Adjusted Pricing.** VCG-like rule (adapted for $\boldsymbol{\lambda}$)
   $\Longrightarrow$ `truth` dominates (for static setups)
2. **Epoch-Based Lazy Updates.** Fix $\boldsymbol{\lambda}$ for $\sqrt{T}$ rounds
   $\Longrightarrow$ hard to affect future (but not impossible)
3. **Randomized Exploration.** Misreport means harm
   $\Longrightarrow$ $\exists$ near-truthful equilibrium (if harm $\geq$ gain)

Theorem. $\boxed{\widetilde{\mathcal{O}}(1) \text{ misreports \& } \widetilde{\mathcal{O}}(1) \text{ misallocations}}$ per epoch

Dual Updates

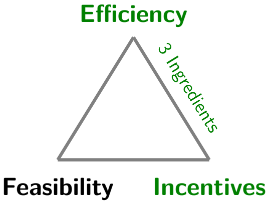

Dual Updates

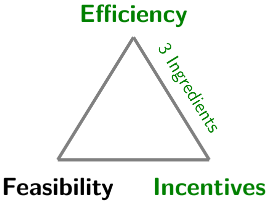

Dynamically tune $\boldsymbol{\lambda}_1, \boldsymbol{\lambda}_2, \ldots$ according to costs

Dual Updates: Online Learning gives $\widetilde{\mathcal{O}}(T^{2/3})$

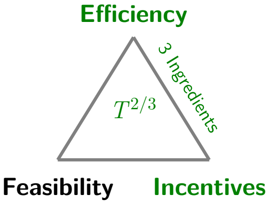

Dynamically tune $\boldsymbol{\lambda}_1, \boldsymbol{\lambda}_2, \ldots$ according to costs

Theorem 1: Sublinear Regret ✓

3 ingredients (primal) + GD / FTRL (dual)
$\Longrightarrow \widetilde{\mathcal{O}}(T^{2/3})$ **regret** ("no-regret" guarantee)

Dual Updates: Online Learning gives $\widetilde{\mathcal{O}}(T^{2/3})$

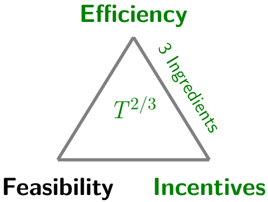

Dynamically tune $\boldsymbol{\lambda}_1, \boldsymbol{\lambda}_2, \ldots$ according to costs

Theorem 1: Sublinear Regret ✓

3 ingredients (primal) + GD / FTRL (dual)
$\Longrightarrow \widetilde{\mathcal{O}}(T^{2/3})$ **regret** ("no-regret" guarantee)

Can we do better?

Dual Updates: Online Learning gives $\widetilde{\mathcal{O}}(T^{2/3})$

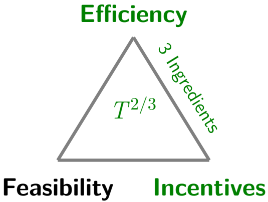

Dynamically tune $\boldsymbol{\lambda}_1, \boldsymbol{\lambda}_2, \ldots$ according to costs

Theorem 1: Sublinear Regret ✓

3 ingredients (primal) + GD / FTRL (dual)
$\Longrightarrow \widetilde{\mathcal{O}}(T^{2/3})$ **regret** ("no-regret" guarantee)

Can we do better?

- Lazy updates **good for primal** (less incentives & ablities)

Dual Updates: Online Learning gives $\widetilde{\mathcal{O}}(T^{2/3})$

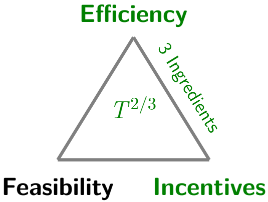

Dynamically tune $\boldsymbol{\lambda}_1, \boldsymbol{\lambda}_2, \ldots$ according to costs

Theorem 1: Sublinear Regret ✓

3 ingredients (primal) + GD / FTRL (dual) $\Longrightarrow \widetilde{\mathcal{O}}(T^{2/3})$ **regret** ("no-regret" guarantee)

Can we do better?

- Lazy updates **good for primal** (less incentives & ablities)
- Lazy updates **bad for dual**

Dual Updates: Online Learning gives $\widetilde{\mathcal{O}}(T^{2/3})$

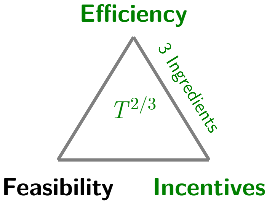

Dynamically tune $\boldsymbol{\lambda}_1, \boldsymbol{\lambda}_2, \ldots$ according to costs

Theorem 1: Sublinear Regret ✓

3 ingredients (primal) + GD / FTRL (dual)
$\Longrightarrow \widetilde{\mathcal{O}}(T^{2/3})$ **regret** ("no-regret" guarantee)

Can we do better?

- Lazy updates **good for primal** (less incentives & ablities)
- Lazy updates **bad for dual**

No Lazy $\quad \boldsymbol{c}_{t-1,i_{t-1}} \longrightarrow \boldsymbol{\lambda}_t \longrightarrow \boldsymbol{c}_{t,i_t} \longrightarrow \boldsymbol{\lambda}_{t+1}$

Dual Updates: Online Learning gives $\widetilde{\mathcal{O}}(T^{2/3})$

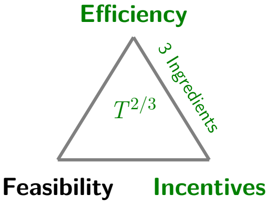

Dynamically tune $\boldsymbol{\lambda}_1, \boldsymbol{\lambda}_2, \ldots$ according to costs

Theorem 1: Sublinear Regret ✓

3 ingredients (primal) + GD / FTRL (dual)
$\Longrightarrow \widetilde{\mathcal{O}}(T^{2/3})$ **regret** ("no-regret" guarantee)

Can we do better?

- Lazy updates **good for primal** (less incentives & ablities)
- Lazy updates **bad for dual**

Start of Epoch ℓ

No Lazy $\boldsymbol{c}_{t-1,i_{t-1}} \to \boldsymbol{\lambda}_t \to \boldsymbol{c}_{t,i_t} \to \boldsymbol{\lambda}_{t+1}$

Lazy

Dual Updates: Online Learning gives $\widetilde{\mathcal{O}}(T^{2/3})$

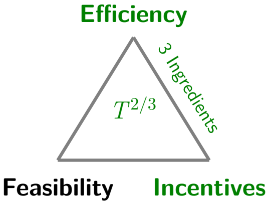

Dynamically tune $\boldsymbol{\lambda}_1, \boldsymbol{\lambda}_2, \ldots$ according to costs

Theorem 1: Sublinear Regret ✓

3 ingredients (primal) + GD / FTRL (dual)
$\Longrightarrow \widetilde{\mathcal{O}}(T^{2/3})$ **regret** ("no-regret" guarantee)

Can we do better?

- Lazy updates **good for primal** (less incentives & ablities)
- Lazy updates **bad for dual**

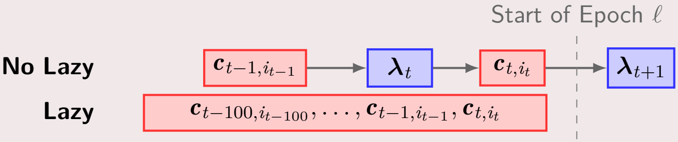

Dual Updates: Online Learning gives $\widetilde{\mathcal{O}}(T^{2/3})$

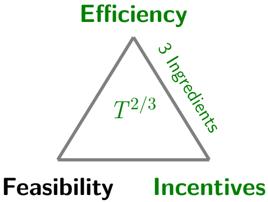

Dynamically tune $\boldsymbol{\lambda}_1, \boldsymbol{\lambda}_2, \ldots$ according to costs

Theorem 1: Sublinear Regret ✓

3 ingredients (primal) + GD / FTRL (dual)
$\Longrightarrow \widetilde{\mathcal{O}}(T^{2/3})$ **regret** ("no-regret" guarantee)

Can we do better?

- Lazy updates **good for primal** (less incentives & ablities)
- Lazy updates **bad for dual**

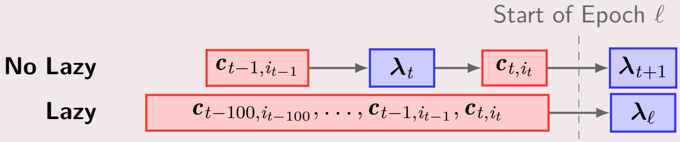

Dual Updates: Online Learning gives $\widetilde{\mathcal{O}}(T^{2/3})$

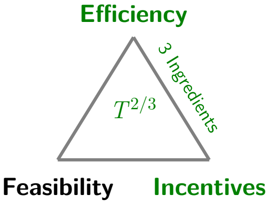

Dynamically tune $\boldsymbol{\lambda}_1, \boldsymbol{\lambda}_2, \ldots$ according to costs

Theorem 1: Sublinear Regret ✓

3 ingredients (primal) + GD / FTRL (dual)
$\Longrightarrow \widetilde{\mathcal{O}}(T^{2/3})$ **regret** ("no-regret" guarantee)

Can we do better?

- Lazy updates **good for primal** (less incentives & ablities)
- Lazy updates **bad for dual** (cannot react to $\boldsymbol{c}_{t,i_t}$ promptly)

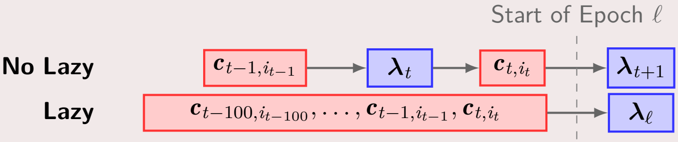

Dual Updates: Online Learning gives $\widetilde{\mathcal{O}}(T^{2/3})$

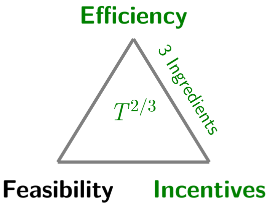

Dynamically tune $\boldsymbol{\lambda}_1, \boldsymbol{\lambda}_2, \ldots$ according to costs

Theorem 1: Sublinear Regret ✓

3 ingredients (primal) + GD / FTRL (dual)
$\Longrightarrow \widetilde{\mathcal{O}}(T^{2/3})$ **regret** ("no-regret" guarantee)

Can we do better?

- Lazy updates **good for primal** (less incentives & ablities)
- Lazy updates **bad for dual** (cannot react to $\boldsymbol{c}_{t,i_t}$ promptly)

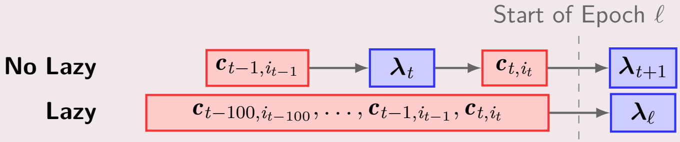

Theorem. "Low-switching online learning" has $\Omega(T^{2/3})$ **regret**

Dual Updates: Even Better via Predictability

Start of Epoch ℓ

Lazy FTRL Past costs $\rightarrow$ Action $\boldsymbol{\lambda}_\ell$ $\Omega(T^{2/3})$ **regret**

Dual Updates: Even Better via Predictability

Start of Epoch ℓ

Lazy FTRL Past costs → Action $\boldsymbol{\lambda}_\ell$ $\Omega(T^{2/3})$ **regret**

Key Insight: (Almost-)Truthfulness $\Longrightarrow$ Predictability

Dual Updates: Even Better via Predictability

Start of Epoch ℓ

Lazy FTRL Past costs → Action $\boldsymbol{\lambda}_\ell$ $\Omega(T^{2/3})$ **regret**

Key Insight: (Almost-)Truthfulness $\Longrightarrow$ Predictability

1. Truthful $\Longrightarrow$ **iid future cost** $\boldsymbol{c}_{t,i_t}$ $(i_t \approx \text{argmax}_i(v_{t,i} - \boldsymbol{\lambda}_\ell^\top \boldsymbol{c}_{t,i}))$

Dual Updates: Even Better via Predictability

Start of Epoch ℓ

Lazy FTRL | Past costs → Action $\boldsymbol{\lambda}_\ell$ | $\Omega(T^{2/3})$ **regret**

Key Insight: (Almost-)Truthfulness $\Longrightarrow$ Predictability

1. Truthful $\Longrightarrow$ **iid future cost $\boldsymbol{c}_{t,i_t}$** ($i_t \approx \operatorname{argmax}_i(v_{t,i} - \boldsymbol{\lambda}_\ell^\top \boldsymbol{c}_{t,i})$)
2. Truthful $\Longrightarrow$ **reliable history** (for distributions $\mathcal{V}_i$ and $\mathcal{C}_i$)

Dual Updates: Even Better via Predictability

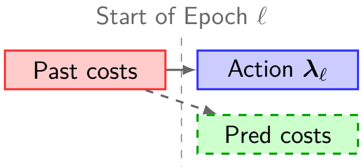

Key Insight: (Almost-)Truthfulness $\Longrightarrow$ Predictability

1. Truthful $\Longrightarrow$ **iid future cost $\boldsymbol{c}_{t,i_t}$** ($i_t \approx \operatorname{argmax}_i(v_{t,i} - \boldsymbol{\lambda}_\ell^\top \boldsymbol{c}_{t,i})$)
2. Truthful $\Longrightarrow$ **reliable history** (for distributions $\mathcal{V}_i$ and $\mathcal{C}_i$)

$\Longrightarrow$ **Predict new costs**

Dual Updates: Even Better via Predictability

Optimistic FTRL

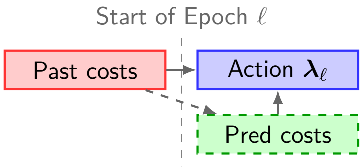

Key Insight: (Almost-)Truthfulness $\Longrightarrow$ Predictability

1. Truthful $\Longrightarrow$ **iid future cost $\boldsymbol{c}_{t,i_t}$** $(i_t \approx \text{argmax}_i(v_{t,i} - \boldsymbol{\lambda}_\ell^\top \boldsymbol{c}_{t,i}))$
2. Truthful $\Longrightarrow$ **reliable history** (for distributions $\mathcal{V}_i$ and $\mathcal{C}_i$)

$\Longrightarrow$ **Predict new costs** for better action $\boldsymbol{\lambda}_\ell$

Dual Updates: Even Better via Predictability

Optimistic FTRL

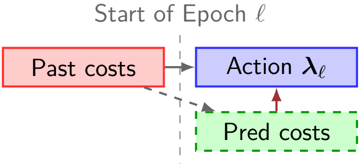

Issue: Circular Dependency

- Yield $\boldsymbol{\lambda}_\ell$ **as-if true costs = pred costs**

Dual Updates: Even Better via Predictability

Optimistic FTRL

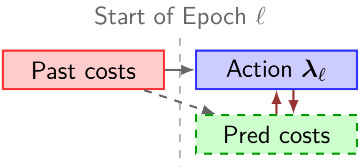

Circular dep

Issue: Circular Dependency

- Yield $\boldsymbol{\lambda}_\ell$ **as-if true costs = pred costs**
- Yield pred costs **as-if true dual = $\boldsymbol{\lambda}_\ell$**

Dual Updates: Even Better via Predictability

Our Framework

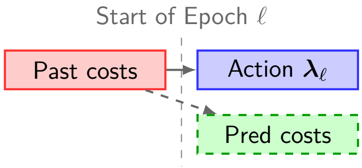

Issue: Circular Dependency

- Yield $\boldsymbol{\lambda}_\ell$ **as-if true costs = pred costs**
- Yield pred costs **as-if true dual = $\boldsymbol{\lambda}_\ell$**

Novel online learning

Dual Updates: Even Better via Predictability

Our Framework

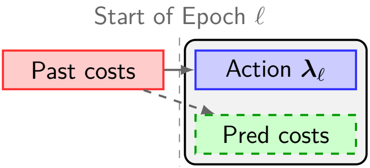

Issue: Circular Dependency

- Yield $\boldsymbol{\lambda}_\ell$ **as-if true costs = pred costs**
- Yield pred costs **as-if true dual = $\boldsymbol{\lambda}_\ell$**

Novel online learning

Dual Updates: Even Better via Predictability

Our Framework

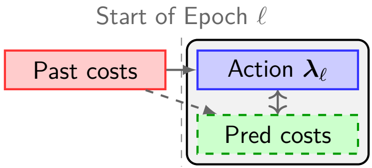

Issue: Circular Dependency

- Yield $\boldsymbol{\lambda}_\ell$ **as-if true costs = pred costs**
- Yield pred costs **as-if true dual =** $\boldsymbol{\lambda}_\ell$

Novel online learning (decide action $\boldsymbol{\lambda}_\ell$ & pred costs simultaneously via fixed-point subroutine; named **O-FTRL-FP**)

Dual Updates: Even Better via Predictability

Our Framework

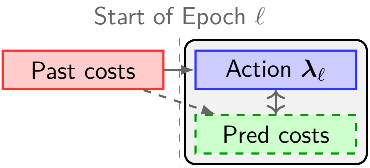

$\widetilde{\mathcal{O}}(\sqrt{T})$ **regret**

Issue: Circular Dependency

- Yield $\boldsymbol{\lambda}_\ell$ **as-if true costs = pred costs**
- Yield pred costs **as-if true dual = $\boldsymbol{\lambda}_\ell$**

Novel online learning (decide action $\boldsymbol{\lambda}_\ell$ & pred costs simultaneously via fixed-point subroutine; named **O-FTRL-FP**) $\Longrightarrow \widetilde{\mathcal{O}}(\sqrt{T})$ **regret**

Dual Updates: Even Better via Predictability

Our Framework

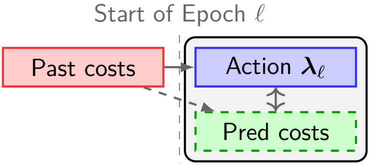

$\widetilde{\mathcal{O}}(\sqrt{T})$ **regret**

Issue: Circular Dependency

- Yield $\boldsymbol{\lambda}_\ell$ **as-if true costs = pred costs**
- Yield pred costs **as-if true dual = $\boldsymbol{\lambda}_\ell$**

Novel online learning (decide action $\boldsymbol{\lambda}_\ell$ & pred costs simultaneously via fixed-point subroutine; named **O-FTRL-FP**) $\Longrightarrow$ $\widetilde{\mathcal{O}}(\sqrt{T})$ **regret**

Recall. Non-strategic lower bound = $\Omega(\sqrt{T})$ regret

Main Results & Takeaway

Main Result

1st dynamic mechanism resolving **trilemma**:

- **Efficiency.** Optimal $\widetilde{\mathcal{O}}(\sqrt{T})$ regret
- **Feasibility.** Zero constraint violation
- **Incentives.** Robust to strategic agents

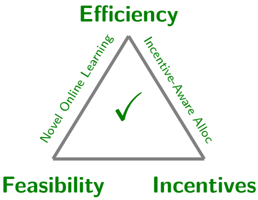

Key Techniques

- **Primal Side: Incentive-Aware Allocation.** Novel mixture of dual-adjusted pricing + lazy updates + random exploration
- **Dual Side: Online Learning for Updates.** Truthfulness ⇒ Predictability + novel framework for circular dependencies

Questions are more than welcomed!

✉ yandai20@mit.edu; 🌍 https://yandaichn.github.io/