Counterbalancing Learning and Strategic Incentives in Allocation Markets

Jamie Kang[1], Faidra Monachou[1], Moran Koren[2], Itai Ashlagi[1]

[1] Stanford University
[2] Harvard University

NeurIPS 2021

MOTIVATION: KIDNEY ALLOCATION WAITLIST

Transplant (tx) quality

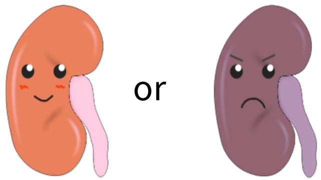

Good Bad

Ex ante unknown

Information about quality

1. **Organ score**

   P {successful tx} based on organ features (e.g., size, donor age)

   Public

2. **Doctors' private opinions**

   Based on their own experience / knowledge

   Private

MOTIVATION: KIDNEY ALLOCATION WAITLIST

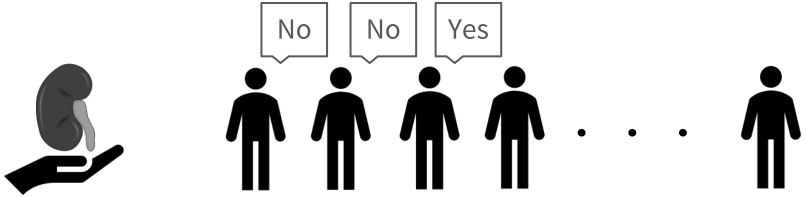

Planner **Waitlist (patients / doctors)**

- Patients waiting for an organ offer
- Upon receiving offer, each patient decides to **accept** or **decline**
- Or in most cases, his or her doctor makes decision
- Social planner decides whether / how to make the offers

Planner's Goal:
Optimize overall tx quality

i.e. Utilize good organs and discard bad organs

Patient / Doctor's Goal:
Optimize my tx quality

BASELINE: FIRST-COME-FIRST-SERVE MECHANISM

- Commonly used -- aka Sequential Offering
- Object offered to each agent sequentially one-by-one

What could go wrong?

X To k-th agent:

availability of object implies previous (k-1) agents have declined it

X Induces **herding** behavior → **incorrect discard** of objects

X In kidney allocation: > 20% discard rate, while ~3.6yr wait time

(De Mel et al. (2020), Mohan et al. (2018), Zhang (2010) for empirical evidence)

MAIN PROBLEM AND RESULTS

Q: Given a **single** indivisible object of **unknown quality,** whether and how to **allocate** it to a queue of **privately informed** and **strategic** and agents?

I.e., How to balance planner's learning and agents' strategic incentives?

MAIN PROBLEM AND RESULTS

A:

1. FCFS can cause welfare loss due to herding
2. Propose a **new class of mechanisms** involving dynamic batched voting to **crowdsource** private information, and show **existence** of such mechanisms that **improve welfare**
3. **Simple greedy** algorithm to achieve this improvement

RELATED LITERATURE

Social Learning

- Banerjee (1992), Bikhchandani et al. (1992)
- In kidney markets: De Mel et al. (2020), Mohan et al. (2018), Zhang (2010)

Voting

- Austen-Smith and Banks (1996), Condorcet (1785)

Information Design / Bayesian Exploration

- Arieli et al. (2018) , Kamenica and Gentzkow (2011), Papanastasiou et al. (2017)
- Glazer et al. (2021), Immorlica et al. (2019), Kremer et al. (2014), Mansour et al. (2016)

MODEL

SET UP

Object

- Single indivisible object
- **Quality**: $\omega \in$ {G, B}

  fixed and ex-ante unknown
- **Prior**: μ = P (w=G)

  commonly known

Agents

For each agent in position *i*

- **Private signal**: $s_i \in \{g, b\}$
- **Precision** of signal:

  q = P (s_i=g | w=G) = P (s_i =b | w=B) ∈ *(1/2,1)*

  commonly known
- **Utility**:

$$\begin{cases} 1 & \text{with object and } \omega = G \\ -1 & \text{with object and } \omega = B \\ 0 & \text{without object} \end{cases}$$

SET UP

Planner's goal

Design a **mechanism** to maximize **Pr {correct allocation outcome}**

1. Asks (a batch of) agents to report private signals
2. Decides whether and how to allocate the object

e.g., FCFS, Lottery…

Allocate if $\omega = G$

Discard if $\omega = B$

We propose a new class of mechanisms

VOTING MECHANISMS

- Idea: **batch-by-batch dynamic voting** to crowdsource information
- For each batch *j*:
  1. Offer to a **batch of K_j** agents
  2. Each agent simultaneously votes to opt in or opt out.
  3. If **majority** opts in: Allocate object uniformly at random.
     Otherwise: Move on to batch j+1.
     Results from batch j become public.

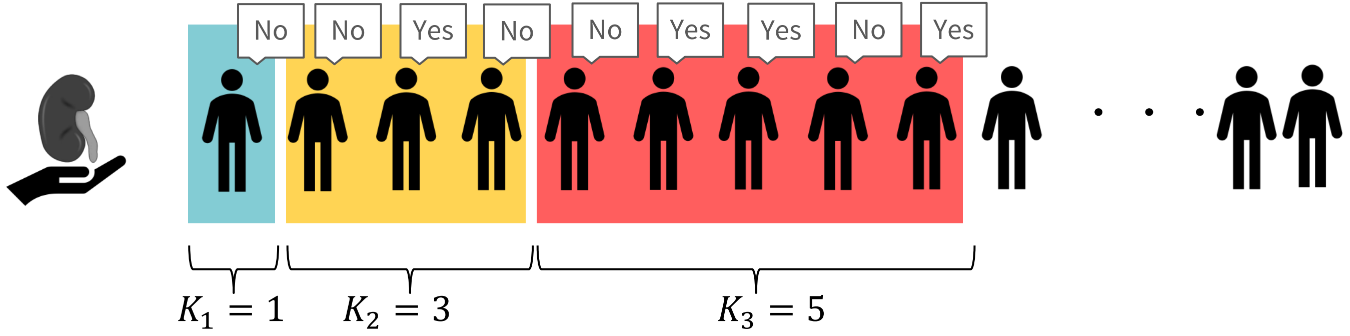

Batch size K_j can be chosen dynamically

VOTING MECHANISMS

- Idea: **batch-by-batch dynamic voting** to crowdsource information
- For each batch *j*:
  1. Offer to a **batch of K_j** agents
  2. Each agent simultaneously votes to opt in or opt out
  3. If **majority** opts in: Allocate object uniformly at random.
     Otherwise: Move on to batch j+1.
     Results from batch j become public.

Sequential → Batch

Randomness

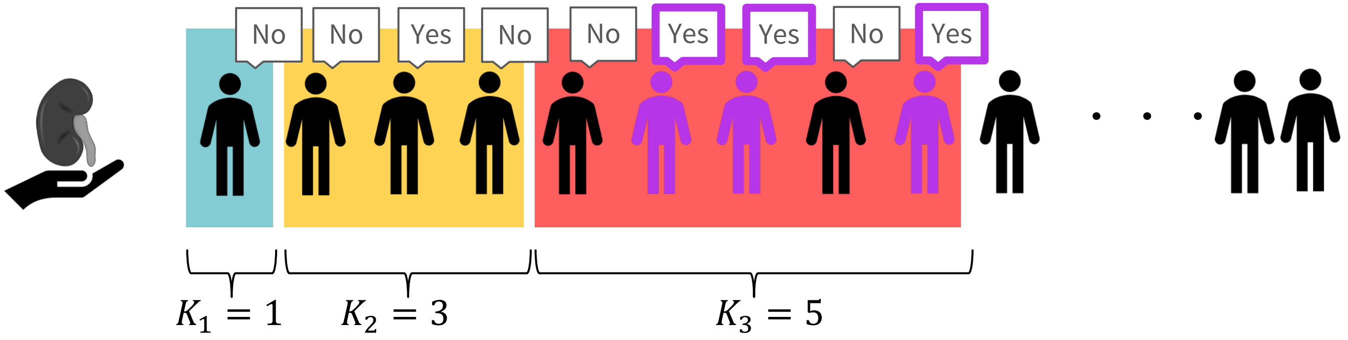

Batch size K_j can be chosen dynamically

VOTING MECHANISMS

FCFS is also a voting mechanism ($K_j = 1$ for all j)

We restrict our attention to the class of voting mechanisms

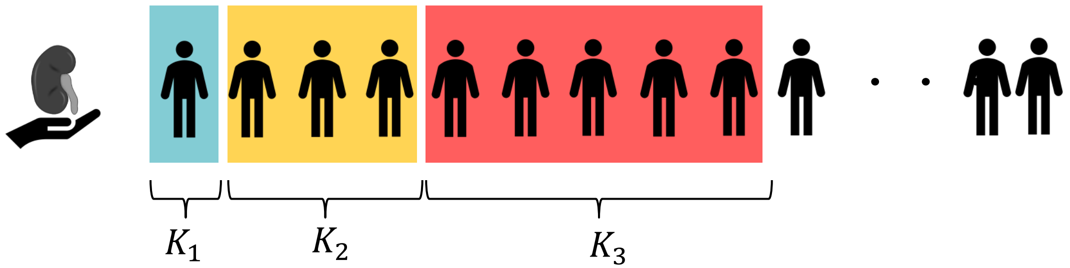

Main problem reduces to:

How to dynamically choose batch size K_j?

RESULTS

FCFS: HAMPERED LEARNING W/ WRONG BATCH SIZE (K=1)

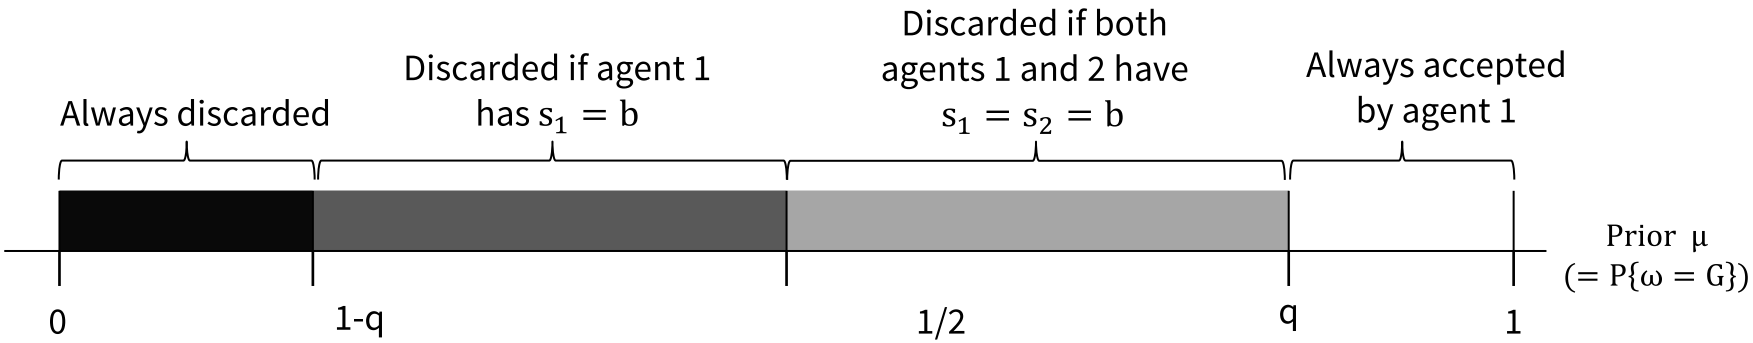

Figure 1. Allocation outcome of the sequential offer mechanism based on the value of prior $\mu \in (0, 1)$ with respect to signal precision q.

- In this extreme case, planner **can learn from only up to two** agents
- Restricted learning leads to **poor correctness** and welfare loss

BALANCING AGENTS' INCENTIVES vs PLANNER'S LEARNING

	Small batch size		Large batch size
✓	**Every vote is pivotal**: in particular, incentivizes agents with $s_i = b$ signals to truthfully opt out	✓	**More data points**: gives confidence that if object is allocated, then it is likely that quality is good
✗	If too small, allocation depends on learning from insufficient sample size	✗	If too large, everyone is incentivized to opt in

- Presence of **incentives puts upper bound on # of private signals** planner can learn from
- Optimal batch size is the **maximum batch size** that agents' **incentives allow** (i.e., IC is tight)

__Kang__, Monachou, Koren, Ashlagi (NeurIPS 2021)

MAIN THEORETICAL RESULTS

Theorem 1.

- For any $\boldsymbol{q} > \boldsymbol{\mu}$, there **always exists a voting mechanism** $V \in \mathcal{V}$ that is **incentive-compatible** and **improves correctness** compared to the sequential offering mechanism V_{SEQ}.

- For any $q \leq \mu$, there is no incentive-compatible voting mechanism and any $V \in \mathcal{V}$ achieves the same correctness as V_{SEQ}.

Corollary 1.

- For any $q > \mu$, such a mechanism can be found using a **greedy algorithm**.

TIGHTER INCENTIVES FOR THE WELL-INFORMED (HIGH q) AND OPTIMISTIC (HIGH μ)

(Formal proofs and results in the paper)

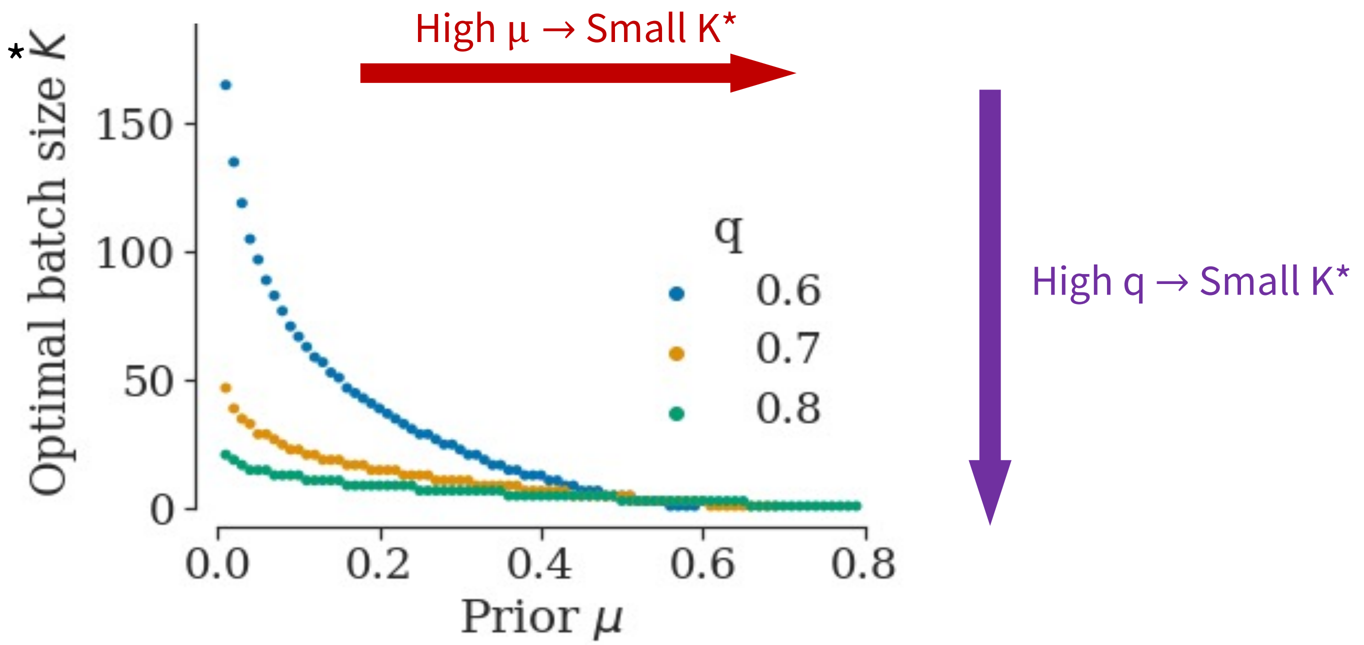

Figure 2. Optimal batch sizes for all possible priors μ for three information regimes $q \in \{0.6, 0.7, 0.8\}$

VOTING WORKS WELL, EVEN IN ITS SIMPLEST FORM

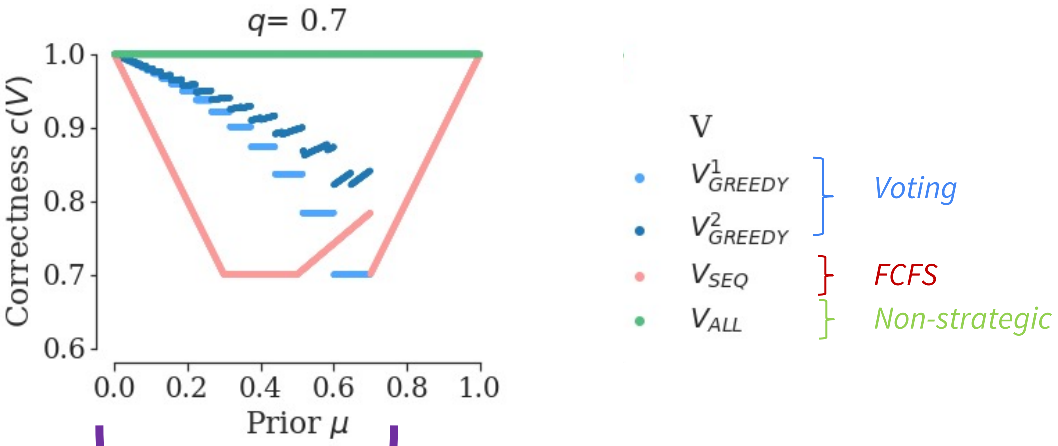

For $\mu < q$, there's always a voting mechanism that outperforms FCFS

Figure 3. Comparison of correctness in different mechanisms simulated with 345 agents.

VOTING WORKS WELL, EVEN IN ITS SIMPLEST FORM

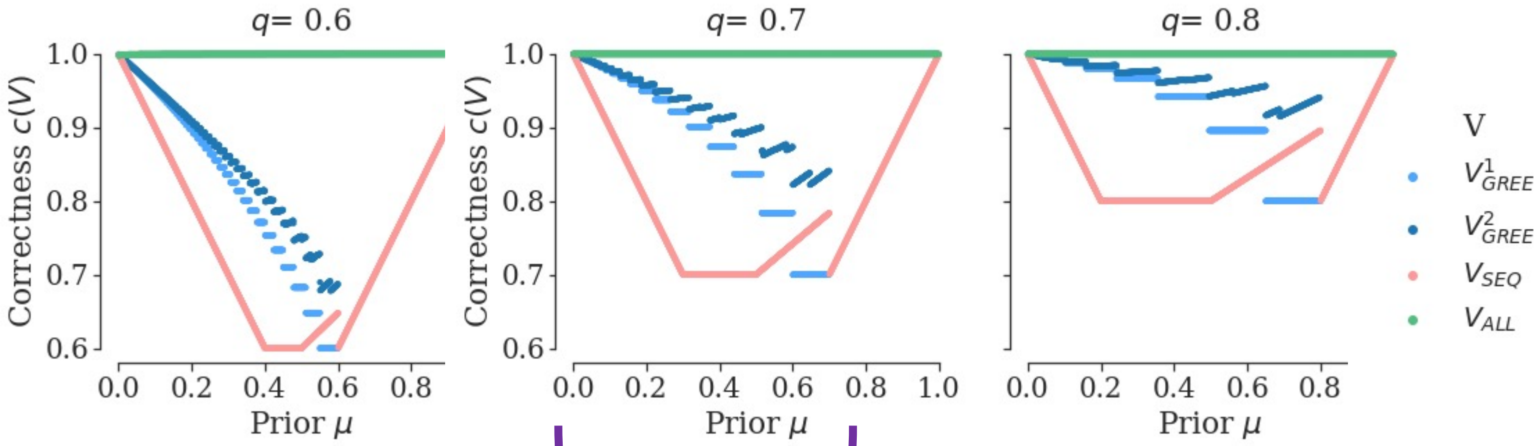

For $\mu < q$, there's always a voting mechanism that outperforms FCFS

Figure 3. Comparison of correctness in different mechanisms simulated with 345 agents.

Kang, *Monachou, Koren, Ashlagi (NeurIPS 2021)*

CONCLUSION

Main takeaways

- **Tension** between: Planner's learning goal *vs* Agents' strategic incentives
- How to **incorporate voting into mechanism design** to mitigate this tension
- In particular, by introducing **batching** and **randomness**

Implications

- Bayesian risk adjustment for organ allocation markets
- Analysis of learning problems with **strategic samples**
- Resembles **exploitation vs exploration**

Limitations

- This is a stylized model!
- Fairness? Voting mechanism (partly) breaks priority for better welfare