Allocating Scarce Societal Resources Based on Predictions of Outcomes

Sanmay Das
Washington University in St. Louis

CSE 591, October 31, 2018

Introduction

- "Resources" that are controlled by or regulated by society are scarce; often cannot rely on market mechanisms
  - Shelter beds and services for homeless households
  - Organs for transplantation
  - Public school spaces, ...
- How can we best allocate these resources to those who need them? Complex problem – we must (at least):
  - Predict outcomes
  - Consider preferences and incentives
  - Define objectives (efficiency, equity, justice/fairness)
- Today: Two case studies
  - **Living donor kidney transplantation**
    - (With Zhuoshu Li, Sofia Carrillo, William Macke, Kelsey Lieberman, Chien-Ju Ho, and Jason Wellen)
  - **Homelessness services**
    - (With Amanda Kube and Patrick Fowler)

Introduction

- "Resources" that are controlled by or regulated by society are scarce; often cannot rely on market mechanisms
  - Shelter beds and services for homeless households
  - Organs for transplantation
  - Public school spaces, ...
- How can we best allocate these resources to those who need them? Complex problem – we must (at least):
  - Predict outcomes
  - Consider preferences and incentives
  - Define objectives (efficiency, equity, justice/fairness)
- Today: Two case studies
  - **Living donor kidney transplantation**
    - (With Zhuoshu Li, Sofia Carrillo, William Macke, Kelsey Lieberman, Chien-Ju Ho, and Jason Wellen)
  - **Homelessness services**
    - (With Amanda Kube and Patrick Fowler)

Introduction

- "Resources" that are controlled by or regulated by society are scarce; often cannot rely on market mechanisms
  - Shelter beds and services for homeless households
  - Organs for transplantation
  - Public school spaces, ...
- How can we best allocate these resources to those who need them? Complex problem – we must (at least):
  - Predict outcomes
  - Consider preferences and incentives
  - Define objectives (efficiency, equity, justice/fairness)
- Today: Two case studies
  - **Living donor kidney transplantation**
    - (With Zhuoshu Li, Sofia Carrillo, William Macke, Kelsey Lieberman, Chien-Ju Ho, and Jason Wellen)
  - **Homelessness services**
    - (With Amanda Kube and Patrick Fowler)

Case Study 1: Living Donor Kidney Transplantation

- About $100,000$ people waiting for kidney transplants in the US (2016)
- About, $19,500$ kidney transplants in recent years, ~ 5500 from living donors
- Unfortunately, willing living donors are often not medically compatible.
- One option for them is to enter a *kidney exchange* program

Kidney Exchange

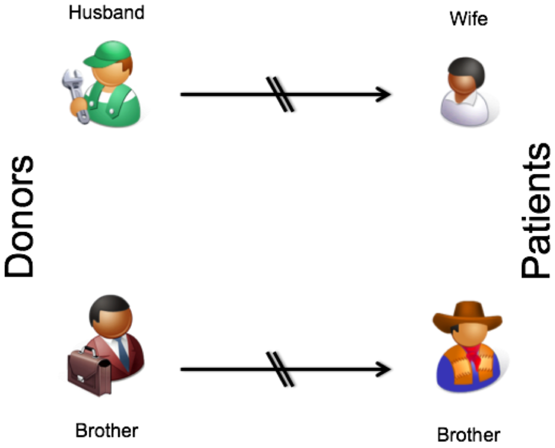

Kidney Exchange

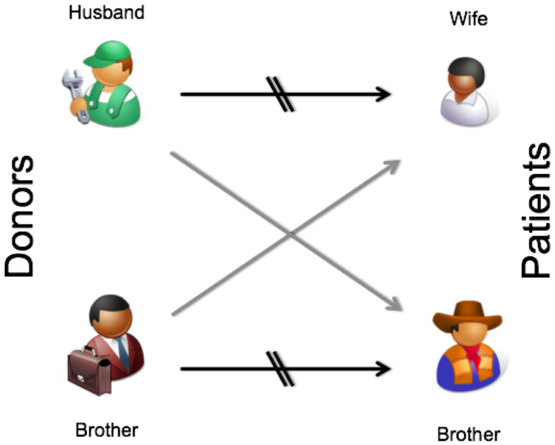

Kidney Exchange

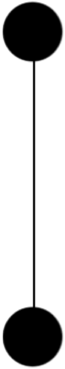

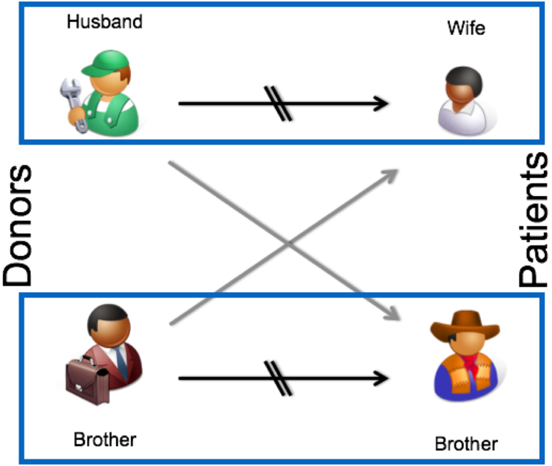

Kidney Exchange in Practice

Problems

- A raft of coordination problems
- Exchange fragmentation

Parts of the solution

- More pooling of pairs (national/international exchanges)
- Desensitization / ABO incompatible transplants
- Today: **Incorporate compatible pairs into exchanges (Gentry et al., 2007)**

Incorporating Compatible Pairs

- Why would a compatible pair want to enter the exchange? (cf. (Anshelevich, Das, and Naamad, 2013))

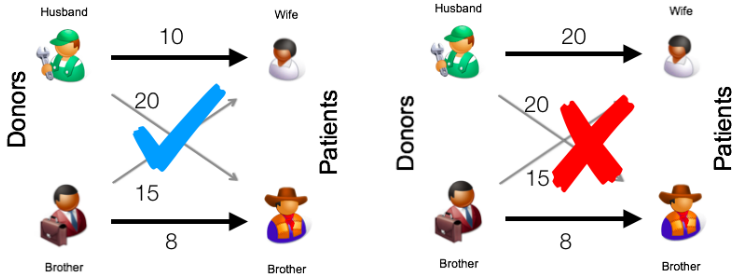

Measuring Match Quality: LKDPI (Massie et al., 2016)

LKDPI Score:

9

This model calculates a risk score for a recipient of a potential live donor kidney.

Live Donor Characteristics:

Donor age:	43
Donor sex:	male
Recipient sex:	female
Donor eGFR:	95
Donor SBP:	130
Donor BMI:	24
Donor is African-American:	No
Donor history of cigarette use:	No
Donor and recipient biologically related:	Yes
Donor and recipient are ABO incompatible:	No
Donor/Recipient Weight Ratio:	0.90 or higher
Donor and recipient HLA-B mismatches:	1
Donor and recipient HLA-DR	

From LKDPI to Graft Survival

- Expected graft survival: estimated as a function of LKDPI $14.78e^{-0.01239LKDPI}$

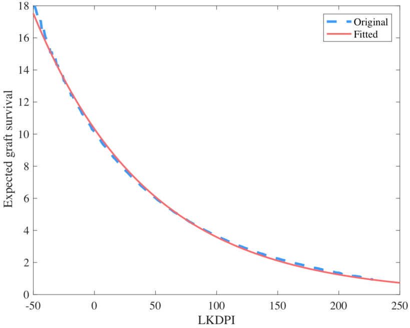

Single Center Analysis

- De-identified data from 2014 - 2016
  - All donor and recipient characteristics for calculating LKDPI

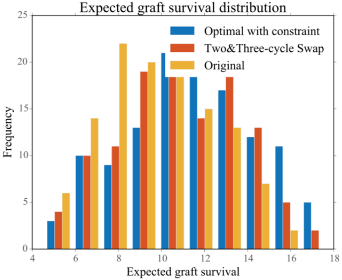

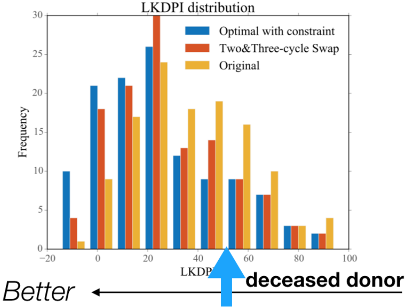

Heterogeneity of Match Quality

	LKDPI original	LKDPI 2&3 swap	LKDPI Optimal
Original 166 dataset	37.15	25.50	22.46

Heterogeneity of Match Quality

	LKDPI original	LKDPI 2&3 swap	LKDPI Optimal
Original 166 dataset	37.15	25.50	22.46
Sample from the whole matrix	40.51	2.67	-2.5

Heterogeneity of Match Quality

	LKDPI original	LKDPI 2&3 swap	LKDPI Optimal
Original 166 dataset	37.15	25.50	22.46
Sample from the whole matrix	40.51	2.67	-2.5
Shuffle all donors per recipient	40.92	4.11	-0.47

Heterogeneity of Match Quality

	LKDPI original	LKDPI 2&3 swap	LKDPI Optimal
Original 166 dataset	37.15	25.50	22.46
Sample from the whole matrix	40.51	2.67	-2.5
Shuffle all donors per recipient	40.92	4.11	-0.47
Shuffle all recipients per donor	40.70	20.6	15.49

Heterogeneity of Match Quality

	LKDPI original	LKDPI 2&3 swap	LKDPI Optimal
Original 166 dataset	37.15	25.50	22.46
Sample from the whole matrix	40.51	2.67	-2.5
Shuffle all donors per recipient	40.92	4.11	-0.47
Shuffle all recipients per donor	40.70	20.6	15.49

Takeaway: Largely donor driven, but with some pairwise idiosyncracies

Simulator

- To analyze the effects of policy changes, we need a faithful simulation of the real process
- Basic simulator model:
  - Generate LKDPI-related characteristics to measure expected graft survival
  - Compatibility based on the simulator from Saidman et al. (2006)

Simulator: Initial Assessment

	LKDPI original	LKDPI 2&3 swap	LKDPI Optimal
Original 166 dataset	37.15	25.50	22.46
Sample from the whole matrix	40.51	2.67	-2.5
Shuffle all donors per recipient	40.92	4.11	-0.47
Shuffle all recipients per donor	40.70	20.6	15.49
Sample from our simulator	39.21	24.50	20.09

Including Compatible Pairs in Kidney Exchange

- Including compatible pairs to thicken the exchange with incompatible pairs
  - Increase in the number of matches for incompatible pairs (quantity)
  - Increase in the expected graft survival for compatible pairs (quality)

Batch Optimization

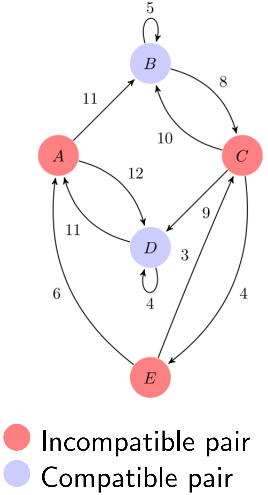

- Simulated population: Any size
  - ◇ Compatible & incompatible pairs
  - ◇ Expected graft survival graph
- Optimization goal
  - ◇ Sum of expected graft survivals: A-D, B-C
  - ◇ Expected number of matches: A-D, B, C-E

Batch Optimization Results

- Increase in number of matches for incompatible pairs (quantity)

	Without compatible	With compatible
Size of pool: 50 (25+25)	33%	64%
Size of pool: 100 (50+50)	40%	76%
Size of pool: 1000 (500+500)	53%	95%

- Increase in expected graft survival for compatible pairs (quality)

	EGS of compatible pairs[1]
Max expected survival	2.04 - 2.36
Max # of matched pairs	1.20 - 1.59

[1]Whose assignments changed

Batch Optimization Results

- Increase in number of matches for incompatible pairs (quantity)

	Without compatible	With compatible
Size of pool: 50 (25+25)	33%	64%
Size of pool: 100 (50+50)	40%	76%
Size of pool: 1000 (500+500)	53%	95%

- Increase in expected graft survival for compatible pairs (quality)

	EGS of compatible pairs[1]
Max expected survival	2.04 - 2.36
Max # of matched pairs	1.20 - 1.59

[1]Whose assignments changed

Dynamic Matching

- Compatible pairs may not be willing to wait any longer than necessary
- Also debate in the literature about the value of patience regardless (Akbarpour, S. Li, and Oveis Gharan, 2017; Ashlagi et al., 2017; Z. Li et al., 2018)
- New model: Impatient compatible pairs and a pool of patient incompatible pairs

Hybrid Static-Dynamic Matching Model

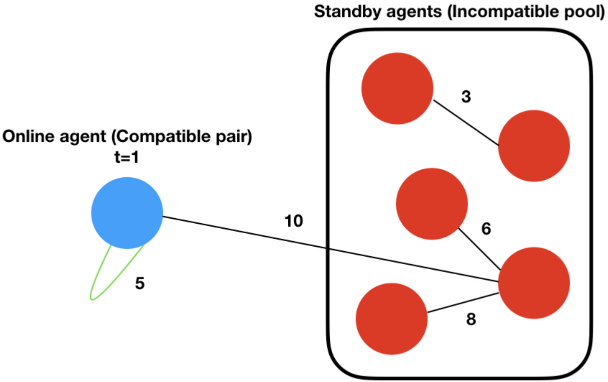

Hybrid Static-Dynamic Matching Model

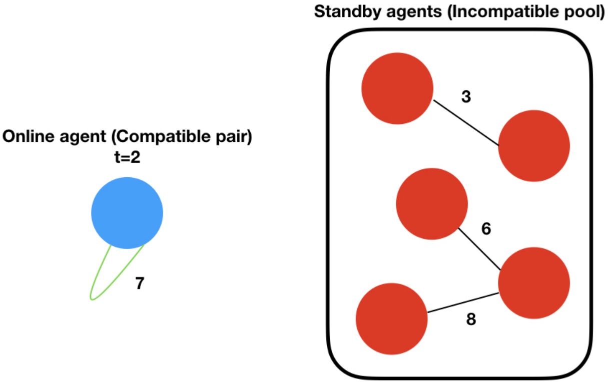

Hybrid Static-Dynamic Matching Model

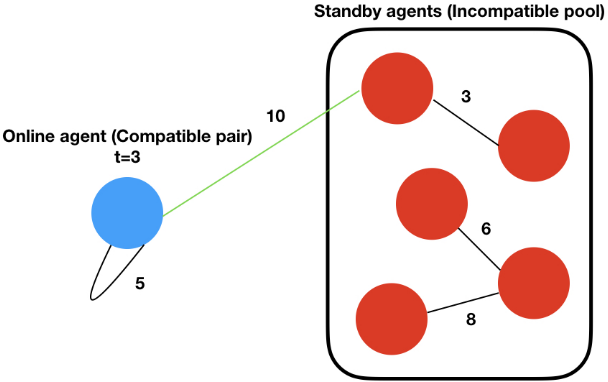

Hybrid Static-Dynamic Matching Model

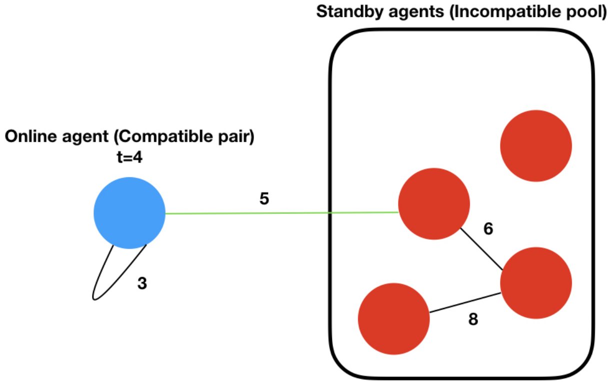

Hybrid Static-Dynamic Matching Model

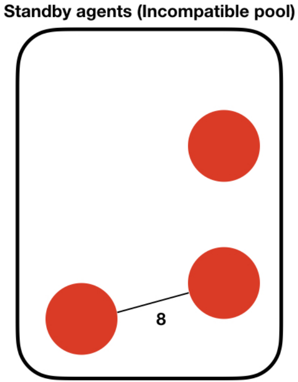

An Oracle for 2-Matching

$$\max \sum_{n=1}^{N} \sum_{i=0}^{I} w_{n,i} x_{n,i}$$

$$\text{s.t.} \sum_{i=0}^{I} x_{n,i} \leq 1, \forall n \in [T]$$

$$\sum_{n=1}^{N} x_{n,i} + \sum_{j=1}^{I} x_{T+i,j} \leq 1, \forall i \in [I]$$

$$x_{n,i} \in \{0,1\}, \forall n \in [N], \forall i \in [I]^*$$

- w's: weights; x's: match variables.
- When $i = 0$, $x_{n,0}$ represents a self-match of agent n.
- When $i > 0$ and $n \leq T$, $x_{n,i}$ represents a match between online n and standby i.
- When $i > 0$ and $n > T$, $x_{n,i}$ represents a match between standby $j = n - T$ and standby i

Dual Formulation and the ODASSE Algorithm

$$
\begin{aligned}
\min \ & \sum_{t=1}^{T} \alpha_t + \sum_{i=0}^{I} \beta_i \\
\text{s.t. } & w_{t,i} - \alpha_t - \beta_i \leq 0, \forall t \in [T], i \in [I]^* \\
& w_{t+j,i} - \beta_j - \beta_i \leq 0, \forall i \in [I], j \in [I] \\
& \alpha_t, \beta_i \geq 0, \forall t \in [T], i \in [I] \\
& \beta_0 = 0
\end{aligned}
$$

- α_t, β_i can be interpreted as estimated values (*shadow survival estimates*) of compatible pairs and incompatible pairs respectively.
- Given optimal β_i^* we can derive the online assignment rule $i^* = \text{argmax}_i\{w_{t,i} - \beta_i^*\}$ (*Online Dual Assignment Using Shadow Survival Estimates*).

Estimating β_i^*

- Run many simulations and get β_i^* values
- Train a linear model on
  - Demographic information of an incompatible pair
  - Initial graph state of incompatible pairs (β_i value when solving the dual on just the incompatible pool).
- Predicted vs. true β^* values.

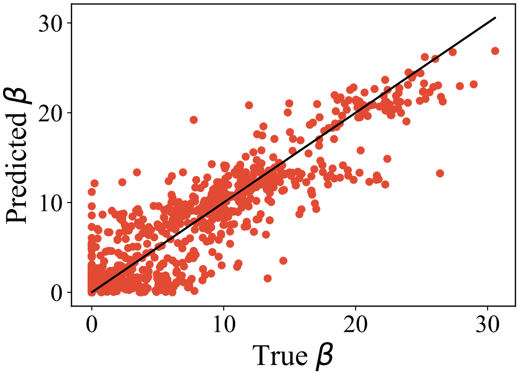

Results

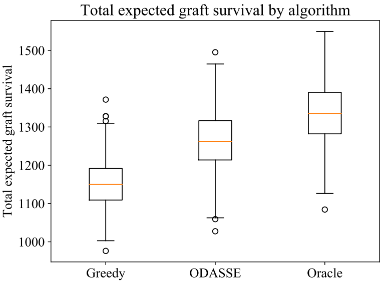

	Original	Greedy	ODASSE	Oracle
Matched proportion of incompatible pairs	53%	61%	72%	76%
Expected graft survival of compatible pairs	9.65	11.13	11.16	11.39
Expected graft survival of incompatible pairs	10.32	9.75	9.80	9.99

Results

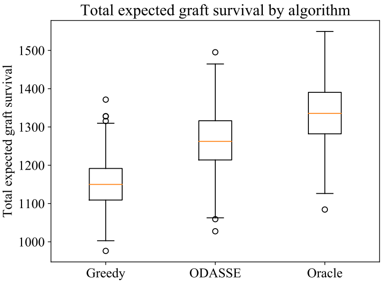

	Original	Greedy	ODASSE	Oracle
Matched proportion of incompatible pairs	53%	61%	72%	76%
Expected graft survival of compatible pairs	9.65	11.13	11.16	11.39
Expected graft survival of incompatible pairs	10.32	9.75	9.80	9.99

Results: Disadvantaged Populations

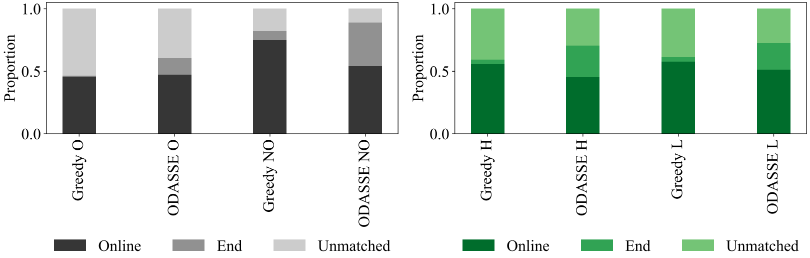

Overall benefits (compared with no compatibles) are disproportionately good for Type O, and proportional for High PRA patients.

Results: Disadvantaged Populations

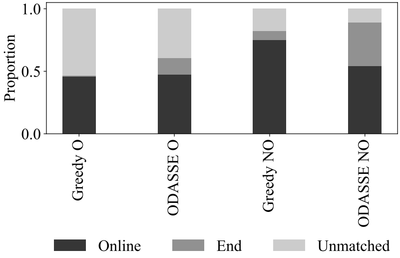

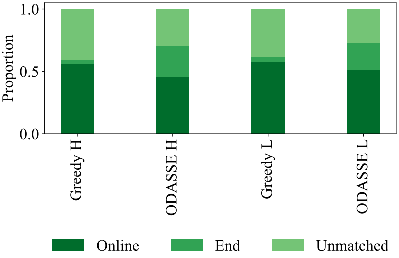

Overall benefits (compared with no compatibles) are disproportionately good for Type O, and proportional for High PRA patients.

Discussion

- Quantifying benefits allows us to think about a richer mechanism that includes compatible pairs in exchanges.
- We estimate substantial benefits in terms of number of incompatible pairs matched and increase in graft survival for compatible pairs.
- Methodological directions:
  - A model with real weights for weighted matching algorithms to work on!
  - A new hybrid static-dynamic matching model.
  - Online primal-dual + learning algorithm

Case Study 2: Homelessness Services

- More than 1.4 million people used services in the US in 2016
- System struggles to keep up with demand
- Yet, limited assessment of efficacy of allocations

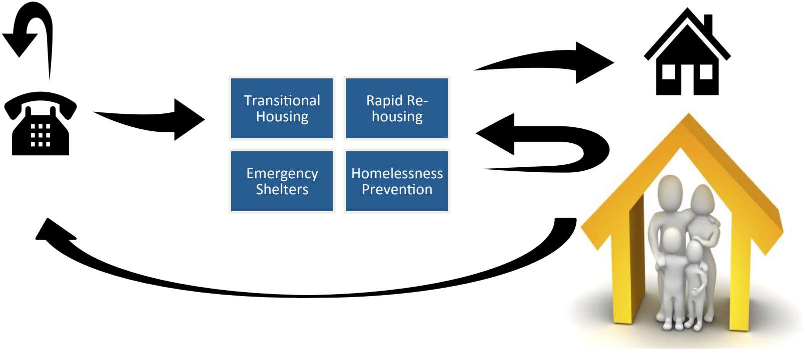

Improving Allocations Using Counterfactual Predictions

- Idea: Personalized intervention / resource allocation
- Estimate how well a household would have done if allocated to one of several different possible interventions
  - Measure: Probability of re-entry within two years of exit
  - Need: Causal / counterfactual prediction
- We use detailed demographic and assessment data from 58 different homeless agencies in a major metro area.
- Use an ensemble method called BART to estimate counterfactual probabilities of re-entry (Chipman, George, McCulloch, et al., 2010; Hill, 2011)
- Optimize allocations on a weekly basis

Improving Allocations Using Counterfactual Predictions

- Idea: Personalized intervention / resource allocation
- Estimate how well a household would have done if allocated to one of several different possible interventions
  - Measure: Probability of re-entry within two years of exit
  - Need: Causal / counterfactual prediction
- We use detailed demographic and assessment data from 58 different homeless agencies in a major metro area.
- Use an ensemble method called BART to estimate counterfactual probabilities of re-entry (Chipman, George, McCulloch, et al., 2010; Hill, 2011)
- Optimize allocations on a weekly basis

Data

Service Type	Number Assigned	Percent Reentered
Emergency Shelter	2897	56.20
Transitional Housing	1927	40.22
Rapid Rehousing	589	53.48
Homelessness Prevention	2061	24.16
Total	7474	43.03

Type	Number	Examples
Binary	3	Gender, Spouse Present, HUD Chronic Homeless
Other Categorical	61	Veteran, Disabling Condition, Substance Abuse
Continuous	4	Age, Income, Calls to Hotline, Duration of Wait
Total Features	68	

Data

Service Type	Number Assigned	Percent Reentered
Emergency Shelter	2897	56.20
Transitional Housing	1927	40.22
Rapid Rehousing	589	53.48
Homelessness Prevention	2061	24.16
Total	7474	43.03

Type	Number	Examples
Binary	3	Gender, Spouse Present, HUD Chronic Homeless
Other Categorical	61	Veteran, Disabling Condition, Substance Abuse
Continuous	4	Age, Income, Calls to Hotline, Duration of Wait
Total Features	68	

Optimal Allocation

Optimization Problem

$$\min_{x_{ij}} \sum_i \sum_j p_{ij} x_{ij}$$

$$\text{subject to} \quad \sum_j x_{ij} = 1 \quad \forall i$$

$$\sum_i x_{ij} \leq C_j \quad \forall j$$

- x_{ij}: whether or not household i is placed in intervention j
- p_{ij}: probability of household i reentering if they are placed in intervention j
- C_j: capacity of intervention j

Results

- We estimate capacities and re-allocate among interventions weekly (for 166 weeks).
- Reduces number of re-entries from 2193 households (43.04%) to 1624 in expectation (31.88%) – a 27.08% reduction!
- BART predicts 2227 re-entries out-of-sample, so empirically relatively unbiased.

Optimal Allocation

Optimization Problem

$$\min_{x_{ij}} \sum_i \sum_j p_{ij} x_{ij}$$

$$\text{subject to} \quad \sum_j x_{ij} = 1 \quad \forall i$$

$$\sum_i x_{ij} \leq C_j \quad \forall j$$

- x_{ij}: whether or not household i is placed in intervention j
- p_{ij}: probability of household i reentering if they are placed in intervention j
- C_j: capacity of intervention j

Results

- We estimate capacities and re-allocate among interventions weekly (for 166 weeks).
- Reduces number of re-entries from 2193 households (43.04%) to 1624 in expectation (31.88%) – a 27.08% reduction!
- BART predicts 2227 re-entries out-of-sample, so empirically relatively unbiased.

Optimal Allocation

Optimization Problem

$$\min_{x_{ij}} \sum_i \sum_j p_{ij} x_{ij}$$

$$\text{subject to} \quad \sum_j x_{ij} = 1 \quad \forall i$$

$$\sum_i x_{ij} \leq C_j \quad \forall j$$

- x_{ij}: whether or not household i is placed in intervention j
- p_{ij}: probability of household i reentering if they are placed in intervention j
- C_j: capacity of intervention j

Results

- We estimate capacities and re-allocate among interventions weekly (for 166 weeks).
- Reduces number of re-entries from 2193 households (43.04%) to 1624 in expectation (31.88%) – a 27.08% reduction!
- BART predicts 2227 re-entries out-of-sample, so empirically relatively unbiased.

Fairness

The optimal allocation hurts as many households as it helps, it just helps **more** overall

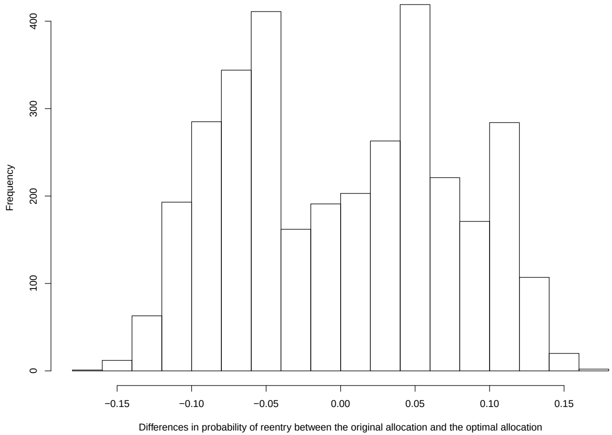

Who is Helped and Hurt?

- We use machine learning techniques to learn whether a household is likely to be helped or hurt in the new allocation.
- Then find the features that are most predictive and analyze them
- The optimal allocation seems to help households that are more in need:
  - Lower monthly incomes
  - Longer waits and more calls to the hotline before being placed
  - More substance abuse problems

Fairness Constraints

- Allocations may be because of policy constraints
  - E.g. require prioritization of those fleeing domestic abuse
- We can require the allocation to not hurt anyone more than a small percentage in expectation
- Add a constraint
  - $$\sum_j p_{ij} x_{ij} \leq \sum_j p_{ij} y_{ij} + 0.05 \; \forall i$$
  - y_{ij} represents whether or not household i was originally placed in intervention j

"Fairer" Allocation

- Now 1904 households (37.38%) reenter the system within two years.
  - Higher than the optimized allocation without the constraint, but still a 14.66% decrease
  - Less room for improvement under constraints

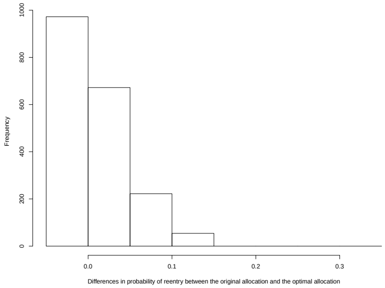

Looking Forward

- Homelessness system itself
  - Different constraints (confidence in counterfactual?)
  - Online matching!
  - Richer sets of resources for allocation (counseling, beds, cash transfers, etc)?
  - Plan for paths through the system (shelter $\rightarrow$ transitional housing, e.g.)
- Bigger picture:
  - Getting the conversation started
  - How can we use data and AI in the service of efficiency, equity, and justice in society?
  - Interplay between (dynamic) optimization and prediction, combined with consideration of long-run incentives is key
  - Ethical and systemic issues must be primary

Looking Forward

- Homelessness system itself
  - Different constraints (confidence in counterfactual?)
  - Online matching!
  - Richer sets of resources for allocation (counseling, beds, cash transfers, etc)?
  - Plan for paths through the system (shelter $\rightarrow$ transitional housing, e.g.)
- Bigger picture:
  - Getting the conversation started
  - How can we use data and AI in the service of efficiency, equity, and justice in society?
  - Interplay between (dynamic) optimization and prediction, combined with consideration of long-run incentives is key
  - Ethical and systemic issues must be primary