Allocating Interventions Based on Counterfactual Predictions: A Case Study on Homelessness Services

Amanda R. Kube, Sanmay Das and **Patrick J. Fowler**
Washington University in St. Louis
amanda.kube@wustl.edu, sanmay@wustl.edu, pjfowler@wustl.edu

Abstract

Modern statistical and machine learning methods are increasingly capable of modeling individual or personalized treatment effects by predicting counterfactual outcomes. These counterfactual predictions could be used to allocate different interventions across populations based on individual characteristics. In many domains, like social services, the availability of different possible interventions can be severely resource limited. This paper considers possible improvements to the allocation of such services in the context of homelessness service provision in a major metropolitan area. Using data from the homeless system, we show potential for substantial predicted benefits in terms of reducing the number of families who experience repeat episodes of homelessness by choosing optimal allocations (based on predicted outcomes) to a fixed number of beds in different types of homelessness service facilities. Such changes in the allocation mechanism would not be without tradeoffs, however; a significant fraction of households are predicted to have a higher probability of re-entry in the optimal allocation than in the original one. We discuss the efficiency, equity and fairness issues that arise and consider potential implications for policy.

1 Introduction

Homelessness represents a long-standing problem with considerable individual and social costs. Homeless services coordinated at the community level (i.e, homeless system) struggle to keep up with demand for housing assistance, and little evidence supports the accuracy of current decision making in the allocation of limited homeless services [Brown *et al.*, 2018; Fowler *et al.*, 2017; Shinn *et al.*, 2013]. Advances in machine learning and AI techniques have made it possible to apply learning algorithms to social problems ranging from police patrol to poaching. Many of these solutions have had success in mitigating the problem to which they were applied [McCarthy *et al.*, 2017; Tambe *et al.*, 2016, e.g.]. In this paper, we test the feasibility of data-driven approaches to inform policies that guide homeless service delivery. Specifically we ask the question of whether one can use individual predictions of success for households assigned to certain types of homeless services to improve outcomes across the population.

Motivation: The use of techniques from artificial intelligence and machine learning (and more broadly, algorithmic approaches) for decision making on resource allocation increasingly raise concerns as applications extend into different social contexts [O'Neil, 2016]. Although demonstrating potential for improvements in efficiency, applications present issues regarding fairness, accountability, and transparency. A number of examples illustrate data-driven allocations that unintentionally introduce systematic biases that perpetuate inequities, such as racial disparities in credit lending, hotspot policing, and crime sentencing [Ensign *et al.*, 2017; Pleiss *et al.*, 2017; Corbett-Davies *et al.*, 2017]. The complexity involved in the development of decision algorithms has called into question the ability to design adequate protections against systematic misuses. In response to these concerns, the European Union recently issued the "General Data Protection Regulation" (GDPR), which imposes restrictions on how individual data can be used for algorithmic decision making in ways that "significantly affect" users. The GDPR coincides with a broader argument for not just full transparency, but rather *human interpretability* regarding how decisions are derived from algorithmic approaches to ensure adequate assessment of fairness.

A counter argument suggests requirements for human interpretability threaten to diminish the potential of AI to solve societal problems. Algorithmic approaches generate novel solutions that escape direct observation; requirements for full explainability of these complex processes limits the inherent value of applications to thorny social problems. In a recent *Wired* op-ed, David Weinberger raises a compelling example related to autonomous vehicles. If they were able to lower the number of fatalities in US vehicle crashes by 90%, would it really be worth losing that benefit because of the difficulty of explaining (or legal liabilities that may be associated with) the remaining crashes? Certainly, the answer to this partly depends on whether the remaining crashes disproportionately affect some portion of the population, and perhaps other considerations. Weinberger goes on to argue that while the governance of AI applied to social problems is critical, it can be achieved through existing processes for resolving policy issues [Weinberger, 2018]. The right approach is then to specify appropriate optimization goals, arrived at through the social processes of policy-making, which could be based on both efficiency and equity considerations. Limited capacity to meet demand for social services further emphasizes the value of policy-making processes for data-driven resource allocations —as is the case with homeless assistance.

A key difference in making resource allocation decisions on the basis of predictions in the social services setting, when predictions are being made based on observational data, is that the importance of causal modeling is magnified. As opposed to the types of problems that Kleinberg *et al.* (2015) call "prediction policy problems", or for example using machine learning predictions of default to manage risk [Butaru *et al.*, 2016], we need useful counterfactual estimates of the effects of different interventions in order to even define the resource allocation problem. There has been significant recent progress in causal modeling from a machine learning perspective [Johansson *et al.*, 2016, e.g.]. For our work we use Bayesian additive regression trees (BART) [Chipman *et al.*, 2007; 2010] which have the benefit of providing coherent probabilistic estimates of heterogeneous treatment effects. Thus, it allows us to predict individual outcomes under counterfactual allocations.

While there is a long history of mechanism design research on assignment problems, school allocation, organ allocation, refugee matching etc (Kominers *et al.* (2017) provide an excellent recent introduction to market design), and much recent interest in the AI and broader computer science community in mechanism design for social good[1], there has been limited prior work on homelessness specifically. The most relevant study is that of Azizi *et al.* (2018), who consider allocation policies specifically for homeless youth. They formulate a dynamic allocation problem between arriving homeless youth and two types of housing resources (rapid rehousing and permanent supportive housing) in order to fairly and efficiently allocate youth to these resources; our focus moves beyond accurate screening to forecast response to multiple interventions using counterfactual approaches. Ours is one of the first studies to consider using machine-learning based estimates of counterfactual outcome probabilities to estimate the value of, and thus inform, allocation decisions for homeless services. We present this work as a proof-of-concept, based on a real administrative dataset across the whole range of homeless populations in a metro area, to address the following question: *By optimizing allocations based on predicted outcomes, how much could we potentially improve outcomes, and what would be the distributional effects of these improvements?*

Problem setup: Community-wide responses to homelessness coordinate services that vary in level of intensity to meet household needs. In the US, services range from time-limited nonresidential supports to ongoing rental assistance with intensive case management [Congress, 2009]. At any given time, the homeless system allocates many households to many interventions, each subject to capacity constraints. Fundamentally, homeless services aim to stabilize households and reduce future demand for assistance. One metric of successful services tracks the number of households that use additional homeless services within two years of initial contact; counts are generated from administrative data that record entries and exists across homeless services [HUD, 2012]. However, routine capacity constraints for needed services within the homeless system conflates success.

The present study takes advantage of unique local administrative records to assess whether households reenter homeless assistance within two years of initial contact, regardless of using services. Data link homeless service records with requests for assistance through a regional homeless hotline. Bayesian additive regression trees (BART) models account for capacities of different interventions over time, as well as build counterfactual estimates of the probability that a household reenters the system if placed in a particular intervention. We formulate the optimization problem of the homeless system as minimizing the expected number of households that reenter the system within two years, subject to capacity constraints on each intervention.

Preview of results: Using administrative data on a weekly basis over the course of 166 weeks, we find that the BART model predicted that, in expectation, 3223 (58.83%) of households would re-enter the system, and 3205 (58.51%) actually did. In the optimized assignment we find, the BART model predicts that only 2483 households (45.33%) would re-enter the system. Thus, there may be substantial benefits achievable (by this re-entry metric) from improving the combined prediction-allocation mechanism. However, these benefits do not come without tradeoffs. They are not even close to pareto-improving. In fact, more households increase their probability of re-entry, according to the predictions, than those that decrease their probability of re-entry. We also formulate and solve a constrained version of the allocation problem, which guarantees that no household increases their probability of re-entry by more than 5 percentage points in the new allocation. In this case, 56.33% of households are predicted to re-enter.

Implications: Our work is intended as a proof of concept and a case study. We bring data to bear on the question of how much AI techniques can improve social service provision, with full awareness that the precise results presented may depend on specific modeling choices, and the reliability of the counterfactual estimates. We expect this work to contribute to the emerging dialogue on intervening based on machine learning predictions. It is very important to consider fairness, ethics, and the long-term dynamics of systems that use these kinds of predictive modules. At the same time, the current state of practice in social services allocation is far from evidence-based; therefore, not engaging these questions with actual data and estimates could be leading to significant societal harm.

2 Background and Data

Homelessness represents a complex public health challenge for communities across the United States. Federal guidelines define homelessness as residence in unstable and nonpermanent accommodations. This includes shelters, places not meant for habitation (eg., cars, park, abandoned buildings), as well as being at imminent risk for eviction. Counts estimate that more than 550,000 people experienced home-

[1]For example, see the ACM EC workshops on Mechanism Design for Social Good in 2017 and 2018.

lessness in the United States on a single night in January, 2016 [Henry *et al.*, 2016], and 1.4 million people used homeless services at some point during the year [Solari *et al.*, 2016]. Families with children under 18 years of age comprised 35% of the homeless population. Experiences of homelessness and associated turmoil carries life long implications, as well as significant social costs [Khadduri *et al.*, 2010; Culhane *et al.*, 2011].

The homeless system represents the primary community-wide response to housing crises. Funds allocated by Congress on an annual basis support the delivery of five types of homeless assistance. Service types vary in intensity, and relatedly, availability. The most intensive service - Permanent Supportive Housing - provides long-term rental assistance plus comprehensive case management to address barriers to stability, such as mental health and substance abuse treatment; it is reserved for the highest risk households and consumes the greatest amount of financial resources. Transitional Housing also offers comprehensive case management but only up to 24 months in congregate settings. Rapid Rehousing allows up to 24 months of rental assistance without additional intensive case management. At the end of two years, households in Transitional Housing or Rapid Rehousing either move on their own or step-up to Permanent Supportive Housing, if available. Emergency Shelters offer immediate accommodations for those with no other place to go, and typically serve a large number of households for a brief period of time. Shelters are intended to stabilize households and divert high-risk families to the longer-term housing interventions. Finally, Homelessness Prevention provides households at imminent risk for homelessness with short-term and non-reoccurring assistance to mitigate housing crises. Local non-profit provider networks determine the delivery of day-to-day services within general structures determined by federal funding priorities.

Despite substantial investments, homeless rates remain stubbornly high in the United States. An enormous challenge is that of matching service types to need. While federal guidelines mandate that local agencies provide services based on risk assessments [Congress, 2009], existing tools fail to discern high and low risk households beyond chance [Brown *et al.*, 2018; Shinn *et al.*, 2013]. Homeless service providers have limited evidence for adapting responses to household characteristics. Moreover, there are no tools that assess the impact of service matches on overall system performance in reducing reentries.[2]

Algorithmic approaches offer substantial promise for addressing the optimization of homeless service delivery. Administrative records systematically track service usage and household characteristics over time, and provide rich sources of information from which to glean insights into service improvements. Therefore, potential exists to evaluate improvements in prediction that support decision making. However, as mentioned above, the application of data-driven approaches for delivery of scarce resources to address homelessness requires careful consideration of fairness. The feasible application of any algorithms must be transparent and assess unintended sources of bias [O'Neil, 2016].

[2]Annual evaluations of homeless system performance monitor rates of return to the homeless system within 24 months; future federal funding depends in part on demonstrating trends toward reductions in reentries.

2.1 Data Collection

Data for the project come from the homeless management information system (HMIS) of a major metropolitan area from 2007 through 2014. The HMIS records all housing services provided to individuals and families seeking federally funded homelessness assistance. Local service providers enter information on requests and receipt of services in real time through a web-based platform in accordance with federal mandates for collection of universal elements. A local non-profit organization contracted with the homeless system hosts the platform and provides support, including user training, technical assistance, and active quality control.

Records provide information on the characteristics and services delivered to households in contact with the homeless system. Household-level characteristics includes an array of information on demographics, housing risk assessments, and eligibility determinations. Services include entry and exit dates from the five federally defined types of homeless assistance: homelessness prevention, emergency shelter, rapid rehousing, transitional housing, and permanent supportive housing. In addition, the metropolitan area coordinates requests for assistance through a homeless hotline, and household-level data record information on every call, including dates and referral for services. Household identifiers allow linkages of information across time. Data sharing agreements with regional homeless systems allow access to de-identified records in accordance with the Washington University Institutional Review Board.

2.2 Data Cleaning and Feature Selection

For this project, we extract data provided by 58 different homeless agencies and link participants across programs by a unique, anonymous identification number. We then aggregate data by household using a unique household identification number. This results in a dataset of households containing household characteristics available upon entry into the system, as well as information on all entries and exits from different homeless services. Permanent supportive housing is meant as an intervention that households transition into after a certain period of time or the conclusion of a particular intervention. Because our dataset focuses on first entries into the homeless system, we exclude permanent supportive housing from our analyses due to the nature of the intervention. The primary outcome (the label we are trying to predict) is reentry into the homeless system. Operationally, reentry is defined as requesting services within two years of exit from the system, regardless of whether services were actually received. This ensures that we capture further need, and not just availability of services. When transitions between services (e.g. homeless shelter to rapid rehousing) occur on the same day, we assume that they represent a continuation of homeless services. We consider households to have exited from the system when the time between leaving one service and entering another exceeds one day. Our analyses include households who entered the homeless system after the start of 2007 and exited before the end of 2012 to provide a minimum two-year follow-up for all households.

Since the data captures homeless services across time, it contains both time-invariant (e.g., race, gender, ethnicity) as well as time-variant (e.g., monthly income, age) features. We select values of time-variant features that are collected at the time of first entry into the homeless system and have adequate amounts of available data for use in our model. Most of the variables we selected were categorical, and missing values are treated as a separate category in these cases.

2.3 Data Characteristics

The analytic dataset includes records on 8080 households. Of these 8080, 4577 (56.65%) reentered the homeless system within two years of exiting. Table 1 shows the number of households assigned to each service type, as well as the percentage reentries within 2 years for each intervention. Of the 4577 who reentered, 1551 (33.89%) were placed in a subsequent service, while 3026 (66.11%) called the hotline for assistance but by the end of the two year period had not been placed in another service.

A single feature vector consists of covariate data for head-of-household, spouse, and children (e.g. race, gender, and disability information) as well as which service type the household was assigned to. The target variable, or label, is a binary indicator of whether or not they reentered the homeless system within 2 years of exiting. Table 2 shows a summary and examples of the features included.

3 Analyzing Interventions

The key decision variable is the choice of intervention to which a household should be allocated. For the larger enterprise proposed in this work to make sense, it is important that different interventions actually have different effects. While Table 1 shows apparent differences in the probability of reentry based on intervention, these differences could be due to unobserved variables or selection bias because of the nonrandom provision of services. Therefore, we start by systematically investigating the differential effects of these housing interventions (homelessness prevention, emergency shelter, rapid rehousing, and transitional housing) on the probability of reentry into homeless services within two years.

This application requires a method that can handle the challenges of counterfactual inference using observational data while simultaneously providing a well-grounded probabilistic model. Bayesian nonparametric modeling for causal inference has a number of advantages that fit this application [Chipman *et al.*, 2010; Hill, 2011; Johansson *et al.*, 2016]. These models provide robust estimates of treatment effects using observational data like administrative service records. They can handle a large number of features or predictors, as well as complex data that include interactions and nonlinearities seen in studies of housing assistance in child welfare. We use BART (Bayesian Additive Regression Trees), an ensemble model that outperforms propensity score and nearest neighbor matching algorithms for causal inference on observational data, especially when the data is complex [Hill, 2011]. BART can also explicitly address heterogeneous response to interventions based on empirically identified features in the data, generating individual treatment effect estimates (or counterfactual predictions) in addition to population-level ones.

3.1 Building the Model

BART [Chipman *et al.*, 2007; 2010] models the data by approximating $f(x) = E(Y|x)$ as a sum of binary regression trees. The sum-of-trees model includes trees of different sizes and allows BART to incorporate both additive and interaction effects of various orders. BART uses a regularization prior to restrain the effect of each tree and then uses a Bayesian backfitting MCMC algorithm to draw samples from the posterior distribution. At the start of the MCMC draws, a chain of single-node trees is instantiated. During each iteration, each tree can increase or decrease its number of nodes or can swap decision rules between a parent node and a child node. Then, BART computes a new sample from the approximated posterior distribution f^* as a sum of the results from the current set of trees. These posterior samples consist of 1000 post-burn-in samples for each observation. Using BART to model the data produces a set of posterior draws *for each household in the dataset*, allowing population-wide as well as household-specific inference. Model fitting and counterfactual inference were done using the R package BayesTree written by the model's creators [Chipman *et al.*, 2010] as well as the package `bartMachine`.

3.2 Population Treatment Effects

We compare service types by doing pairwise inference. We select data for each pair and build a BART model based on this data. We use BART to approximate the posterior distribution of reentry based on this model for the factual service type as well as the counterfactual (if all covariates remain the same but service type changes). Then, we take the mean and 2.5% and 97.5% quantiles of the difference between counterfactual samples and factual samples in order to find treatment effects and 95% estimated credible intervals for service type. We do this for all pairs of service types as well as for rapid rehousing compared to any other service type.

Pairwise differences show that population-wide treatment effects for emergency shelter, transitional housing, and homelessness prevention are not largely different from one another (the 95% Estimated Credible Intervals for prevention versus transitional housing and Emergency shelter versus transitional housing include 0; the treatment effect of homelessness prevention versus emergency shelter is 0.03 with 95% Estimated Credible Interval = [0.003,0.05]). The only pairs for which there seem to be meaningful treatment differences are those that included rapid rehousing: rapid rehousing versus transitional housing (TE = 0.32, 95% Estimated Credible Interval = [0.10,0.42]), rapid rehousing versus emergency shelter (TE = 0.26, 95% Estimated Credible Interval = [0.12,0.38]), and rapid rehousing versus homelessness prevention (TE = 0.16, 95% Estimated Credible Interval = [0.05,0.24]) implying that all services are more effective than rapid rehousing at reducing the probability of reentry within two years.

Service Type	Number Assigned	Percent Reentered
Emergency Shelter	3277	66.19
Transitional Housing	2141	46.29
Rapid Rehousing	785	82.42
Homelessness Prevention	1877	41.02
Total	8080	56.65

Table 1: Summary of assignment to services across the dataset as well as reentry statistics for each type of service

Type	Number	Examples
Binary Features	3	Gender, Spouse Present, HUD Chronic Homeless
Non-Binary Categorical Features	63	Veteran Status, Disabling Condition, Substance Abuse
Continuous Features	4	Age, Monthly Income, Calls to Hotline, Duration of Wait
Total Features	70	

Table 2: Summary of features

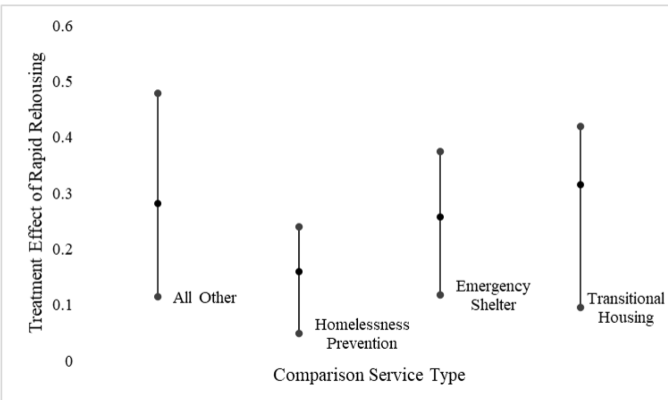

Figure 1: Plot of average treatment effects as well as upper and lower quantiles for rapid rehousing versus other service types

Treatment effects can be interpreted as average increases in probability of reentering the homeless system given a household is assigned to the first listed service rather than the second listed service. Results in comparison to rapid rehousing are pictured in Figure 1. On average, those assigned to rapid rehousing see a 28.14 percentage point increase in probability of reentering the homeless system compared to having been assigned to any other service, although the credible interval is wide-ranging. Started in 2009, rapid rehousing is a recent addition to the list of homeless services [Congress, 2009; HUD, 2012]. Problems associated with starting a new service may explain the ineffectiveness of rapid rehousing found in the current analysis. However, an analysis of the quarterly two year reentry percentage of rapid rehousing across the study period shows that the reentry percentage of rapid rehousing has stayed consistent across time, and thus, is not improving. Next, we turn to understanding differences in effects for different households, which is essential to finding an optimal service allocation for individual households.

3.3 Subpopulation Average Treatment Effects

Next, we would like to determine if there is a simple characterization of features that lead to lower probabilities of reentry. In order to do so, we calculate treatment effects for subpopulations of the data separated based on certain covariate values. We focus on the effect of rapid rehousing versus any other service because this treatment effect is large in absolute value and also has high variance, as can be seen in Figure 1. Using BART, we calculate the treatment effects of rapid rehousing compared to any other service for each household. Then, we use regression trees to predict this treatment effect using all features but ignoring service type. Regression trees were chosen because they give interpretable feature importance scores. We use these scores to decide which features have the most effect on the model fit, and subsequently, focus on key features for the subpopulation analysis.

The largest ten treatment effects along with their corresponding features and values are listed in Table 3. [3] These can be interpreted as the average percentage point increase in probability of reentry given a household with the given value of the listed feature is assigned to rapid rehousing rather than another service. For example, we find that households with a monthly income of less than 200 dollars have a 34.54 pp (95% Estimated Credible Interval of 14.62 pp to 48.75 pp) increase in the probability of reentry if they are assigned to rapid rehousing instead of another service. The subpopulation treatment effects are all close to the overall treatment effect with an average of 27.03 and a standard deviation of 5.02. This shows that the differences in treatment effects are not attributable to one or a few clear, easily interpretable factors, but may be due to nonlinear interactions picked up by BART.

4 Optimal Allocation Using Estimated Personalized Treatment Effects

In order to frame the optimal allocation problem, we need two main sets of variables estimated from the data. First are the actual predictions of probability of reentry for households given they are placed in each of the possible inter-

[3] Subpopulation effects for the Prior Residence feature are not included in this list nor in 4 due to a lack of interpretability. Though the dataset contains codes for this variable, it does not contain the meanings of those codes.

Feature	Value	Percentage of Population	Treatment Effect		
			Average	Lower 5%	Upper 95%
Head of Household Received Substance Abuse Services	Yes	8.32	39.56	19.63	49.64
Head of Household Has Substance Abuse Problem	Drug Abuse	8.39	35.03	12.61	49.16
Gender	Male	42.03	34.76	15.37	48.74
Monthly Income	Less than 200	23.99	34.54	14.62	48.75
Wait Before Entry	Less than 30 days	47.90	33.56	16.13	48.67
Head of Household Has Substance Abuse Problem	Both Alcohol and Drug Abuse	6.39	32.82	13.00	49.11
Head of Household Has Substance Abuse Problem	Alcohol Abuse	3.97	32.68	13.36	46.59
Calls Before Entry	One to Five	69.36	31.54	15.47	48.59
Calls Before Entry	None	3.99	30.99	21.06	44.25
Housing Status At Entry	Client Doesn't Know	62.60	30.57	11.16	48.64

Table 3: Ten largest subpopulation effects where an effect of 39.56 corresponds to a 39.56 percentage point increase in probability of reentry. The total population effect is 28.14 (95% Estimated Credible Interval = [11.55,47.96]).

ventions. For this, we use out-of-sample BART predictions. Second are the capacities of the different interventions. In order to estimate these, we aggregate data on a weekly basis, and match the number of entering households into the interventions to the capacities of those interventions in that week. One week is granular enough to give some flexibility to the optimizer, while also not leading to waits that are outside the tolerance of the system. We note here that we solve the problem in a static manner every week, although there could of course be interesting dynamic matching issues at play [Akbarpour *et al.*, 2017; Anshelevich *et al.*, 2013].

4.1 The Optimization Problem

Let x_{ij} be a binary variable representing whether or not household *i* is placed in intervention *j*. Then, the Integer Programming problem is given by

$$\begin{aligned} \min_{x_{ij}} & \sum_i \sum_j p_{ij} x_{ij} \\ \text{subject to} \quad & \sum_j x_{ij} = 1 \quad \forall i \\ & \sum_i x_{ij} \leq C_j \quad \forall j \end{aligned}$$

where p_{ij} is the probability of household *i* reentering if they are placed in intervention *j* and C_j is the capacity of intervention *j*.

We use this IP framework and Gurobi optimization software to find an optimal allocation for households who entered the system during each week. Only households who entered the homeless system between October, 2009 (after initial implementation of the rapid rehousing intervention) through December, 2012 were included in the optimization resulting in 166 separate weeks optimized.

Over the 166 weeks, 3205 out of 5478 households (58.51%) actually reentered the homeless system. Using BART predictions to estimate how many households would reenter in expectation produces an estimate of 3223 households (58.83%), suggesting that the predicted reentry probabilities given by BART are reliable. Using these predicted probabilities to find an optimal allocation, predicted reentries reduce to 2483 households (45.33%). Thus, the optimal allocation framework reduces the predicted number of reentries into the homeless system by 22.95% over this period, a truly substantial potential improvement in outcomes.

4.2 Fairness Considerations

An immediate question is whether the optimal allocation is capturing some inherent inefficiency in the allocation system, and is therefore pareto-improving or at least improving allocations for a substantial portion of the population. This turns out to not be the case. In the optimal allocation, 1731 (31.60%) individual households are allocated to a service in which they have a lower probability of reentry than the service in which they actually participated. Another 1912 (34.90%) are allocated to the same service they were originally assigned. Importantly, 1835 (33.50%) households are allocated to a service in which they have a *higher* probability of reentry. Therefore, the optimal number of expected reentries is achieved by, in effect, hurting more households than it helps in the original allocation. At the same time, the benefits to those who are helped are so strong that they outweigh the costs to those households who are hurt in an additive welfare model. Figure 2 quantifies this by showing the distribution of changes in the probability of reentry between the two allocations.

To further explore differences between those who benefit from the optimal allocation and those who are predicted to do worse, we use a random forest to predict whether a household will have a higher or lower probability of reentry after the optimal allocation using all original features and ignoring service type. We then are able to get measures of variable importance from the random forest model. Figure 3 shows the "mean decrease in accuracy" measure (a standard permutation test for random forest feature importance) for the 30 most influential features. This analysis shows that the two most influential variables for deciding which households will have a lower probability of reentry and which will have a higher probability are length of stay at prior residence and monthly income. Table 4 shows summary statistics for the ten most influential features for the group who improved, the group who was harmed, and the group who did not change.

Perhaps the most striking fact to emerge from this analysis is that the optimal allocation seems to demonstrate a "rich get richer" effect. Households with higher monthly incomes, shorter waits and fewer calls to the hotline before being placed, and those who are less likely to have substance abuse problems are more likely to be placed in interventions

Continuous Feature	Mean (SD) for Group Who Improved	Mean (SD) for Group Who Was Harmed	Mean (SD) for Group Who Did Not Change
Monthly Income	1561.51 (1413.06)	743.53 (2118.75)	967.23 (1217.36)
Calls Before Entry	4.39 (7.82)	6.23 (8.83)	5.59 (7.70)
Wait Before Entry	409.99 (540.15)	452.53 (553.36)	419.98 (556.18)
Age of Head of Household	44.30 (12.50)	40.61 (12.65)	41.28 (12.71)

Categorical Feature	Values	Percentage of Group Who Improved	Percentage of Group Who Was Harmed	Percentage of Group Who Did Not Change
Length of Stay At Prior Residence	One night or less	13.69	2.89	5.33
	Two to six nights	15.89	2.62	5.75
	One week or more but less than one month	47.54	10.30	20.76
	Client doesn't know	22.88	84.20	68.15
Gender	Male	30.16	52.53	37.24
	Female	69.84	47.47	62.76
Housing Status At Entry	Homeless	27.50	5.61	9.94
	At imminent risk of losing housing	30.62	8.94	16.11
	Homeless only under other federal statutes	15.54	1.74	5.28
	Fleeing domestic violence	3.58	0.87	1.88
	Client doesn't know	22.76	82.83	66.79
Head of Household Has Substance Abuse Problem	No	78.57	50.03	60.20
	Alcohol abuse	3.76	7.19	6.49
	Drug abuse	6.99	17.87	11.98
	Both alcohol and drug abuse	5.89	11.34	10.67
	Missing	4.79	13.57	10.67
HUD Chronic Homeless	No	91.33	96.08	96.34
	Yes	8.67	0.98	2.62
	Missing	0.00	2.94	1.05
Head of Household Received Substance Abuse Services	No	90.24	66.59	77.72
	Yes	4.97	19.84	11.61
	Missing	4.79	13.57	10.67

Table 4: Summary statistics for the 10 most influential features for determining which households will benefit from the optimal allocation

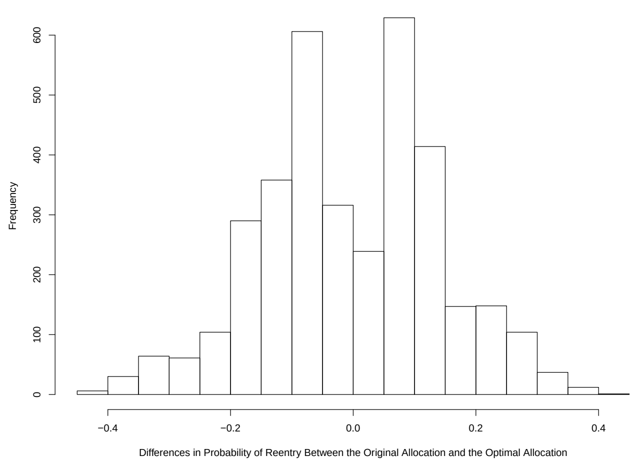

Figure 2: Histogram of improvement in reentry probability under the unconstrained optimized allocation (the 1912 individuals whose probability of reentry was unchanged are not included)

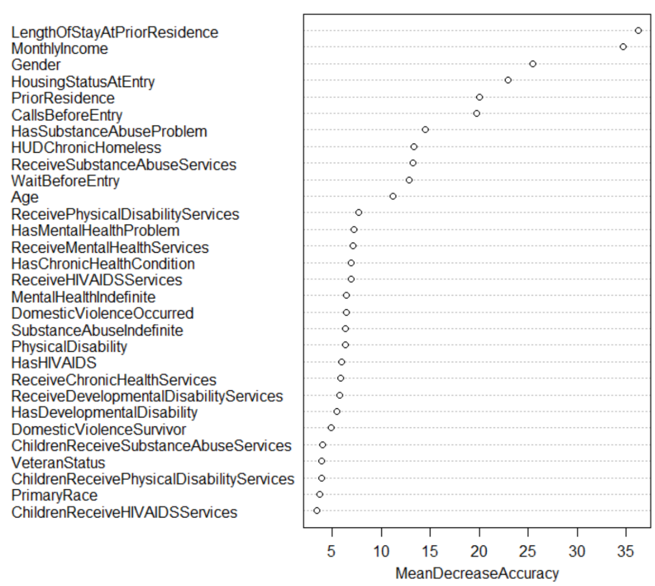

Figure 3: Plot of the mean decrease in accuracy of features for predicting whether the optimal allocation will increase or decrease a household's probability of reentry

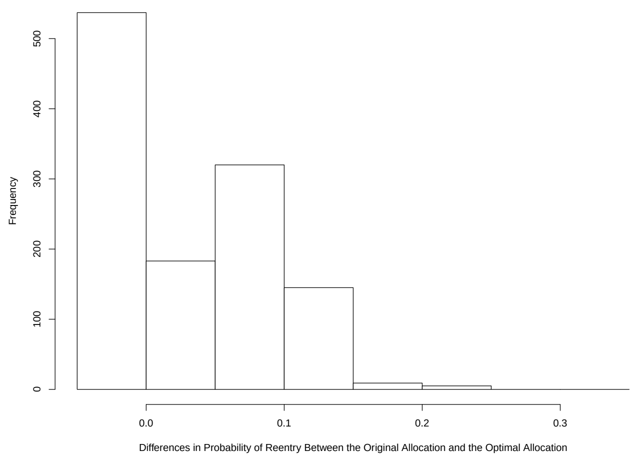

Figure 4: Histogram of improvement in reentry probability under the constrained optimized allocation (the 4279 individuals whose probability of reentry was unchanged are not included)

that are better for them in expectation. This is consistent with the explanation that part of the reason for the inefficiency of the allocation is because of either explicit or implicit policy constraints that require more at-risk households to be given more extensive services.

4.3 Constraining Increased Probability of Reentry

Our results above provide some evidence that the inefficiency of the system with respect to the objective of minimizing reentry may be because of the fact that households that are more at risk have to be prioritized for more intensive services, in keeping with social justice goals. Therefore, this may be a "price of fairness." Assuming the original allocation is fair or close to it in this sense, we could ask whether making (approximate) fairness constraints explicit in the optimization still has the potential to improve outcomes. As an example, we consider what happens if we add a constraint that prevents any household from suffering too high a predicted cost from the change in allocation:

$$\sum_j p_{ij} x_{ij} \leq \sum_j p_{ij} y_{ij} + 0.05 \, \forall i$$

where each y_{ij} is a binary variable representing whether or not household i was originally placed in intervention j. This constraint keeps households from being allocated to a service in which their predicted probability of reentry is more than 5 percentage points higher than that of the service they participated in originally.

When we include this constraint, the solution to the optimization problem yields an allocation with a predicted 3086 households (56.33%) reentering the system within two years. This is obviously higher than the optimized allocation without the constraint, but still a 4.28% decrease compared to the predicted reentry number for the original allocation. Looking again at individual households, 662 households (12.08%) are allocated into a service where they had a lower probability of reentry, 4279 (78.11%) are allocated into the service they were originally assigned to, and 537 (9.80%) are allocated into a service in which they had a higher probability of reentry. Because of the added constraint, no households suffer a penalty of more than 5 pp in the new allocation (see Figure 4).

5 Discussion

This paper tests the feasibility of using data-driven counterfactual approaches to inform policies that guide homeless service provision. Our contributions are along two main dimensions. First, we use Bayesian methods for causal inference to learn about the effects of different interventions. Our findings highlight the ineffectiveness of rapid rehousing in reducing reentry into the homeless system within 2 years. In this dataset, 82.42% of households assigned to rapid rehousing reenter the homeless system within 2 years! A quarterly analysis of reentry over time shows that this number remained stable across the study period, suggesting that the intervention is not improving outcomes. Moreover, although most households respond worse to rapid rehousing, the effects function through complex combinations of household features. Thus, no clear pattern emerges for improving service delivery.

Our findings add to emerging evidence that supports efforts to expand the use of prevention in the homeless system. After accounting for pretreatment differences, homelessness prevention —a time-limited, minimal service —performed as well as longer, more intensive (and more expensive) interventions that provide housing plus case management. Policies that invest in homelessness prevention offer greater potential for improving homeless service delivery than current initiatives to extend expand rapid rehousing. Evidence from this study, as well as others, question current initiatives to expand rapid rehousing. Policies that invest in homelessness prevention offer greater potential for improving homeless service delivery.

Second, we analyze the potential for different allocation mechanisms to improve outcomes using counterfactual estimates of probability of reentry into the system. We estimate that optimal assignments, done on a weekly basis, could reduce the number of reentries into the system significantly. However, a significant number of households are also hurt by the changed allocation (albeit less than the others are helped). Thus, data-driven benefits for the homeless system as a whole do not necessarily improve outcomes for all. We provide evidence to suggest that this may be due to explicit or implicit fairness constraints that prioritize more at-risk households. In an attempt to reduce the harmful effects to part of the population, we impose an additional constraint to prevent households from suffering too much of an increase in the probability of reentry, satisfying a notion of approximate fairness (assuming the original allocation is fair). This still reduces the number of reentries into the system when compared to the actual allocation, but including the constraint reduces the overall benefits from optimizing the assignment of households to interventions.

Our analysis suggest that fairness or justice considerations must be addressed before these types of allocations should be implemented. One potential solution is allowing workers to override certain allocation decisions. This idea has previously been adopted as part of a screening instrument used in New York City [Shinn *et al.*, 2013]. Shinn and colleagues also mention that analysis of the reasons behind these overrides can help to inform future models of this type. The addition of potential override reasons to an allocation model of this type could help to increase fairness, tune future versions of the model, as well as make the transition to an allocation program smoother by allowing homeless service workers to maintain control over allocations.

The findings must be considered in context of study limitations. The observational nature of the data makes it difficult to rule out all potential confounding variables that we were not aware of or to which we did not have access. However, the dataset included all variables measured consistently by the HMIS for which there was enough available data.

Avenues for future work include further analyzing traits of households who were reallocated to services in which they have a higher or lower probability of reentry. It is very important to make sure that allocation systems such as this are not disproportionately harming specific groups. Additionally it would be interesting to look at which new allocations result in lower or higher probabilities of reentry. For example, are more people who end up with higher probabilities of reentry being allocated to emergency shelters rather than homelessness prevention? Answering questions like this will help us learn how to decrease the number of households harmed by this type of service allocation.

References

[Akbarpour *et al.*, 2017] Mohammad Akbarpour, Shengwu Li, and Shayan Oveis Gharan. Thickness and information in dynamic matching markets, 2017. Working paper. Initial version appeared at ACM Conference on Economics and Computation (EC-14).

[Anshelevich *et al.*, 2013] Elliot Anshelevich, Meenal Chhabra, Sanmay Das, and Matthew Gerrior. On the social welfare of mechanisms for repeated batch matching. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 60–66, 2013.

[Azizi *et al.*, 2018] Mohammad Javad Azizi, Phebe Vayanos, Bryan Wilder, Eric Rice, and Milind Tambe. Designing fair, efficient, and interpretable policies for prioritizing homeless youth for housing resources. In *Proceedings of the International Conference on the Integration of Constraint Programming, Artificial Intelligence, and Operations Research*, 2018. To appear.

[Brown *et al.*, 2018] Molly M. Brown, Camilla Cummings, Jennifer Lyons, Andrés Carrión, and Dennis P. Watson. Reliability and validity of the vulnerability index and service prioritization decision assistance tool (vi-spdat) in real-world implementation. *Journal of Social Distress and the Homeless*, 2018. Advance online publication.

[Butaru *et al.*, 2016] Florentin Butaru, Qingqing Chen, Brian J. Clark, Sanmay Das, Andrew W. Lo, and Akhtar R. Siddique. Risk and risk management in the credit card industry. *Journal of Banking and Finance*, 72:218–239, November 2016.

[Chipman *et al.*, 2007] Hugh A. Chipman, Edward I. George, and Robert E. McCulloch. Bayesian ensemble learning. *Advances in Neural Information Processing Systems*, 19:265–272, 2007.

[Chipman *et al.*, 2010] Hugh A Chipman, Edward I George, Robert E McCulloch, et al. Bart: Bayesian additive regression trees. *The Annals of Applied Statistics*, 4(1):266–298, 2010.

[Congress, 2009] United States Congress. Homeless Emergency Assistance and Rapid Transition to Housing (HEARTH) Act, 2009.

[Corbett-Davies *et al.*, 2017] Sam Corbett-Davies, Emma Pierson, Avi Feller, Sharad Goel, and Aziz Huq. Algorithmic decision making and the cost of fairness. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 797–806. ACM, 2017.

[Culhane *et al.*, 2011] Dennis P Culhane, Jung Min Park, and Stephen Metraux. The patterns and costs of services use among homeless families. *Journal of Community Psychology*, 39(7):815–825, 2011.

[Ensign *et al.*, 2017] Danielle Ensign, Sorelle A Friedler, Scott Neville, Carlos Scheidegger, and Suresh Venkatasubramanian. Runaway feedback loops in predictive policing. *arXiv preprint arXiv:1706.09847*, 2017.

[Fowler *et al.*, 2017] Patrick J. Fowler, Kenneth Wright, Katherine E. Marcal, Ellis Ballard, and Peter Hovmand. Capability traps impeding homeless services: A community based system dynamics evaluation. Working paper, 2017.

[Henry *et al.*, 2016] Meghan Henry, Rian Watt, Lily Rosenthal, and Azim Shivji. The 2016 Annual Homelessness Assessment Report (AHAR) to Congress: Part 1 Point-in-Time Estimates of Homelessness, 2016. The Department of Housing and Urban Development.

[Hill, 2011] Jennifer L Hill. Bayesian nonparametric modeling for causal inference. *Journal of Computational and Graphical Statistics*, 20(1):217–240, 2011.

[HUD, 2012] HUD. Homeless emergency assistance and rapid transition to housing: Continuum of care program; interim final rule. *Federal Register*, 77(147):"45422–45467", 2012.

[Johansson *et al.*, 2016] Fredrik Johansson, Uri Shalit, and David Sontag. Learning representations for counterfactual inference. In *International Conference on Machine Learning*, pages 3020–3029, 2016.

[Khadduri *et al.*, 2010] Jill Khadduri, Josh Leopold, Brian Sokol, and Brooke Spellman. *Costs Associated with First-Time Homelessness for Families and Individuals*. Dept. of Housing and Urban Development, 2010.

[Kleinberg *et al.*, 2015] Jon Kleinberg, Jens Ludwig, Sendhil Mullainathan, and Ziad Obermeyer. Prediction policy problems. *The American Economic Review*, 105(5):491–495, 2015.

[Kominers *et al.*, 2017] Scott Duke Kominers, Alexander Teytelboym, and Vincent P Crawford. An invitation to market design. *Oxford Review of Economic Policy*, 33(4):541–571, 2017.

[McCarthy *et al.*, 2017] Sara Marie McCarthy, Phebe Vayanos, and Milind Tambe. Staying ahead of the game: Adaptive robust optimization for dynamic allocation of threat screening resources. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, pages 3770–3776. AAAI Press, 2017.

[O'Neil, 2016] Cathy O'Neil. *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. Crown Publishing Group, New York, NY, USA, 2016.

[Pleiss *et al.*, 2017] Geoff Pleiss, Manish Raghavan, Felix Wu, Jon Kleinberg, and Kilian Q Weinberger. On fairness and calibration. In *Advances in Neural Information Processing Systems*, pages 5684–5693, 2017.

[Shinn *et al.*, 2013] Marybeth Shinn, Andrew L Greer, Jay Bainbridge, Jonathan Kwon, and Sara Zuiderveen. Efficient targeting of homelessness prevention services for families. *American Journal of Public Health*, 103(S2):S324–S330, 2013.

[Solari *et al.*, 2016] Claudia Solari, Azim Shivji, Tanya de Sousa, Rian Watt, and Mark Silverbush. The 2016 Annual Homelessness Assessment Report (AHAR) to Congress: Part 2 Estimates of Homelessness in the United States, 2016. The Department of Housing and Urban Development.

[Tambe *et al.*, 2016] Milind Tambe, Kenneth Pence, and Paul Speer. Optimal allocation of police patrol resources using a continuous-time crime model. In *Decision and Game Theory for Security: 7th International Conference, GameSec 2016, New York, NY, USA, November 2-4, 2016, Proceedings*, volume 9996, page 139. Springer, 2016.

[Weinberger, 2018] David Weinberger. Don't make artificial intelligence artificially stupid in the name of transparency. *Wired*, Jan 2018.