

Washington University in St. Louis

Washington University Open Scholarship

McKelvey School of Engineering Theses & Dissertations

McKelvey School of Engineering

Spring 5-2018

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

Amanda R. Kube

Washington University in St. Louis

Follow this and additional works at: https://openscholarship.wustl.edu/eng_etds



Part of the [Applied Statistics Commons](#), [Engineering Commons](#), and the [Social Work Commons](#)

Recommended Citation

Kube, Amanda R., "Allocating Interventions Based on Counterfactual Predictions: A Case Study on Homelessness Services" (2018). *McKelvey School of Engineering Theses & Dissertations*. 300.
https://openscholarship.wustl.edu/eng_etds/300

This Thesis is brought to you for free and open access by the McKelvey School of Engineering at Washington University Open Scholarship. It has been accepted for inclusion in McKelvey School of Engineering Theses & Dissertations by an authorized administrator of Washington University Open Scholarship. For more information, please contact digital@wumail.wustl.edu.

Washington University in St. Louis
School of Engineering and Applied Science
Department of Electrical and Systems Engineering

Thesis Examination Committee:
Sanmay Das, Co-chair
Zachary Feinstein, Co-chair
Ulugbek Kamilov

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

by

Amanda Rose Kube

A thesis presented to the Graduate School of Arts and Sciences
of Washington University in partial fulfillment of the
requirements for the degree of

Master of Science

May 2018
Saint Louis, Missouri

Acknowledgments

Firstly, I would like to thank my advisor, Dr. Das. With his help and support over the past year, I have learned and grown more than I could have imagined. I would also like to thank Dr. Fowler for his expertise and guidance in the field to which my work is applied. A special thanks goes to my mother. Without her endless love and encouragement, I could never have achieved any of this.

Bibliographic Note: This thesis is based on a paper of the same title by Kube, Das, and Fowler submitted to IJCAI 2018.

Amanda Rose Kube

Washington University in Saint Louis
May 2018

Contents

Acknowledgments	ii
List of Tables	iv
List of Figures	v
Abstract	vi
1 Introduction	1
1.1 Motivation	1
1.2 Previous Work	2
1.3 Problem Setup	5
1.4 Preview of Results	5
1.5 Implications	6
2 Background and Data	7
2.1 Background	7
2.2 Data Collection	8
2.2.1 Data Cleaning and Feature Selection	9
2.3 Data Characteristics	10
3 Analyzing Interventions	11
3.1 Building the Model	12
3.2 Population Treatment Effects	12
3.3 Subpopulation Average Treatment Effects	14
4 Optimal Allocation Using Estimated Personalized Treatment Effects	16
4.1 The Optimization Problem	16
4.2 Fairness Considerations	17
4.3 Constraining Increased Probability of Reentry	18
5 Discussion	20

List of Tables

2.1	Summary of assignment to services across the dataset as well as reentry statistics for each type of service	10
2.2	Summary of features	10
3.1	Ten Largest Subpopulation Effects	15

List of Figures

3.1	Plot of Average Treatment Effects as well as Upper and Lower Quantiles for Homelessness Prevention Versus Other Service Types	13
4.1	Histogram of Improvement in Reentry Probability Under the Unconstrained Optimized Allocation (the 1888 individuals whose probability of reentry was unchanged are not included)	18
4.2	Histogram of Improvement in Reentry Probability Under the Constrained Optimized Allocation (the 3700 individuals whose probability of reentry was unchanged are not included)	19

ABSTRACT OF THE THESIS

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

by

Amanda Rose Kube

Master of Science in Data Analytics and Statistics

Washington University in St. Louis, May 2018

Research Advisor: Professors Sanmay Das and Zachary Feinstein

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 possible interventions can be severely resource limited. This thesis 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, I 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 reentry in the optimal allocation than in the original one. I discuss the efficiency, equity and fairness issues that arise and consider potential implications for policy.

Chapter 1

Introduction

Homelessness represents a long-standing problem with considerable individual and social costs. The homeless system struggles to keep up with demand for services, and there is little empirical support that assesses the accuracy of current decision making in the allocation of limited housing resources [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, Kar et al., 2017, Chan et al., 2017, Yadav et al., 2016b, eg]. In this thesis, I test the feasibility of data-driven approaches to inform policies that guide homeless service delivery. Specifically I ask the question of whether one can use individual predictions of success for families assigned to certain types of homeless services to improve outcomes over the whole population.

1.1 Motivation

The use of techniques from artificial intelligence and machine learning (and more broadly, algorithmic approaches) to make decisions about resource allocation in different societal contexts is both increasingly prevalent and increasingly a matter of concern. Although applications of algorithmic approaches increasingly demonstrate potential improvements in efficiency, examples also raise concerns regarding fairness, accountability, and transparency. A number of examples illustrate the unintended introduction of systematic biases in data-driven allocation of resources in that they perpetuate inequities, such as racial bias 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 the 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 *human interpretability* regarding how decisions are derived from algorithmic approaches to ensure adequate assessment of fairness.

A counter argument to the requirements for human interpretability is that such requirements 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 application 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? Of course, 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 process of policy-making, which could be based on both efficiency and equity considerations.

1.2 Previous Work

There has been much recent interest in the AI and broader computer science community in mechanism design for social good (for example, there was a workshop on this topic at ACM EC 2017). Some topics of interest have included threat screening, poaching, police patrol, and homelessness [McCarthy et al., 2017, Tambe et al., 2016, Kar et al., 2017, Chan et al., 2017, Yadav et al., 2016b]. One useful tool for aiding homeless shelters was developed by Yadav et al. (2016b). This software is named HEALER (Hierarchical Ensembling based Agent which pLans for Effective Reduction in HIV Spread) and is used in a participating homeless shelter to help inform youth about the spread of HIV. HEALER uses Facebook to gather information about the homeless youths’ social networks

and then uses these networks to develop a sequential list of youths who, if invited to participate in the HIV program, have the potential to spread the information to the most other people through the influence of their social networks [Yadav et al., 2016b]. Yadav and colleagues (2016) state that HEALER solves the Dynamic Influence Maximization under Uncertainty (DIME) problem, which is unique in that it selects intervention participants sequentially rather than maximizing in one batch. Previous attempts at solving this problem included PSINET which is a POMDP (Partially Observable Markov Decision Process) based algorithm which runs slowly and runs out of memory when used on large networks. In contrast, HEALER provides a low-cost way of computing social networks by using Facebook and is able to solve the DIME problem more efficiently than PSINET by using HEAL (Hierarchical Ensembling Algorithm for pLanning) another POMDP solver which is quicker and utilizes less memory than PSINET. In another recent study, Chan et al. (2017) develop a decision aid tool for use in allocating homeless youth to services based on NST (a popular risk measure) as well as other features such as age, history of substance abuse, mental, or physical illness, and risk of harm. They show that this tool can help improve outcomes for homeless youth by helping workers place youth in housing programs and identify youth within those programs who may require extra services. Though these studies are tremendously useful, they focus on a small section of the homeless population, homeless youth. To the best of my knowledge, no research team has worked on optimizing the allocation of homeless households (individuals as well as families) to services based on counterfactual probability estimates.

The problem of matching people to resources itself, however, has been subject to much research. There is a long history of mechanism design research on assignment problems including school allocation, organ allocation, refugee matching, etc. For example, Kominers et al. (2017) provide an excellent recent introduction to market design. Multiple studies have developed algorithms for finding allocations that provide the best social utility overall. Anshelevich et al. (2012) compare greedy and maximum-weight repeated batch matching to see which algorithm produces matches with optimal social utility after multiple batches/rounds of matching. Ultimately they find that introducing a threshold under which the compatibility of the match is too low and is thus ignored during the current batch results in a significant increase in social welfare after computing all repeated batch matches. Additionally, work has been done comparing the social utility of stable matches (matches in which no pair would both rather be paired with each other than with the pair to which they were allocated) to matches in which social utility is maximized regardless of stability [Anshelevich et al., 2013b]. However, in these nonstable matchings, there are often situations where certain pairs opt out of the matching in order to pair with each other since this pairing would

benefit them more than their current pairing. To remedy this, Anshelevich et al. (2013b) propose a system in which there is a cost to switching. The switching cost deters switching from the optimal pairing and thus helps to increase the social utility of the match. The problem of allocating homeless households to services can be seen as a repeated batch matching problem with unidirectional utility (i.e. there is no perceived benefit to a service of a particular household being allocated to it). In the case of homelessness, one metric often used to assess the success of homeless services is how many households are in need of more help from the homeless system after having already received previous services. In this application, the social utility of a match can be operationalized as the probability of a household reentering the homeless system given it is allocated to that service.

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 “policy prediction 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. Though observational data contains the outcome for one value of an intervention, potential outcomes from other interventions are also needed in order to make inferences about which interventions are better overall (this is a version of the Rubin Causal Model discussed in Rubin [1974]). There has been significant recent progress in causal modeling from a machine learning perspective. For example, Johansson et al. (2016) discuss counterfactual inference as a domain application problem using machine learning models with similar distributions for both treated and untreated populations. Matching methods (eg. propensity score matching) are also often incorporated in machine learning models to facilitate causal inference. These methods are used to create matches between treatment and control observations based on values of covariates. These matches are used to create treatment and control groups with similar distributions of covariate values in order to reduce potential third variable effects (see Stuart [2010] for a thorough review of matching methods and their uses). For my work, I use Bayesian additive regression trees (BART) [Chipman et al., 2007, 2010] which have the benefit of providing coherent probabilistic estimates of heterogeneous treatment effects [Hill, 2011]. Thus, it allows me to predict individual outcomes under counterfactual allocations.

Despite this past research, this is one of the first studies to consider using machine-learning based estimates of counterfactual outcome probabilities to inform allocation decisions. I present this

work as a proof-of-concept, based on a real administrative dataset, to address the following question: *By optimizing allocations based on counterfactual predictions, how much could we potentially improve outcomes, and what would be the distributional effects of these improvements?*

1.3 Problem Setup

One measure of homeless service provision success is whether a household experiences a repeat episode of homelessness within two years of exiting the homeless system. At any time, there are multiple families entering the system, and multiple types of interventions they could be allocated to (for example, an emergency shelter or a more heavyweight intervention like permanent supportive housing), each subject to capacity constraints. I use data from the homeless system for a metro area to assess the differences in population-wide probability of reentry between pairs of programs. I then estimate capacities of different intervention programs over time and build counterfactual estimates of the probability that a particular household reenters the system if it is placed in each intervention. I 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.

1.4 Preview of Results

Using data from all available households, I find that homelessness prevention leads to a 46.79 percentage point reduction in probability of reentering the system within two years compared to all other services. The effect of prevention is most pronounced when comparing against rapid rehousing which is found to be ineffective on average (in this dataset, 97.20% of those assigned to rapid rehousing reentered the homeless system within two years). Using data on a weekly basis over the course of 166 weeks, I find that the BART model predicts, in expectation, 3146 (62.27%) of the households would reenter the system, and 3147 (62.28%) actually did. The similarity between the out-of-sample predictions and the true reentry statistics shows the BART predictions are well-calibrated. In the optimized assignment, the BART model predicts that only 2479 households (49.07%) would reenter the system. Thus, there may be substantial benefits achievable (by

this reentry 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 reentry, according to the predictions, than those that decrease their probability of reentry. In order to improve the fairness of the allocations, I also formulate and solve a constrained version of the allocation problem, which guarantees that no household increases their probability of reentry by more than 5 percentage points in the new allocation. In this case, 56.59% of households are predicted to reenter.

1.5 Implications

This work is intended as a proof of concept and a case study. I use data to inform 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. I 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.

Chapter 2

Background and Data

2.1 Background

Homelessness represents a complex public health challenge for communities across the United States. Federal guidelines define homelessness as residence in unstable and non-permanent 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 homelessness 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, existing tools fail to discern high and low risk households beyond chance [Shinn et al., 2013]. Providers have limited insight into adapting responses to household characteristics. Moreover, there are no tools that assess the impact of service matches on overall system performance in reducing reentries.¹

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, the 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.

2.2 Data Collection

The data used for this project were collected and managed by a homeless management information system (HMIS). The HMIS recorded all housing services provided to individuals and families seeking federally funded homelessness assistance in a major metropolitan area. Information gathered included household-level details on demographics, risks for housing problems, and documentation of all interactions with the homeless system including dates as well as services.

¹Annual 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.

Homelessness assistance included five major services defined and guided by federal policies. These included homelessness prevention, emergency shelter, rapid rehousing, transitional housing, and permanent supportive housing. Local service providers entered information on service usage in real time through a web-based platform in accordance with federal mandates for collection of universal elements. The platform was hosted and supported by a non-profit organization contracted with the homeless system to provide training, technical assistance, and quality control. Permissions for use of de-identified data were obtained from the local homeless. Records were available from 2007 to 2014.

2.2.1 Data Cleaning and Feature Selection

For this project, I extract data provided by 58 different homeless agencies and link participants across programs by a unique, anonymous identification number. I 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. The primary outcome (the label I am 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 I 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, I assume that they represent a continuation of homeless services. I consider households to have exited from the system when the time between leaving one service and entering another exceeds one day. My 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. I 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 my model. Most of the variables I selected were categorical, and missing values are treated as a separate category in these cases.

Service Type	Number Assigned	Percent Reentered
Emergency Shelter	3313	65.92
Permanent Supportive Housing	256	43.75
Transitional Housing	2150	46.19
Rapid Rehousing	919	97.20
Homelessness Prevention	1190	34.20
Total	7828	58.64

Table 2.1: Summary of assignment to services across the dataset as well as reentry statistics for each type of service

2.3 Data Characteristics

The final dataset includes records on 7828 households. Of these 7828, 4590 (58.64%) reentered the homeless system within two years of exiting. Table 2.1 shows the number assigned to each service type as well as the percentage of those assigned to that service that later reentered within 2 years. Of the 4590 who reentered, 1519 (33.09%) were placed in a subsequent service while 3071 (66.91%) 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 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 2.2 shows a summary and examples of the features included.

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 2.2: Summary of features

Chapter 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 2.1 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, I start by systematically investigating the differential effects of these housing interventions (homelessness prevention, emergency shelter, rapid rehousing, transitional housing, permanent supportive 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. I 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

I compare service types by doing pairwise inference. I select data for each pair and build a BART model based on this data. I 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, I 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. I do this for all pairs of service types as well as for homelessness prevention compared to any other service type.

All pairs that included homelessness prevention did not include zero in the 95% estimated credible interval. However, one pair (prevention versus permanent supportive housing) implied the permanent supportive housing was slightly more effective than prevention (TE = 0.02, 95% Estimated Credible Interval = [0.01,0.03]) while all other pairwise results implied prevention was more effective than each other service. These results are pictured in Figure 3.1. Other pairs that did not include zero in the estimated credible interval were transitional housing versus rapid rehousing

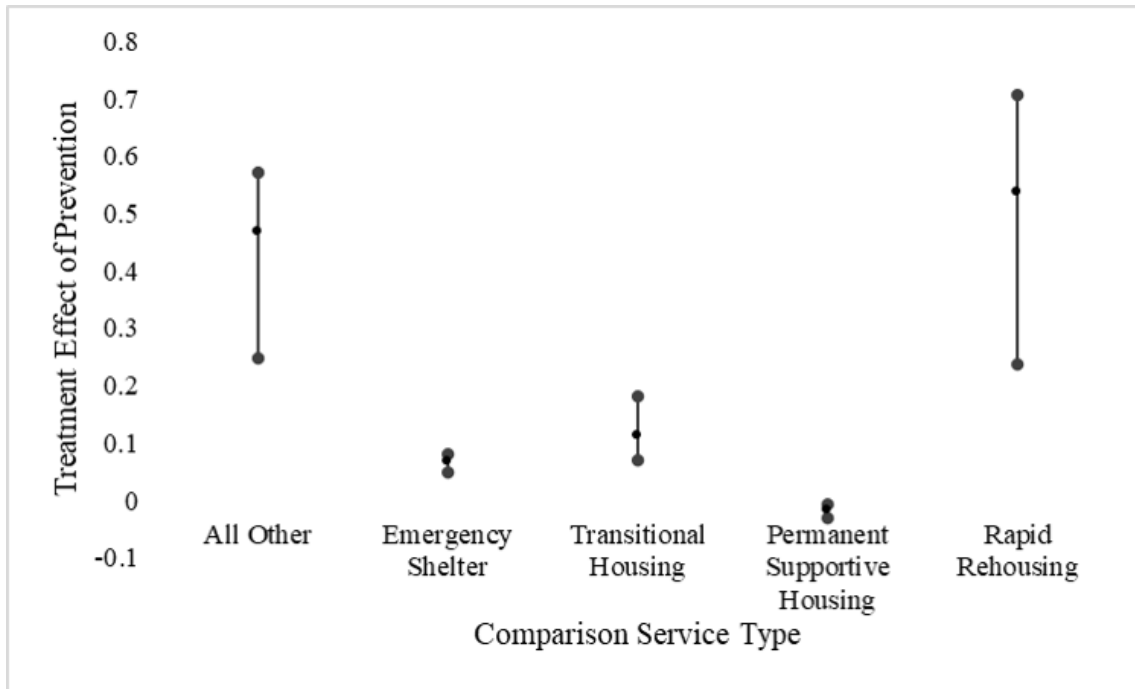


Figure 3.1: Plot of Average Treatment Effects as well as Upper and Lower Quantiles for Homelessness Prevention Versus Other Service Types

(TE = 0.41, 95% Estimated Credible Interval = [0.10,0.60]), emergency shelter versus rapid rehousing (TE = 0.34, 95% Estimated Credible Interval = [0.15,0.54]), and permanent supportive housing versus rapid rehousing (TE = 0.47, 95% Estimated Credible Interval = [0.26,0.63]) implying that all services are more effective than rapid rehousing at reducing the probability of reentry within two years. Started in 2009, rapid rehousing is a recent addition to the list of homeless services [Shinn et al., 2013]. Problems associated with starting a new service may explain the ineffectiveness of rapid rehousing found in the current analysis. Additionally, pairwise results for permanent supportive housing show that it performs better on average than all other services (permanent supportive housing versus emergency shelter TE = 0.08, 95% Estimated Credible Interval = [0.05,0.08]; permanent supportive housing versus transitional housing TE = 0.05, 95% Estimated Credible Interval = [0.03,0.06]) which is unsurprising due to its being the most heavy-weight of all interventions. These treatment effects can be interpreted as average decreases in probability of reentering the homeless system given a household is assigned to the first listed service rather than the second listed service. My results show that assignment to homelessness prevention has the best predicted outcomes of almost any service. On average, those assigned to homelessness prevention

see a huge 46.79 percentage point reduction in probability of reentering the homeless system compared to having been assigned to any other service, although the credible interval is wide-ranging. This effect is largely driven by the relative efficacy of prevention versus rapid rehousing, but there are also clearly significant benefits of prevention compared with assignment to either emergency shelters or transitional housing. Overall, my results on the effectiveness of homelessness prevention, given that it is a relatively lightweight intervention, are somewhat surprising; I discuss this further in Section 5. Next, I 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

I would like to determine if there is a simple characterization of features that lead to lower probabilities of reentry. In order to do so, I calculate treatment effects for subpopulations of the data separated based on certain covariate values. I focus on the effect of homelessness prevention versus any other service because this treatment effect is large and also has high variance, as can be seen in Figure 3.1. Using BART, I calculate the treatment effects of prevention compared to any other service for each household. Then, I 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. I use these feature importance scores to decide which features have the most effect on the model fit and focus on these features for the subpopulation analysis. The largest ten treatment effects along with their corresponding features and values are listed in Table 3.1². These can be interpreted as the average percentage point reduction in probability of reentry given a household with the given value of the listed feature is assigned to prevention rather than another service. For example, I find that households in which the head of household is 60 years of age or older have a 48.05 pp (Credible Interval of 25.78 pp to 58.90 pp) decrease in the probability of reentry if they are assigned to homelessness prevention instead of another service. The subpopulation treatment effects are all close to the overall treatment effect with an average of 45.64 and a standard deviation of 4.52. 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.

²Subpopulation Effects for the Prior Residence feature are not included in this list 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	Treatment Effect		
		Average	Lower 5%	Upper 95%
Housing Status at Entry	Homeless only under other federal statutes	53.28	37.03	59.02
Housing Status at Entry	At imminent risk of losing housing	51.87	34.93	68.83
Length of Stay At Prior Residence	One week or more but less than one month	51.26	35.72	58.94
Housing Status at Entry	Fleeing domestic violence	49.51	15.62	59.08
Head of Household Received Substance Abuse Services	Missing	49.26	36.79	53.60
Head of Household Has Mental Health Problem	Missing	49.26	36.79	53.60
Monthly Income Amount	1400+	48.83	31.43	58.83
Age	60+	48.05	25.78	58.90
Gender	Female	48.03	30.54	58.62
Calls	More than 10	47.83	29.50	57.49

Table 3.1: Ten Largest Subpopulation Effects

Chapter 4

Optimal Allocation Using Estimated Personalized Treatment Effects

In order to frame the optimal allocation problem, I need two main sets of variables estimated from the data. First are the actual predictions of probability of reentry for households given that they are placed in each of the possible interventions. For this, I use out-of-sample BART predictions. Second are the capacities of the different interventions. In order to estimate these, I 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. I note here that I 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., 2013a].

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 Mixed Integer Programming (MIP) problem is given by

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

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 .

I use this MIP 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, 3147 out of 5053 households (62.28%) actually reentered the homeless system. Using BART predictions to estimate how many households would reenter in expectation produces an estimate of 3146 households (62.27%), suggesting that the predicted reentry probabilities given by BART are reliable. Using these predicted probabilities to find an optimal allocation, predicted reentries reduce to 2479 households (49.07%). Thus, the optimal allocation framework reduces the predicted number of reentries into the homeless system by 21.20% 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, 1516 (30.00%) 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 1888 (37.36%) are allocated to the same service that they were originally. Importantly, 1649 (32.63%) households are

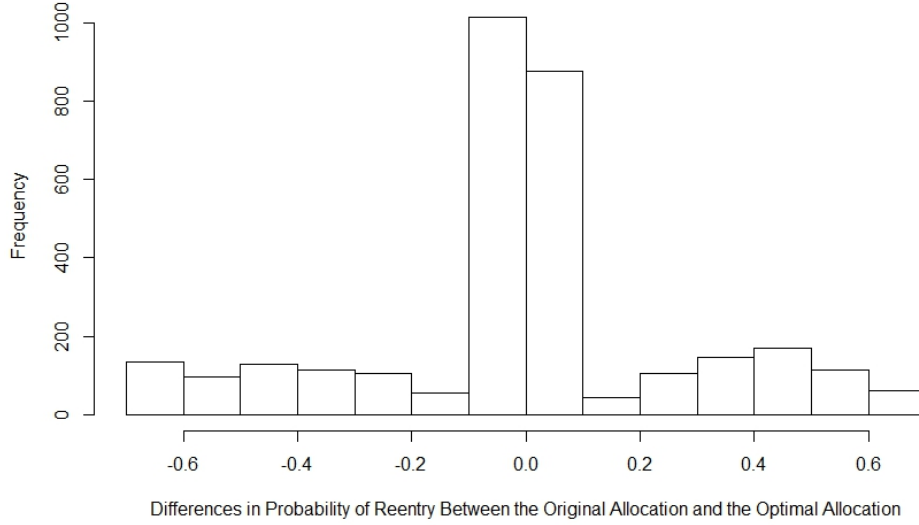


Figure 4.1: Histogram of Improvement in Reentry Probability Under the Unconstrained Optimized Allocation (the 1888 individuals whose probability of reentry was unchanged are not included)

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 completely outweigh the costs to those households who are hurt in an additive welfare model. Figure 4.1 quantifies this by showing the distribution of changes in the probability of reentry between the two allocations.

4.3 Constraining Increased Probability of Reentry

One way to potentially deal with fairness concerns like those raised above is to make them explicit in the optimization. As an example, I consider what happens if I add a constraint that prevents any household from suffering too high a predicted cost from the change in allocation:

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




Figure 4.2: Histogram of Improvement in Reentry Probability Under the Constrained Optimized Allocation (the 3700 individuals whose probability of reentry was unchanged are not included)

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 I include this constraint, the solution to the optimization problem yields an allocation with a predicted 2859 households (56.59%) reentering the system within two years. This is obviously higher than the optimized allocation without the constraint, but still a 9.11% decrease compared to the predicted reentry number for the original allocation. Looking again at individual households, 577 households (11.42%) are allocated into a service where they had a lower probability of reentry, 3700 (73.22%) are allocated into the service they were originally assigned to, and 776 (15.36%) 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 – in fact Figure 4.2 shows that the majority that do worse suffer very small penalties.

Chapter 5

Discussion

This paper tests the feasibility of using data-driven counterfactual approaches to inform policies that guide homeless service provision. Contributions are made along two dimensions. First, I use careful causal analysis to learn about the effects of different interventions. My findings highlight the value of homelessness prevention as well as the ineffectiveness of rapid rehousing in reducing reentry into the homeless system within 2 years. In this dataset, homelessness prevention, a short-term, minimal service, performed as well as, and in many cases better than, longer, more intensive (and more expensive) interventions that provide housing and case management. Although some households respond more to prevention, the benefits function through complex combinations of household features. My findings add to emerging evidence that supports efforts to expand the use of prevention in the homeless system. Additionally, as seen in Figure 3.1, much of the effect of prevention in reducing reentry is driven by its efficacy in comparison to rapid rehousing. In this dataset, 97.20% of households assigned to rapid rehousing reenter the homeless system within 2 years! This is further evidence supporting the notion that despite the recent influx of money to fund rapid rehousing (in 2009 \$1.5 billion was given as part of the American Recovery and Reinvestment Act to fund both Homelessness Prevention and Rapid Rehousing) this intervention is not having the desired effect [Shinn et al., 2013].

Second, I analyze the potential for different allocation mechanisms to improve outcomes, using counterfactual estimates of probability of reentry into the system. I 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. In an attempt to reduce the harmful effects to part of the

population, I impose an additional constraint to prevent households from suffering too much of an increase in the probability of reentry. 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.

This brings up many fairness considerations that must be addressed before these types of allocations could 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 current study also has limitations. The observational nature of the data makes it difficult to say that there are no potential confounding variables that I was not aware of or did not have access to. However, the dataset included all variables measured consistently by the HMIS for which there was enough available data. Another limitation of my dataset is that there may have been more homelessness prevention data that I did not have access to when creating the dataset. The addition of more data on homelessness prevention from other sources has the potential to change the effects of prevention seen in the current analyses.

Avenues for further study include 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 this type of service allocation.

This proof-of-concept using real data from a homeless system provides much needed information to begin dialog and further research into responsible use of algorithmic approaches for social service delivery. Policy discussion must carefully consider both technical and ethical considerations when demographic and other personal data are being used to make decisions about services.

Even if algorithmic allocation decisions like the ones considered here prove untenable, knowing the possible level of benefit will hopefully spur investigations that further unpack the mechanisms underlying heterogeneity in response to housing interventions. Understanding these mechanisms offers great promise to improve policies guiding service delivery.

References

- 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).
- Elliot Anshelevich, Meenal Chhabra, Matthew Gerrior, and Sanmay Das. On the social welfare of mechanisms for repeated batch matching. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 3*, pages 1321–1322. International Foundation for Autonomous Agents and Multiagent Systems, 2012.
- 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, 2013a.
- Elliot Anshelevich, Sanmay Das, and Yonatan Naamad. Anarchy, stability, and utopia: creating better matchings. *Autonomous Agents and Multi-Agent Systems*, 26(1):120–140, 2013b.
- 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.
- 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.
- Hau Chan, Eric Rice, Phebe Vayanos, Milind Tambe, and Matthew Morton. Evidence from the past: Ai decision aids to improve housing systems for homeless youth. 2017.
- Hugh A. Chipman, Edward I. George, and Robert E. McCulloch. Bayesian ensemble learning. *Advances in Neural Information Processing Systems*, 19:265–272, 2007.
- 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.
- 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.
- 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.

- Sanmay Das, Allen Lavoie, and Malik Magdon-Ismael. Manipulation among the arbiters of collective intelligence: How Wikipedia administrators mold public opinion. In *Proceedings of the ACM Conference on Information and Knowledge Management*, pages 1097–1106, 2013.
- Danielle Ensign, Sorelle A Friedler, Scott Neville, Carlos Scheidegger, and Suresh Venkatasubramanian. Runaway feedback loops in predictive policing. *arXiv preprint arXiv:1706.09847*, 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. 2017. Working paper.
- Yuan Gao, Sanmay Das, and Patrick Fowler. Homelessness service provision: A data science perspective. In *AAAI Workshop on AI for Social Good*, 2017.
- 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.
- Jennifer L Hill. Bayesian nonparametric modeling for causal inference. *Journal of Computational and Graphical Statistics*, 20(1):217–240, 2011.
- Fredrik Johansson, Uri Shalit, and David Sontag. Learning representations for counterfactual inference. In *International Conference on Machine Learning*, pages 3020–3029, 2016.
- Debarun Kar, Benjamin Ford, Shahrzad Gholami, Fei Fang, Andrew Plumptre, Milind Tambe, Margaret Driciru, Fred Wanyama, Aggrey Rwetsiba, Mustapha Nsubaga, et al. Cloudy with a chance of poaching: adversary behavior modeling and forecasting with real-world poaching data. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*, pages 159–167. International Foundation for Autonomous Agents and Multiagent Systems, 2017.
- 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.
- Jon Kleinberg, Jens Ludwig, Sendhil Mullainathan, and Ziad Obermeyer. Prediction policy problems. *The American Economic Review*, 105(5):491–495, 2015.
- Scott Duke Kominers, Alexander Teytelboym, and Vincent P Crawford. An invitation to market design. *Oxford Review of Economic Policy*, 33(4):541–571, 2017.
- Mary E Larimer, Daniel K Malone, Michelle D Garner, David C Atkins, Bonnie Burlingham, Heather S Lonczak, Kenneth Tanzer, Joshua Ginzler, Seema L Clifasefi, William G Hobson, et al. Health care and public service use and costs before and after provision of housing for chronically homeless persons with severe alcohol problems. *Journal of the American Medical Association*, 301(13):1349–1357, 2009.

- Sara Marie Mc Carthy, 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.
- 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.
- Donald B Rubin. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of educational Psychology*, 66(5):688, 1974.
- 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.
- 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.
- Elizabeth A Stuart. Matching methods for causal inference: A review and a look forward. *Statistical Science*, 25(1):1–21, 2010.
- 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.
- David Weinberger. Don’t make artificial intelligence artificially stupid in the name of transparency. *Wired*, Jan 2018.
- Amulya Yadav, Hau Chan, Albert Jiang, Eric Rice, Ece Kamar, Barbara Grosz, and Milind Tambe. Pomdps for assisting homeless shelters—computational and deployment challenges. In *International Conference on Autonomous Agents and Multiagent Systems*, pages 67–87. Springer, 2016a.
- Amulya Yadav, Hau Chan, Albert Xin Jiang, Haifeng Xu, Eric Rice, and Milind Tambe. Using social networks to aid homeless shelters: Dynamic influence maximization under uncertainty. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*, pages 740–748. International Foundation for Autonomous Agents and Multiagent Systems, 2016b.

Service Allocation Using Prediction, Kube, M.S. 2018