5-5-2022

Data-Driven Decision-Making: Using Counterfactual Predictions to Allocate Scarce Homeless Services Fairly and Efficiently

Amanda Rose Kube
Washington University in St. Louis

WASHINGTON UNIVERSITY IN ST. LOUIS

Division of Computational and Data Sciences

Dissertation Examination Committee:
Patrick Fowler, Chair
Chien-Ju Ho, Co-Chair
Sanmay Das
Douglas Luke
William Yeoh

Data-Driven Decision-Making: Using Counterfactual Predictions to Allocate Scarce Homeless Services Fairly and Efficiently

by

Amanda Kube, M.S.

A dissertation presented to
The Graduate School
of Washington University in
partial fulfillment of the
requirements for the degree
of Doctor of Philosophy

May 2022
St. Louis, Missouri

Table of Contents

List of Figures

List of Tables

Acknowledgments

Firstly, I would like to thank my advisors, Drs. Sanmay Das and Patrick Fowler for their expertise and guidance in both computation the field to which my work is applied. I would also like to thank all the members of my Dissertation Committee including Dr. Chien-Ju Ho, Dr. Douglas Luke, and Dr. William Yeoh for their time, support, and advise. I extend a special thanks Dr. Yevgeniy Vorobeychik and the rest of the Fairness in AI Working Group at Washington University for extended discussions on portions of this work. Lastly, I am thankful for Sadie and Stella who made working from home bearable.

I am grateful for support from NSF awards 1527037, 1560191, 1910392, 1927422, 1939677, 2127752, and 2127754, the Administration for Children and Families – Children's Bureau Award 90CA1815, and Amazon through an NSF-Amazon FAI award. I appreciate the assistance of Municipal Information Systems Inc. (MISI) for coordinating data access, as well as the households whom it represents, and I am grateful to numerous local homeless providers, who graciously contributed insights throughout my doctoral work.

Amanda Kube

Washington University in St. Louis

May 2022

Dedicated to my mom and sister. Your constant support means the world. And to Matt who believes in me even when I don't believe in myself.

ABSTRACT OF THE DISSERTATION

Data-Driven Decision-Making: Using Counterfactual Predictions to Allocate Scarce Homeless Services Fairly and Efficiently

by

Amanda Kube

Doctor of Philosophy in Computational and Data Sciences

Washington University in St. Louis, 2022

Patrick Fowler, Chair

Chien-Ju Ho, Co-Chair

Artificial intelligence, machine learning, and algorithmic techniques in general, provide two crucial abilities with the potential to improve decision-making in the context of allocation of scarce societal resources. They can flexibly and accurately model treatment response at the individual level, potentially allowing us to better match available resources to individuals. In addition, they can reason simultaneously about the effects of matching sets of scarce resources to populations of individuals. This thesis leverages these abilities to study algorithmic allocation of scarce societal resources in the context of homelessness. In communities throughout the United States, there is constant demand for an array of homeless services intended to address different levels of need. Allocations of housing services must match households to appropriate services that continuously fluctuate in availability, while inefficiencies in allocation could "waste" scarce resources as households will remain in-need and reenter the homeless system, increasing the overall demand for homeless services. This complex allocation problem introduces novel technical and ethical challenges.

First, using administrative data from a regional homeless system, we formulate the problem of "optimal" allocation of resources given data on households with need for homeless services.

The optimization problem aims to allocate available resources such that predicted probabilities of household reentry are minimized. The key element of this work is its use of a counterfactual prediction approach that predicts household probabilities of reentry into homeless services if assigned to each service. To address the inherent fairness considerations present in any context where there are insufficient resources to meet demand, we discuss the efficiency, equity, and fairness issues that arise in our work and consider potential implications for homeless policies. Next, we turn our focus to interpretability and ease of use for homeless caseworkers by using counterfactual predictions to develop decision rules for allocation of resources to homeless households. We use calculated treatment effects of homeless services to develop simple allocations that reduce rate of reentry into the homeless system. We compare these to the original allocation on reentry rate, estimated financial cost, and potential biases in group fairness.

Finally, we examine justice in data-aided decisions in the context of a scarce social resource allocation problem. We empirically elicit decision-maker preferences for whether to prioritize more vulnerable households versus households who would best take advantage of more intensive interventions. We find that, for a subset of about one-third of decision-makers, these preferences change from vulnerability-oriented to outcome-oriented when they are previously exposed to outcome predictions in a different context. In a separate task, when decision-makers assign homeless households to scarce services based on a random presentation of household descriptions with or without algorithmically derived risk predictions, we find that risk predictions reinforce decision-maker preferences. Among those who prioritize the most vulnerable, presenting the risk predictions in addition to the household descriptions leads to a significant increase in allocations to the more vulnerable household, whereas among those who prioritize households who could best take advantage of more intensive interventions, presenting the risk predictions leads to a significant decrease in allocations to the more vulnerable household. These findings emphasize the importance of explicitly aligning data-driven decision aids with the allocation goals – an essential element of social policies that frequently determine whom to serve with scarce resources.

Preface

This thesis contains three individual pieces of co-authored work completed during the course of my doctoral research.

- Chapter 4 represents joint work with Sanmay Das and Patrick Fowler that was accepted for publication by the Journal of Artificial Intelligence Research (JAIR). A preliminary version of that article appeared in the Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI 2019) [54].

- Chapter 5 also represents joint work with Sanmay Das and Patrick Fowler that is currently a working paper.

- Chapter 6 represents joint work with Sanmay Das, Patrick Fowler, and Yevgeniy Vorobeychik that has been submitted to The Twenty-Third ACM Conference on Economics and Computation (EC 2022) and is awaiting review.

I am the primary author on all three papers. I conducted all analyses and the majority of writing. My co-authors provided necessary advise and guidance, reviewed results, and conducted edits to the prose of the final manuscripts.

Chapter 1

Introduction

Homelessness represents a long-standing social problem with considerable individual and collective costs. 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. Annual counts estimate that more than 550,000 people experience homelessness on a single January night across the United States, while approximately 1.5 million people use homeless services at some point during each year, and families with children under 18 years of age comprise more than one-third of homeless households [43].

1.1 Burdens of Homelessness

Experiences of homelessness and associated disruption carry life long implications, as well as significant social costs in lost productivity, compromised health, and social service expenses [26, 36, 49]. Half of homeless families participating in a longitudinal study of homelessness in Alameda County were reported to Child Protective Services at some point during the 37 month follow-up, but only one-fifth of such reports were substantiated [67]. Unsubstantiated reports increased in months leading up to shelter stays and spiked immediately after shelter entry, especially for African American families. First time shelter stays are also associated with increased hospital stays and emergency department visits, particularly prior to shelter entry [75]. Among the top diagnoses for these visits were pregnancy complications, mental health, and drug related diagnoses. Among adults entering shelter for the first time, exits to stable housing are associated with reduced risk of mortality while extended shelter use is associated with increased mortality hazard [62]. Thus, homelessness is a major public health problem in the United States; however, despite efforts at the federal level to reduce rates of homelessness, rates have increased over the past 4 years. [43].

Homeless services coordinated at the community level (as we will discus further in Chapter 2) have limited resources and therefore struggle to keep up with demand for housing assistance. The Family Options Study (FOS), a randomized trial comparing the efficacy of different housing services to care-as-usual, provided evidence of a mismatch between availability and targeting of housing resources and characteristics and preferences of families in shelter [71]. Half of families willing to participate in FOS lost access to at least one potential housing service due to availability, with one service in particular, transitional housing, as the most constrained due to turnover rates or wrong size of available units. In addition, many families found offers unattractive and therefore turned down services. For example, some families had to choose between housing and family integrity in cases where the housing offered did not allow teen or adult males. Therefore, housing services are considered scarce societal resources, and their proper allocation to those in need becomes a complicated social, ethical, and computational problem.

1.2 Assessment and Allocation of Services

There is little evidence to support the efficiency of current decision making in the allocation of limited housing services as the allocation decisions themselves are under-studied [14, 35, 72]. One of the most commonly used tools to aid allocation decisions is the Vulnerability Index-Service Prioritization Decision Assistance Tool (VI-SPDAT) a rapid, interview-style risk assessment tool. This tool has been shown to have poor reliability and validity and was not recommended as the sole instrument for housing prioritization by the assessing researchers [14]. The Next Step Tool (NST) for Homeless Youth, a commonly used risk tool specific to youth seeking homeless services was found to be effective for assisting high risk youth, but greatly augmented by additional predictive features [17]. The information available for decision-making is far from perfect and there is poor understanding of what services work for whom [41, 71]. In the context of scarcity, providers make complex decisions under great uncertainty with small margins of error. Poor decisions that either under- or over-serve homeless households waste scarce resources and miss opportunities for meeting the needs of those not served at all.

Insights drawn from community-based system dynamics suggest improving prevention strategies as a potential leverage point from which to reduce this burden [35]. Predicting those at risk for homelessness or housing instability can help to better target homeless prevention services to the housing insecure [16, 72, 81]. In addition, the ability to predict high-cost/persistent users homeless services can aid in allocation of more supportive services to the chronically homeless. There is evidence that having children separated at shelter entry, having experienced recent parental unemployment, and having previous housing instability are related to repeated or persistent homelessness [38, 39]. In addition, families with parents who experienced homelessness or foster care as children or who experienced homelessness as an adult before shelter are more likely to repeated or persistent users of homeless services [38]. A key research question this thesis addresses is whether advanced computational techniques from the realm of machine learning can be used to accurately generate such predictions and therefore provide additional insight into targeting homeless services to households.

1.3 Promise of AI for Societal Impact

Advances in machine learning and AI techniques have made it possible to apply learning algorithms to generate possible solutions to social problems ranging from raising HIV awareness [84] to wildlife conservation [30]. For example, child protective services incorporate risk scores for abuse and neglect when placing children out of the home [13]; landlords use credit and rental histories to predict evictions when accepting new tenants [1, 61, 66, 70]; judges consider the predicted probability of recidivism when making parole decisions[3, 50].

In the algorithmic decision-making literature on social service provision, the typical approach is to prioritize decisions based on risk scores. For example, Chouldechova and colleagues consider risk assessment in the context of child maltreatment to decide on which calls to a child protection hotline should be investigated further [13, 21]. These represent classic triage situations, and deal with the problem of which cases to select given a limited budget and a risk assessment. Another context of algorithmic decision-making concerns online resource allocation – when a resource becomes available, which of various agents waiting in a queue should be allocated that resource? The most relevant studies along these lines are that of Chan et al. 2017 and Azizi et al. 2018, who consider allocation policies specifically for homeless youth. Chan et al. 2017 focus on possible improvement over the current score-based allocation system with improved human-machine collaboration using AI Decision Aids. Azizi et al. 2018 take this thought further, formulating a dynamic allocation problem between arriving homeless youth and two types of housing resources (rapid rehousing and permanent supportive housing) and consider the issues involved in fair and efficient online allocation of youth to these resources.

The market and mechanism design literature features considerable research on assignment problems, school allocation, organ allocation, refugee matching, etc. (Kominers, Teytelboym, and Crawford 2017 provide an excellent recent introduction to market design). A key focus there has been on the preferences and incentives of the participants, as well as the level of control of the mechanism in allocation decisions. In traditional assignment problems, it is assumed that the principal, the agent who chooses the payoff structure, has full control over all allocation decisions [55]; much of the literature on two-sided matching seeks stable matchings that respect pairwise preferences [69]; work on school choice considers student preferences and school priorities differently [1]; the kidney exchange literature seeks to maximize the number of matches of incompatible pairs [29, 68].

One of the main benefits of the approach taken in this thesis is the possibility of increasing efficiency by exploiting gains from heterogeneity in match quality between households and services. This issue has been explored infrequently in the market design literature, perhaps because of the historical focus on ordinal preferences rather than cardinal utilities [5, 6], which better aligns with systems where agents have considerable control in terms of accepting and rejecting their assignment or matching. However, consideration of cardinal utilities has come up recently in the context of compatible living donor kidney transplantation [58] where one can take advantage of differences in match quality between the organ and the patient and in refugee matching [e.g. 10, 28, 74] where one can optimize over utilities of matchings between refugees and resettlement venues.

Recent approaches to refugee matching, are the closest point of comparison to our work [10, 74]. Trapp et al. 2021 use a combination of machine learning and integer programming to optimize employment outcomes for resettled refugees using historical data. Following Bansak et al. 2018, Trapp et al. 2021 take advantage of the randomness present in current refugee assignment to ensure that selection bias does not affect their modeling. Though we also use machine learning and integer programming on historical data, our setting offers a different challenge given that housing services in our data are not assigned to households at random. Therefore, our observational data, which is routinely collected as part of service provision, is confounded by caseworker decisions. This magnifies the importance of causal modeling. As opposed to the types of problems that Kleinberg et al. 2015 call "prediction policy problems", or for example using machine learning predictions of loan default to manage risk [15], we need useful counterfactual estimates of the effects of different services in order to begin defining the resource allocation problem.

1.4 Fairness Considerations

While algorithms hold promise for improving efficiencies in allocating social resources, a growing body of evidence simultaneously warns against the potential misuses of algorithmic decision-making that could perpetuate racial, social, and economic inequities. Algorithms trained on data that capture disparities inherently reproduce biased predictions [9, 11, 25, 31, 59]. For instance, a healthcare screening system under-enrolled Black patients into needed services compared with Whites by more than half [64]; the algorithm predicted need for care based on healthcare expenditure data that historically exclude Blacks given racial disparities in access to care. Moreover, healthcare access disparities are at risk of widening due to vicious cycles that emerge as data-driven screening systems incorporate biased decisions into future predictions [33]. Thus, a potential exists for automating inequities [34, 63].

The European Union recently passed legislation in response to concerns about ethics, fairness, and privacy. The "General Data Protection Regulation" (GDPR) 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. However, requirements for human interpretability could also diminish the potential of AI to solve societal problems. Algorithmic approaches generate novel solutions that may not correspond to human intuition; requirements for full explainability of these complex processes limits the inherent value of applications to thorny social problems.

David Weinberger presents a compelling example related to autonomous vehicles in a *Wired* op-ed [83]. If self-driving automobiles lowered the number of vehicular fatalities 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 partly depends on whether the remaining crashes disproportionately affect some portion of the population as well as other considerations. Weinberger goes on to argue that while the regulation of AI applied to social problems is critical, it can be achieved through existing processes for resolving policy issues [83]. Governance provides formal and informal methods for establishing rules and norms applied to collective problems, which also include sustainable approaches for mutual accountability. According to Weinberger, the right approach towards AI regulation involves specification of appropriate optimization goals arrived through the social processes of policy-making that consider both efficiency and equity. However, with a few exceptions [e.g. 21] there has not been much empirical investigation probing the tradeoffs that emerge when incorporating fairness considerations into algorithmic decisions, especially in the context of scarcity.

Efforts to promote fair AI systems reveal the complexities involved in data-driven decision-making. One strategy, exemplified by the Moral Machine, aims to train machines in human ethical decision-making. The crowdsourcing platform has elicited more than 40 million decision preferences by presenting humans from nearly 250 countries with a series of unavoidable crash scenarios [7]. Recording whether humans choose to swerve or stay on course provides extensive data with which one could develop decision-making strategies for fully autonomous vehicles. Similar methods attempt to elicit preferences for food donation, organ transplantation, and homeless services recommender systems [37, 56, 80]; however, it remains to be seen whether these systems will achieve fairness or preserve existing undesirable societal biases.

Another strategy attempts to define and assess the fairness of algorithmic decision-making systems. For example, the principle of Anti-Classification ensures protected attributes are not used to build models, Classification Parity requires certain measures of predictive performance to be equal across groups, and Calibration required outcomes to be independent of protected attributes after controlling for estimated risk [23]. However, each of these has its own strengths and weaknesses dependent on the domain of application, and there is no single definition of a fair system that has been developed to date [23]. Despite multiple philosophical and mathematical definitions of fairness, Hannan, Chen, and Joseph found that while variables such as the service being allocated, demographics of the decision-maker, and demographics of the person being allocated the service affect human perceptions of fairness, those perceptions do not correspond directly with any one theoretical definition as presented in philosophical literature.

Research on COMPAS (Correctional Offender Management Profiling for Alternative Sanctions), a proprietary algorithm used by courts in sentencing that predicts defendant risk of recidivism, shows that judges assess recidivism risk inconsistently, and frequently, include prejudices that the algorithm avoids [50]; yet, impossibility results show that algorithms cannot meet all reasonable metrics of fairness at the same time [23, 52, 65]. Fairness tradeoffs emerge depending on which aspect of fairness the decision-making system is designed to achieve, which reflect and potentially perpetuate human biases. Additionally, studies of COMPAS reveal the complexities involved in integrating human and computer decision-making. Evidence shows that non-expert humans perform as well as algorithms in assessing recidivism risk when trained with the essential information and given immediate feedback on their accuracy [31]. However, algorithms identify the most relevant information on which to base accurate decisions more efficiently than humans, especially across many features [59]. Such findings suggest the potential value of integrating algorithmic and human strengths to maximize fairness, but attempts thus far have proved challenging. On one side, training computers on the pattern of errors made by humans and algorithms fails to improve the accuracy of the integrated decisions [73]. On the other side, presenting humans with COMPAS-generated risk predictions fails to improve their accuracy. In contrast, the presence of predictions by themselves has been known to trigger a cognitive bias (i.e., anchoring) that worsens accuracy [76, 79]. Moreover, a longstanding literature exists on framing effects that shows how the presentation of information influences subsequent decisions [77]. Such results highlight the intricacies of integrating algorithmic and human decision-making and raise warnings about unintended consequences that diminish accuracy and further threaten fairness.

1.5 The Present Work

This work investigates the ability to mitigate the burden on the homeless system by using evidence from administrative data to increase the number of households who exit services to stable housing. We leverage the ability of machine learning and AI to model complex systems such as the homeless system to predict returns to homelessness. We make predictions given each possible allocation of the household to a homeless service, giving counterfactual predictions of reentry for each combination of household and service. These counterfactual predictions are the foundation on which the rest of our thesis builds. We investigate potential mathematical and theoretical optimizations of the allocation process as well as the ability to use predictions to further understand the efficacy of homeless services for particular groups. In addition, we ground our work in discussions of the fairness and tradeoffs of these allocation methods. We go on to empirically investigate allocation preferences of human decision-makers and how presenting them with counterfactual predictions may affect their allocation decisions. This work sets the stage for further study of how this information can be used to augment current decision-making in the homeless system and introduces considerations for using predictive methods in other scarce resource settings.

The rest of this thesis proceeds as follows:

- Chapter 2 gives an overview of current national policies on homeless service provision as well as more specific information about homeless services funded by the federal government. It also gives site specific information about the homeless system of St. Louis, Missouri which is the site from which all data for this work was gathered.
- Chapter 3 explains how data was gathered and pre-processed as well as what variables are included in the data. Then, it explains the predictive model used and gives metrics of predictive accuracy.
- Chapter 4 provides the groundwork for the rest of this thesis by showing the feasibility of allocating homeless services based on predictions and beginning to investigate and discuss fairness considerations.
- Chapter 5 builds on Chapter 4 by presenting measures of the efficacy of potential homeless services both across the population of homeless households as well as for individual subpopulations of interest. These measures are used to develop simple,

interpretable decision rules for allocation that are evidence-based yet do not require caseworkers to view and interpret predictive scores. Allocations made using these decision rules are compared to current allocations as well as random allocations in efficiency, financial cost, and fairness.

- Chapter 6 presents results of an empirical study of justice in data-aided decision-making for homeless service allocation. We elicit decision-maker preferences for whether to prioritize more vulnerable households versus households who would best take advantage of more intensive interventions and test how/if prioritizations change when presented with algorithmically derived risk predictions.

- Chapter 7 contains a discussion of the work presented in the previous three chapters as well the implications for future research in this field.

Chapter 2

Background and Current Policy

The homeless system represents the primary community-wide service 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. Similarly to Permanent Supportive Housing, 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 services. 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. During the study period, providers offered services to eligible households on a first-come-first-served basis.

Despite substantial investments, homeless rates remain high in the United States [36]. An enormous challenge is that of matching service types to need. While federal guidelines mandate that local agencies provide services based on risk assessments [78], existing tools fail to discern high and low risk households reliably and accurately [14, 72]. Homeless service providers have limited evidence for adapting responses to observed and unobserved household characteristics [35]. Moreover, there are no tools that assess the impact of service matches on overall system performance in reducing reentries.[1]

2.1 Continuum of Care

The Department of Housing and Urban Development (HUD) is devoted to ending homelessness in the United States. A major change to their programs came in 2009 with The Homeless Emergency Assistance and Rapid Transition to Housing Act (HEARTH Act) [78]. The HEARTH Act of 2009, among other things, consolidated what was once 3 separate homelessness assistance grant programs under the previous McKinney-Vento Homeless Assistance Act (Supportive Housing program, Shelter Plus Care program, and Section 8 Moderate Rehabilitation SRO program) into one grant program which they named the Continuum of Care Program (CoC) [46]. According to HUD, the CoC program is

[1] Annual evaluations of homeless system performance monitor overall rates of return to the homeless system within 24 months, but do not evaluate allocations; future federal funding depends in part on demonstrating trends toward reductions in reentries.

designed to:

- promote community-wide planning and strategic use of resources to address homelessness;
- improve coordination and integration with mainstream resources and other programs targeted to people experiencing homelessness;
- improve data collection and performance measurement;
- and allow each community to tailor its program to the particular strengths and challenges within that community [46].[2]

Federal CoC funds are allocated competitively each year among applicants and can be used to support the following 5 programs: permanent housing (PH), transitional housing (TH), supportive services only (SSO), The Homeless Management Information System (HMIS), and homelessness prevention [46]. PH refers to both permanent supportive housing (PSH) and rapid rehousing (RRH) programs. SSO programs provide services, often in the form of outreach to both sheltered and unsheltered homeless, referrals to housing and other services, and ongoing support but do not provide housing themselves. HMIS refers to a the establishment, operation, and maintenance of technology used to collect information on those receiving services and the service provision itself for a particular CoC. Each CoC is responsible for creating and maintaining its own HMIS. CoCs with designation has a High-Performing Community (HPC) are the only CoCs allowed to house homelessness prevention (HP) programs. CoCs must reapply for this designation each year and up to 10 HPCs are selected. Once designated as an HPC, CoCs may fund HP programs for those at risk of homelessness. HP includes services such as short- and medium- term rental assistance and relocation and stabilization services [46].

[2]Formatting added

2.2 Coordidnated Entry

Under the CoC Program, each CoC must establish a Coordinated Entry (CE) system. HUD defines Coordinated Entry as, "an approach to coordination and management of a crisis response system's resources that allows users to make consistent decisions from available information to efficiently and effectively connect people to interventions that will rapidly end their homelessness [45]." Coordinated Entry standardizes the way households "access, are assessed for and referred to the housing and services that they need for housing stability" [47]. CE ensures that those with the highest need are prioritized and that all CoC resources are use effectively through an assessment process that each client participates in. Prioritization cannot be based on protected factors such as: race, religion, national origin, sex, age, familial status, disability, amount of disability-related service required, sexual orientation, gender identity, or marital status. CE prohibits screening out households unless there are state or local restrictions preventing projects from serving people with particular convictions. HUD refers to the new policy as a move toward a "person-centric" rather than "project-centric" focus [45]. There are four Core Elements to Coordinated Entry: access points for those in crisis, standardized assessment of needs, prioritization, and referral to housing and supportive services.

2.2.1 Access

A key component of CE is ensuring equal access to services [45]. CE programs are required to ensure that the system can be accessed throughout the geographic region that the CoC serves. Outreach programs can serve as an access point by either directly reaching out to unsheltered homeless or by building a network of providers and associated personnel who are likely to come into contact with those seeking services (i.e. social services staff, fire fighters, etc.). Emergency access points must be available at all times, regardless of CoC operating hours. These access points serve those seeking emergency services by providing short-term shelter.

All access points must provide equal access to emergency services and use the same assessment and prioritization criteria. However, CoC are allowed to have separate access points for the following subpopulations: adults without children, adults with children, unaccompanied youth, those "fleeing, or attempting to flee, domestic violence, dating violence, sexual assault, stalking, or other dangerous or life-threatening conditions (including human trafficking)", and households "at imminent risk of literal homelessness" [45]. HUD allows specialized access points for veterans if run by the VA or its partners. Other specialized access points (ie partnering with mental health clinics) are allowed as long as all specialized access points provide access to anyone who presents at their access point regardless of their belonging to the specialized group being served by that access point. In addition, CoCs must market their services to those who are unlikely to reach out on their own and must ensure that all of their marketing and communication is accessible to those with disabilities.

2.2.2 Assessment

CoCs are required to use a standardized assessment process across all access points. This assessment must document a person or household's barriers to stable housing as well as characteristics that can help determine their needs and how they should be prioritized for placement. The same assessment process must be used for everyone with the exception of the five subpopulations listed earlier: adults without children, adults with children, unaccompanied youth, those fleeing domestic violence, and at-risk households. These 5 groups may use differing assessments (ie all unaccompanied youth must us the same assessment but that assessment can be different from that used for adults with children) [45]. Participants may choose not to answer any assessment questions they wish and this should not affect their access to services or referrals. However, in some cases, missing responses may limit their referral options, and staff are required to make participants aware of this if that is the case.

CoCs have freedom in choosing the their assessment tools. HUD does not endorse any specific tool and suggests that assessment tools be selected from the many publicly available tools and can be customized to fit the particular community and subpopulation [45]. Though HUD does not suggest or endorse a particular tool, and no tool has become universal across CoCs they do give the following guidelines:

> Any tool used by a CoC for its coordinated entry process should have, to the greatest extent possible, the following qualities:
>
> - Tested, valid, and appropriate; Reliable (provide consistent results)
> - Comprehensive (provide access to all housing and supportive services within the CoC)
> - Person-centered (focused on resolving the person's needs, instead of filling project vacancies)
> - User-friendly for both the person being assessed and the assessor
> - Strengths-based (focused on the person's barriers to and strengths for obtaining sustainable housing)
> - Housing First–oriented (focused on rapidly housing participants without preconditions)
> - Sensitive to lived experiences (culturally and situationally sensitive, focused on reducing trauma and harm)
> - Transparent in the relationship between the questions being asked and the potential options for housing and supportive services [45].

CoCs should make an effort to only collect as much information as is needed to serve its clients and can do so in phases (ie collect information needed at intake and then gather more later when necessary). The assessment tool or tools used should ensure that the most vulnerable people rise to the top of the priority list within and across all subpopulations [45]. However, HUD itself acknowledges that the predictive value of current assessment tools is untested [45].

2.2.3 Prioritization

Based on information gathered during Assessment, CoCs must prioritize households for services. This prioritization must be consistent throughout the CoC and the criteria used for prioritization must be "specific and definable" and made publicly available [45]. HUD states that prioritization criteria may include any of the following: "significant health or behavioral health challenges or functional impairments that require a significant level of support for the person to maintain permanent housing, high use of crisis or emergency services to meet basic needs, including emergency rooms, jails, and psychiatric facilities, extent to which people, especially youth and children, are unsheltered, vulnerability to illness or death, risk of continued homelessness, vulnerability to victimization, including physical assault, trafficking, or sex work, [and/or] other factors determined by the community and based on severity of needs [45]."

As with assessment tools, HUD does not suggest a particular prioritization scheme. However, they do suggest starting with an assessment tool that produces a prioritization score. Unfortunately, they also state, "no single scoring or other prioritization method has been proven to reliably predict what housing and supportive services project(s) will end homelessness for a specific person [45]." This indicates that both assessment and prioritization are areas of CE where further research into computation methods for predicting outcomes for homeless households may provide improvement to the current system.

2.2.4 Referral

Lastly, Referral to services is based upon the Prioritization step as well as additional information from the Assessment step. During referral, those with the highest levels of need, and thus the highest priority, are offered housing and services first. Staff make suggestions of potential housing and additional services and the participant chooses which, in any, services they would like to accept. Before, participants can enroll in those services, staff must ensure that they meet eligibility criteria for those services [45]. If so, they are enrolled in the services and take off the priority list. If not, they are not enrolled and maintain their place at the top of the priority list. Eligibility criteria are set by the individual projects that provide services and are not considered part of CE. However, CoCs are encouraged to incorporate eligibility (or presumptive eligibility) into the CE process during assessment [45].

2.3 St. Louis, MO

In order to understand how these policies are implemented in the system studies in this thesis, we turn to a local example. As a major Midwestern city with HPC designation, St. Louis provides an interesting case study. St. Louis City and County have separate CoCs but a joint coordinated entry system, allowing them to better serve a community that is transient across city and county lines. We now outline the CE process as it is implemented in St. Louis City and County [22].

2.3.1 Access

Vulnerable households in St. Louis have access to service entry points through walk-ins or by telephone. There are service entry points, termed "front doors," made available by the

CoCs throughout the Greater St. Louis geographic region that are accessible via walk-in. All front doors are required to provide the same assessment described in more detail below. If a front door is unable to serve a client (vulnerable household requesting services), they will contact another front door on the client's behalf. Front door providers are available to all households at risk of or experiencing homelessness and must provide any and all CE assessments required to place a household on the prioritization list. Front doors may choose to serve specific special populations, but all special populations must be able to access at least two front door providers (ie no front door can be the exclusive access point for any special population).

2.3.2 Assessment

During assessment, clients are provided information about the CE system and what services are available to them. They then decide which programs they wish to participate in and what information they wish to provide during assessment. Clients assessed by front doors complete a CE Participation Agreement and receive a Participant Rights Packet. If they decide to proceed with assessment, but are not part of the population served by the front door, staff will call another front door on the clients behalf.

If a client is eligible for any of the services provided by the St. Louis CoCs and wishes to proceed with CE assessment, complete any intake forms or gather required demographics. Then, complete assessments including the VI-SPDAT if there is not already a score on file or if the score on file needs to be updated due to significant life changes or elapsed time. The VI-SPDAT and full SPDAT are the designated assessment tools for the St. Louis CE system and the scores they provide determine priority for services.

The VI-SPDAT and SPDAT have become commonly used tools for the assessment portion of CE and are in use thousands of communities across the US, Canada, and Australia [82]. The VI-SPDAT (Vulnerability Index Service Prioritization Decision Assistance Tool) is an open-access script used to collect information on clients and provide a prioritization score. Information collected includes basic demographic information, family composition, history of homelessness, health, mental health, and risk. Though it is widespread and has been in use, in some form, for over a decade [82], there is evidence that the VI-SPDAT does not reliably measure need. Scores do not accurately predict reentry into the homeless system and have been shown to have poor reliability and validity [14]. Therefore, assessment is an area in need of further study and improvement. Despite its shortcomings, scores from the VI-SPDAT are used in prioritization and referral decisions across the St. Louis City and County CoCs.

2.3.3 Prioritization and Referral

All clients must complete a triage assessment and determine the most appropriate referral. Referrals may be to shelter, prevention, or diversion(RRH, TH, PSH). Emergency shelter beds are assigned on a first-come, first-served basis. Emergency Service Entry Points can make referrals to open shelter spaces at any time. Homelessness prevention is prioritized using minimum eligibility criteria (ie income under defined limits, eviction or disconnect notice). Diversion programs all receive clients through the weekly matching process. Each week, all funded agencies come together for housing matching meetings. A list of available housing openings is gathered and an updated prioritization list is generated. During the meeting, a list of clients, the same length as the number of openings, who are next to receive housing openings is presented. Then, possible matches between the clients and the openings are presented and those attending the meeting may request to swap clients between openings to "ensure the best possible match for each client's needs [22]." Within 48 hours of the weekly housing matching meeting, a housing provider must schedule the eligibility interview with the client which should take place within 7 days of the referral. A client is taken off the prioritization list when the client moves into the permanent housing unit.

2.4 Housing as a Computational Problem

There are several areas, where computational tools can be used to augment this process: triage (deciding between shelter, prevention, and diversion), prioritization (ensuring the neediest households are served first), and referral (finding the best matches during weekly matching meetings). Each of these can be formulated as a prediction problem using data gathered during assessment.

These tools, however, cannot be universal across CoCs as individual policies and eligibility requirements as well as the needs and characteristics of the community vary. Therefore, it is useful to think of the computational problem more generally. Accurately and reliably predicting, and therefore operationalizing, "need" is essential in all three of the areas mentioned above. We argue that predictions of need/vulnerability using information that can be gathered during assessment should be a major focus of computational efforts in the realm of homeless service delivery. A tool that outputs estimates of vulnerability can help with prioritization, but can also help with triage and matching when combined with counterfactual inference. The research presented in this thesis focuses on using predictions of vulnerability to both determine the households most likely to need further services and to aid in determining how to match households with services that lead to stable housing.

Chapter 3

Data and Model Building

3.1 Data

Data for this work come from the homeless management information system (HMIS) of St. Louis, MO 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 include an array of information on demographics, housing risk, 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 deidentified records in accordance with the relevant Institutional Review Board, which made a non-human subjects determination. Regardless, all information was transferred, stored, and analyzed according to best practices in data security. This includes ethics training in research for all research team members.

3.1.1 Data Cleaning and Feature Selection

For this project, we extract data provided by 75 different homeless agencies and link participants across programs by a unique, anonymous identification number. We then aggregate data by household over time 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. We exclude permanent supportive housing for the present study because the service was rarely used as an initial response for first time entries into the homeless system during the study period.

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. We do this using hotline call records to determine whether a household requested additional housing assistance after the initial service. 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 and do not count this as a reentry. 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.

Type	Number	Examples
Binary Features	3	Gender, Spouse Present, HUD Chronic Homeless
Non-Binary Categorical Features	19	Veteran Status, Disabling Condition, Substance Abuse
Continuous Features	13	Age, Monthly Income, Calls to Hotline, Duration of Wait
Total Features	35	

Table 3.1: Summary of features included in BART model

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 modeling. Most of the variables we selected were categorical, and missing values are treated as a separate category in these cases. Table 3.1 shows a summary and examples of the features included. A more complete summary of the dataset is included in Table A.1 in Appendix A.

3.1.2 Data Characteristics

The dataset includes records on 13940 households. The target variable, or label, is a binary indicator of whether households reentered the homeless system, defined as requesting and/or receiving homeless services within 2 years of initial exit. Of the 13940 households, 3987 (28.60%) reentered the homeless system within two years; among reentries, 2066 (51.82%) were placed in a subsequent service, while 1921 (48.18%) called the hotline to request services, but by the end of the two year period had not been placed in another service. Reasons for failing to receive additional services varied; most commonly, services were unavailable and clients were referred to other services (79.13%) or clients did not follow up on referrals (17.67%).

Table 3.2 shows the number of households initially assigned to each homeless service type, as well as the percentage of reentries within 2 years for each service. Models use a single feature vector, which consists of service assignment plus additional covariate data collected at first entry into the system.

Service Type	Number Assigned	Percent Reentered
Emergency Shelter	4431	43.11
Transitional Housing	2449	34.38
Rapid Rehousing	844	40.40
Homelessness Prevention	6216	14.38
Total	13940	28.60

Table 3.2: Summary of service assignment and homeless system reentry within two years by type of service

Eligibility for Prevention

In addition to labeling whether or not a household reentered within 2 years, each household was given a label describing its eligibility for Homelessness Prevention. Households are considered ineligible for Homelessness Prevention if their current housing circumstances do not provide them with adequate shelter. However, this is not always clear from administrative data. In the following work, we use two different definitions of prevention eligibility: generous eligibility in Chapter 4 and conservative eligibility in Chapters 5 and 6. By the conservative definition, anyone who clearly has adequate shelter is considered eligible for prevention. By the generous definition, anyone who clearly does not have adequate shelter is considered ineligible for prevention. Specifically, if a household is at risk or at imminent risk of losing housing or is currently stable or if their current housing status is unknown but their prior residence is either a rental or owned by the head of household with or without a subsidy, they are considered eligible for Homelessness Prevention under the conservative definition of eligibility. If a household's prior residence is a psychiatric facility, detox center, hospital, jail or prison, hotel or motel, staying with a friend or family member, foster care or group home, rental by client, or owned by client or if their prior residence was unknown but they are not considered homeless by federal definitions, they are considered eligible for Homelessness Prevention under the generous definition of eligibility.

Of the 13940 households in the dataset, 5238 (37.58%) are considered conservatively eligible and 10492 (75.27%) are considered generously eligible for Homelessness Prevention. In actuality, the true number of prevention eligible households lies somewhere between these two numbers. Therefore, they serve as upper and lower bounds on the true eligibility constraint.

3.2 Model Building

While Table 3.2 shows apparent differences in the probability of reentry across homeless services, these differences could be due to unobserved variables or selection bias given the nonrandom provision of services. Therefore, it is important to systematically investigate 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. It is assumed that there are differential treatment effects, such that some households would benefit more from specific interventions. Therefore, counterfactual analysis is used to calculate the probability that each individual household will reenter the homeless system within two years given they are placed into each intervention. 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 Additive Regression Trees (BART), an ensemble model that outperforms propensity score and nearest neighbor matching algorithms for causal inference on observational data, especially when the data are complex [44], is a promising method for mitigating this challenge, [19, 20].

Bayesian nonparametric modeling for causal inference has a number of advantages that fit this application [19, 44, 48]. Such models are capable of providing 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 prior studies of homeless service delivery [72]. In the following section, we compare the predictive performance of BART on our dataset to that of several other popular machine learning algorithms: random forests, logistic regression, LASSO, and gradient boosted trees.

3.2.1 Model Comparison

We compared the out-of-sample predictive performance of BART to four commonly used machine learning algorithms using 10-fold cross validation. First, we implemented BART using the default parameters provided by the model creators [19]. Then, we implemented simple logistic regression and LASSO using 10-fold cross validation to choose the value of lambda, the regularization parameter. We also implemented random forests with 500 trees, a minimum node size of 1, and considering 6 variables for each split [12]. Lastly, we implemented gradient boosted trees with 100 trees with a maximum depth of 1, 10 observations per node minimum, and a learning rate of 0.1. As we are comparing to BART using the default parameters, these hyperparameters were chosen because they are commonly used default parameters/implementations for each method. We assess predictive performance using multiple metrics: AUC (Area Under the ROC Curve), Misclassification Error, Precision, Recall, and Calibration which we operationalize as Expected Reentries/True Reentries. The results of this analysis are shown in Table 3.3. We also assess calibration individually for each service type operationalized in the same manner in Table 3.4. BART either outperforms or performs equally to each of the other methods and is well-calibrated across services. As stated previously, it also mitigates the issue of confounder bias that may be present in our observational data and allows for the estimation of household-specific treatment effects. For these reasons we conduct all counterfactual prediction using BART. All model fitting and counterfactual inference that follows is done using the R package `BayesTree` written by the model's creators [19].

Method	AUC	Misclassification Error	Precision	Recall	Calibration
BART	0.7534	0.2506	0.6136	0.3393	0.9999
Logistic Regression	0.7386	0.2576	0.6171	0.2670	0.9996
LASSO	0.7386	0.2583	0.6254	0.2465	0.9995
Random Forests	0.7444	0.2516	0.6110	0.3361	0.8864
Gradient Boosted Trees	0.7462	0.2564	0.6104	0.2920	0.9999

Table 3.3: Comparison of prediction performance of several commonly used methods using multiple metrics

Method	Emergency Shelter	Transitional Housing	Rapid Rehousing	Homelessness Prevention
BART	0.9990	1.0009	0.9961	1.0022
Logistic Regression	1.0001	0.9989	0.9980	0.9999
LASSO	0.9921	1.0059	0.9753	1.0183
Random Forests	0.9423	0.8285	0.9824	0.7860
Gradient Boosted Trees	0.9747	1.0371	0.9382	1.0419

Table 3.4: Comparison of the calibration of each method by service type

Chapter 4

Fair and Efficient Allocation of Scarce Resources Based on Predicted Outcomes: Implications for Homeless Service Delivery

In this chapter, we explore the feasibility of data-driven approaches to inform policies that guide homeless service delivery. Specifically, we ask the question of whether individual predictions of success for certain types of homeless services can be leveraged to reduce the rate of re-entry into the homeless system across the population of households seeking assistance.

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 social services, specifically interventions for homeless households. 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: Local homeless systems coordinate community-wide services that address housing crises. In the United States, services range in intensity from time-limited nonresidential supports to ongoing rental assistance with intensive case management [78]. Each service is capacity constrained, given the constant widespread demand for affordable housing. Thus, homeless providers allocate many households to many services that each vary in availability at any given time. Homeless services aim to stabilize households and reduce future demand for assistance.

National policies currently focus evaluation of homeless service delivery on whether households use additional homeless services within two years of entry into the system; counts are generated from administrative data that record entries and exists across homeless services [46]. However, routine capacity constraints make it challenging to measure success, since those in need may not be able to receive services. Missing information impedes service improvements in most communities across the United States [35].

In this work, we take advantage of unique local administrative records as explained in Chapter 3 that capture community-wide demand and receipt of homeless assistance across time. The data we use link homeless service records with *requests* for assistance through a regional homeless hotline. Operators at the central hotline field all requests for services, as well as make referrals to appropriate and available services. Households call back if they are in need of additional services, and a digital trail captures subsequent requests, regardless of eligibility or delivery of services. This extensive data collection exceeds federal requirements and allows for a comprehensive assessment of homeless services impossible for most communities.

Implications: Our work serves as a proof of concept through a case study. We bring administrative 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. This work contributes to the emerging dialogue on social service delivery based on machine learning predictions. We emphasize the importance of considering fairness, ethics, and the long-term dynamics of systems that use these kinds of predictive models, while at the same time believing that engaging these questions with actual data and estimates can contribute to resolving the lack of evidence guiding current social service delivery.

4.1 Analyzing Services

4.1.1 Counterfactual Estimation of Heterogeneity in Match Quality

Using BART, we built models to produce out-of-sample counterfactual estimates of reentry probabilities if households received each homeless service (i.e., prevention, rapid rehousing, shelter, transitional housing).[3]. For those households that are ineligible for prevention, we did not consider prevention as a potential service. For most of the 10492 households eligible for prevention, homelessness prevention produced the lowest probability of reentering the system within two years (10030 households are predicted to do best in prevention). Three households were predicted to do best in emergency shelter, 323 in transitional housing, and 136 in rapid rehousing. Most of these households were predicted to have the highest probability of reentry if placed in emergency shelter (7126 households) with less predicted to do worst in transitional housing (1990 households), rapid rehousing (1374 households), and prevention (2 households).

[3]These counterfactual estimates for all 13940 households are made available in the following repository: https://github.com/amandakube/Allocating-Homelessness-Interventions—Counterfactual-Predictions

For most of the 3448 households ineligible for prevention, transitional housing produced the lowest probability of reentering the system (2324 households). One-hundred-ninety-two households were predicted to do best in emergency shelter and 932 in rapid rehousing. Again, most of these households were predicted to have the highest probability of reentry if placed in emergency shelter (2119 households) with less predicted to do worst in transitional housing (476 households) and rapid rehousing (853 households).

Relative Ordering of Services	Number of Households	Average Probability of Reentry in ES	Average Probability of Reentry in TH	Average Probability of Reentry in RRH	Average Probability of Reentry in Prev
Prevention Eligible					
Prev, TH, RRH, ES	4363	0.32	0.26	0.29	0.21
Prev, RRH, TH, ES	2415	0.30	0.28	0.26	0.20
Prev, RRH, ES, TH	1601	0.25	0.26	0.22	0.17
Prev, TH, ES, RRH	1105	0.25	0.22	0.27	0.17
Prev, ES, RRH, TH	348	0.22	0.24	0.23	0.16
TH, Prev, RRH, ES	208	0.42	0.33	0.39	0.36
Prev, ES, TH, RRH	198	0.20	0.21	022	0.15
TH, Prev, ES, RRH	71	0.41	0.33	0.43	0.36
RRH, Prev, TH, ES	68	0.53	0.49	0.45	0.47
TH, RRH, Prev, ES	44	0.52	0.45	0.46	0.48
RRH, Prev, ES, TH	35	0.42	0.44	0.38	0.39
RRH, TH, Prev, ES	28	0.57	0.51	0.50	0.52
RRH, ES, Prev, TH	4	0.50	0.56	0.50	0.52
ES, Prev, RRH, TH	2	0.65	0.71	0.67	0.66
ES, TH, RRH, Prev	1	0.84	0.85	0.87	0.88
RRH, TH, ES, Prev	1	0.75	0.72	0.72	0.75
Prevention Ineligible					
TH, RRH, ES	1561	0.46	0.39	0.43	-
TH, ES, RRH	763	0.44	0.39	0.46	-
RRH, TH, ES	558	0.45	0.41	0.39	-
RRH, ES, TH	374	0.38	0.40	0.36	-
ES, RRH, TH	102	0.43	0.47	0.45	-
ES, TH, RRH	90	0.40	0.41	0.43	-

Table 4.1: Number of households having each of the orderings of services from least to greatest probability of reentry (Emergency Shelter = ES, Transitional Housing = TH, Rapid Rehousing = RRH, Homelessness Prevention = Prev)

For each household, we determined which services are predicted to outperform others and developed a relative ordering of service effectiveness. Table 4.1 illustrates this ordering of service effectiveness. Summing across households, almost one-third (31.3%) do best in prevention followed by transitional housing, rapid rehousing, and shelter. Another 17.3% would benefit most in prevention, followed by rapid rehousing, transitional housing, and shelter. For a small proportion of households (11.2%), transitional housing followed by rapid rehousing, and shelter would be best as they are ineligible for prevention. These patterns demonstrate the heterogeneity in treatment effects we hope to leverage to improve the efficiency of allocations.

The probabilities estimated by BART allow us to perform an initial examination of the possibility of optimizing homeless service delivery. If all households were placed in the service in which they have the lowest predicted probability of reentry, we predict 25.00% of households would reenter in expectation. This is a 12.59 percent decrease from the 28.60% who actually reentered. However, it represents an oversimplification of the allocation problem, which in reality is subject to capacity constraints on the number of households that can be served by a particular service at any given time. In the following section, we formulate the optimal allocation problem including these service capacity constraints.

4.2 Optimal Allocation Using Estimated Probabilities

In order to frame the optimal allocation problem, we need two main sets of variables estimated from the data. First are the predictions of probability of reentry for households given they are placed in each of the possible services. For this, we use out-of-sample BART predictions. Second are the capacities of the different services mentioned in the previous section - that is, the number of households that can be accommodated at a given time due to space or monetary limitations. In order to estimate these, we aggregate data on a weekly basis, and set the capacity of a service equal to the number of households who truly entered into the service 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 [2, 4].

4.2.1 The Optimization Problem

We solve an Integer Program for each week of data. Our objective is to minimize the expected number of reentries, ensuring that every household is assigned exactly one service and that no service is assigned more households than its estimated capacity as described above. Let x_{ij} be a binary variable representing whether or not household i is placed in service j. Then, the Integer Programming problem is given by

$$
\begin{aligned}
\min_{x_{ij}} \quad & \sum_i \sum_j p_{ij} x_{ij} \\
\text{subject to} \quad & \sum_j x_{ij} = 1 \quad \forall i \in \mathbb{Z} \\
& \sum_i x_{ij} \leq C_j \quad \forall j \in \mathbb{Z} \\
& x_{ij} \in \{0, 1\}
\end{aligned}
$$

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

We use this IP framework and Gurobi optimization software to find an optimal allocation for households who entered the system during each week.

In the following section, we show this can be re-formulated as a weighted bipartite b-matching problem, known to admit a polynomial time solution.

Reduction to Weighted Bipartite b-matching

Weighted Bipartite b-Matching is the following problem: Given a weighted bipartite graph G with positive, real-valued edge weights, find a subgraph H of G with maximum total weight such that every vertex i in H is incident to at most b_i edges [18].

Given an instance of the current optimization problem, we create an instance of Weighted Bipartite b-Matching as follows. First, create a bipartite graph G such that there are four nodes representing the four services on the right and a single node representing each household on the left. Between each household node i and each service node j, create an edge and give that edge weight $1 - p_{ij}$. For each household node i, let the degree constraint b_i of node i be 1. For each service node j, let the degree constraint b_j of node j be C_j. Then, the allocation of households to services that minimizes expected re-entries while respecting capacity constraints is given by a maximum weighted bipartite b-matching on graph G.

Claim: An optimal weighted bipartite b-matching solution of maximum weight on graph G gives an allocation of households to services that solves the current optimization problem.

Proof: Assume there exists an optimal weighted bipartite b-matching solution of maximum weight on graph G that does not give an allocation of households to services that minimizes expected re-entries while respecting capacity constraints.

We know that each household i is going to be matched to exactly one service j since each household node in G has capacity 1 and not fulfilling that capacity can only reduce the total weight of the solution. Similarly, each service j must be at capacity, since leaving any household unmatched would only result in a solution of smaller total weight. Therefore, if there is an improvement to be made to the optimal allocation of households to services, it must be due to swapping some pair of edges. Now, suppose household h is assigned to service k and h' to service k'. Suppose swapping them so that h were assigned to k' and h' to k would improve the re-entry minimization objective. Then it must be the case that $p_{hk}+p_{h'k'} > p_{hk'}+p_{h'k}$. Which implies $(1-p_{hk})+(1-p_{h'k'}) < (1-p_{hk'})+(1-p_{h'k})$. Therefore, swapping them would increase the total weight of the weighted bipartite b-matching solution.

This contradicts the assumption that our solution to the weighted bipartite b-matching problem was of maximum weight. Therefore, the allocation of households to services that minimizes expected re-entries while respecting capacity constraints must be given by a maximum weighted bipartite b-matching on graph G. ∎

This shows that the solution to our optimization problem can be found in polynomial time. In practice, the optimization is extremely fast in Gurobi (0.03 seconds on average), and time requirements are dominated by running BART, therefore we use the IP formulation.

Optimization Results

Only households who entered the homeless system between October, 2009 (after initial implementation of the rapid rehousing service) through December, 2012 were included in the optimization. This results in tracking 10043 households across 166 separate weeks optimized.[4] For households ineligible for prevention, their predicted probability if placed in prevention is set to 1 so as to eliminate the possibility of the algorithm placing that housing in prevention.

Over the 166 weeks, 2765 out of 10043 households (27.53%) actually reentered the homeless system. Summing BART predictions to estimate how many households would reenter in expectation produces an estimate of 2855 households (28.43%), suggesting that the predicted reentry probabilities given by BART are reliable. Using these predicted probabilities to find

[4]Two simultaneous changes in homeless service delivery precluded additional follow-up. First, new data management software failed to match households in the system before and after 2015. Second, local homeless providers simultaneously shifted services to comply with federal requirements for coordinated entry into homeless services; the result, in effect, unpaired prevention from other homeless services.

an optimal allocation, predicted reentries reduce to 2611 households (26.00%). Thus, the optimal allocation framework reduces the predicted number of reentries into the homeless system by 5.57% over this period. Also recall that the best that could be achieved by assigning each household to its optimal service, without any capacity constraints, was a reentry rate of 25.00%, so our allocation gets us much closer to the best possible reentry rate for this formulation.

4.2.2 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.

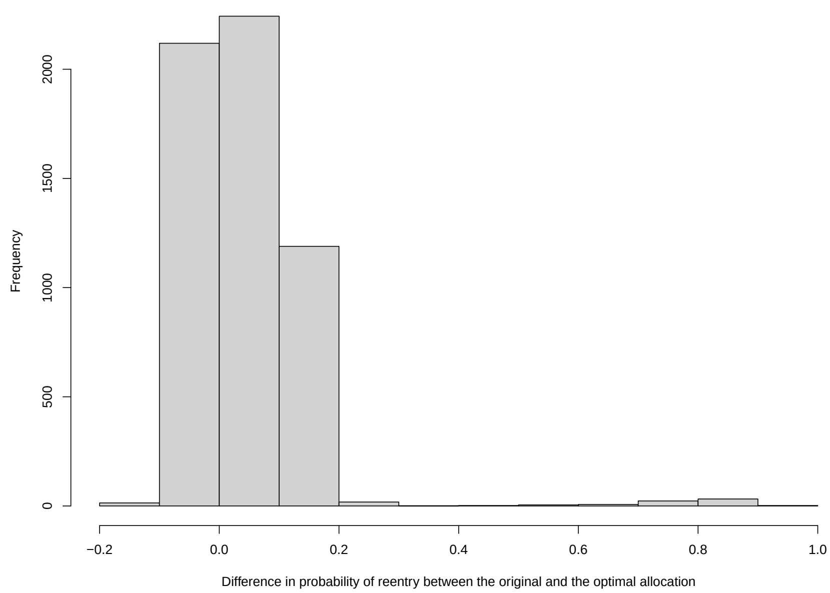

Figure 4.1: Histogram of improvement in reentry probability under the unconstrained optimized allocation (the 4388 households whose probability of reentry was unchanged are not included)

Figure 4.1 shows the distribution of changes in predicted probability of reentry based on our BART model in the optimal service versus predicted probability of reentry for the actual service allocation. In the optimal allocation, 3522 (35.07%) individual households are allocated to a service in which they have a lower probability of reentry than the service in which they actually participated (shown by the area of the histogram to the right of 0). Another 4388 (43.69%) are allocated to the same service they were originally assigned. Importantly, 2133 (21.24%) households are allocated to a service in which they have a *higher* probability of reentry (shown by the area of the histogram to the left of 0). Therefore, a substantial fraction of households are being hurt by the reassignment, even though more are being helped.

Original \ Optimal	Emergency Shelter	Transitional Housing	Rapid Re-housing	Homelessness Prevention
Emergency Shelter	0	0.08	0.06	0.12
Transitional Housing	-0.03	0	0.01	0.06
Rapid Re-housing	0.02	0.09	0	0.10
Homelessness Prevention	-0.03	0.03	-0.00	0

Table 4.2: Average percentage point difference in probability of reentry for households moving between services in the optimal and original allocations. Positive numbers represent decreases in probability of reentry.

Table 4.2 shows the average percentage point difference in probability of reentry for households moving from one service in the original allocation to a different service in the optimal allocation. The mainly positive non-zero off diagonals suggest potential improvements from optimization that range from small (e.g., rapid rehousing to shelter) to larger changes, especially reassignment to transitional housing. Although BART shows homelessness prevention represents the best option for most households, the percentage point gains are relatively modest; those who are moved out of prevention typically have worse outcomes.

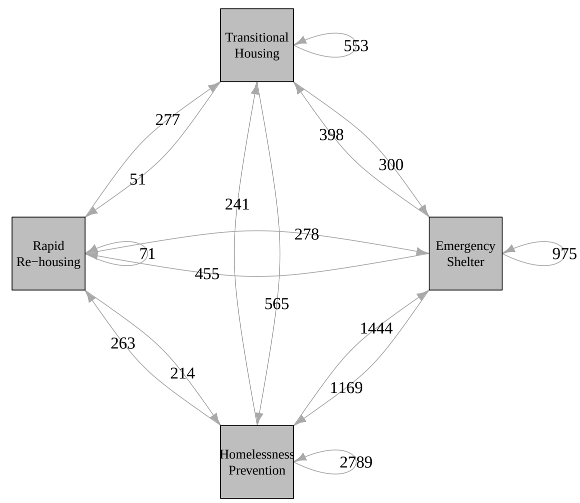

(a) Full network map

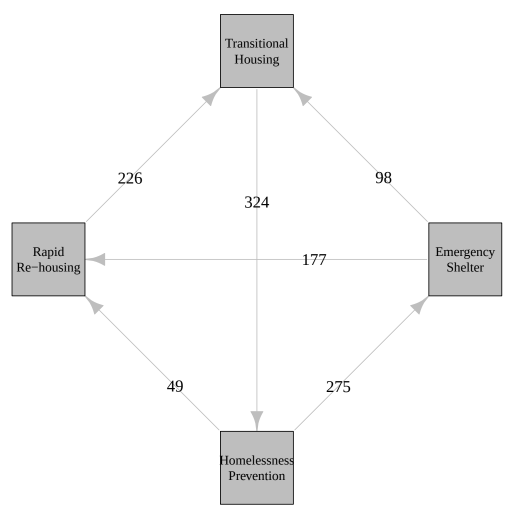

(b) Network map with net flow

Figure 4.2: Network maps of the number of households moving from each service to another in the optimal allocation

Figure 4.2 shows the mechanism of improvement, given capacity constraints. It maps the changes in allocation between the different services in the optimal allocation, as compared with the original. Figure 4.2a shows the number of households who moved from each service to another in the optimal allocation and Figure 4.2b shows the net flows of households moving between services. It is clear that the main mechanisms of improvement are flows where a significant number of households are being placed in transitional housing rather than rapid rehousing and in prevention rather than transitional housing; in order to make room for these, households move from prevention to shelters. This flow indicates a potentially complex mechanism for improving outcomes, since it is not simply a two-way swap between services.

We explore further who benefits in optimization to assess potential inequities. We build random forest models using the default hyperparameter values listed for the classification problem of predicting whether a household has a higher or lower probability of reentry after optimal allocation. We chose random forests due to the ease of producing measures of variable importance from a random forest model. The models have access to the entire original set of features, but ignore service type. The relative importance of each feature for prediction (calculated using the mean decrease in accuracy of features – a standard permutation test used in random forest feature importance) provides insights into the key characteristics that differentiate those who improve or worsen their reentry probability. The out-of-bag error for the random forest model was 0.09 and the AUC was 0.97. Figure 4.3 plots the 30 most influential variables. Some of the most important features are prior residence, housing status at entry, and the number of hotline calls prior to entry.

Perhaps the most striking discovery to emerge from the analysis is that the optimal allocation seems to help those who stand out as being *more* in need. Households benefited most by reallocation disproportionately are homeless upon entry and make frequent calls to the hotline for help; they also are more likely to reside in non-federally funded homeless services (primarily

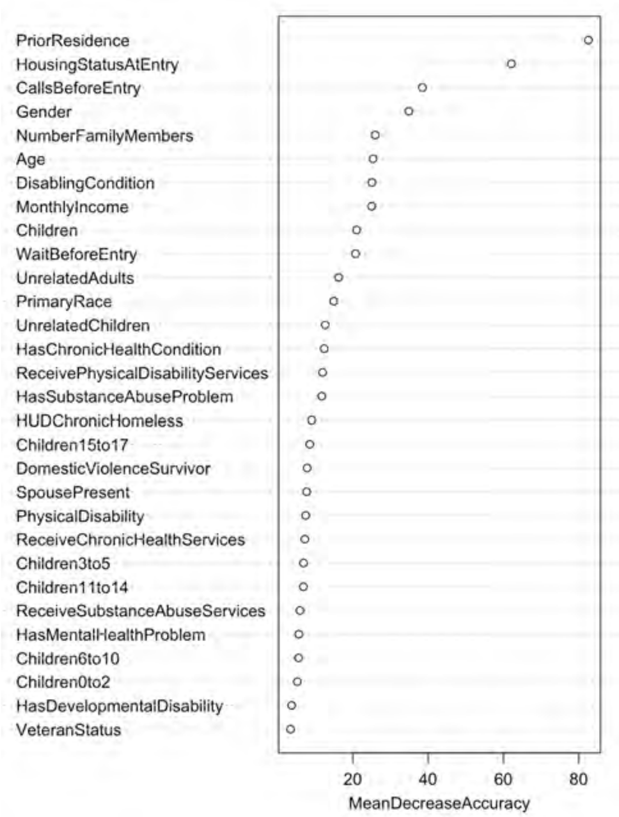

Figure 4.3: Plot of the mean decrease in accuracy of features for predicting whether the optimal allocation will increase or decrease a household's probability of reentry

provided through local religious organizations), substance abuse treatment facilities, or with family. Moreover, reallocation benefits households more likely to report a disability who wait longer for entry into services. Households harmed by optimization, on the other hand, are more likely to be at risk or at imminent risk upon entry, first time hotline callers with briefer waits for services, and in their own or rental units; household heads also are somewhat older and more likely to have children. The no change group also experience stable housing in their own units upon entry. Table 4.3 summarizes comparisons of household characteristics by reallocation outcomes. All differences in continuous variables between the group who improved versus harmed were tested using a Student's t-test and have p-valus at or below 3.21e-14.

Feature	Improved Group (n = 3522) %/Mean (SD)	Harmed Group (n = 2133) %/Mean (SD)	No Change Group (n = 4388) %/Mean (SD)	Total (n = 10043) %/Mean
Prior Residence: Emergency Shelter	42.84	10.77	46.40	12.85
Prior Residence: Transitional housing for homeless persons	41.17	8.33	50.50	5.97
Prior Residence: Substance abuse treatment facility or detox center	81.19	1.63	17.19	16.72
Prior Residence: Staying or living in with family member	59.17	11.67	29.16	11.78
Prior Residence: Rental by client no ongoing housing subsidy	4.07	33.63	62.30	24.45
Prior Residence: Owned by client no ongoing housing subsidy	2.94	43.12	53.94	12.88
Housing Status At Entry: Homeless	79.76	9.66	10.58	10.72
Housing Status At Entry: At imminent risk of losing housing	5.41	36.31	58.28	34.63
Housing Status At Entry: At-risk of homelessness	3.82	38.05	58.12	11.46
Housing Status At Entry: Stably Housed	21.60	26.83	51.57	2.86
Housing Status At Entry: Client doesn't know	58.49	6.20	35.31	40.33
Calls Before Entry	3.65 (5.91)	1.35 (4.15)	2.63 (5.50)	2.72 (5.46)
Gender of Head of Household: Female	31.33	21.68	47.00	68.31
Gender of Head of Household: Male	43.14	20.30	36.57	31.69
Number of Family Members	1.53 (1.10)	2.66 (1.62)	2.01 (1.31)	1.98 (1.38)
Age of Head of Household	38.31 (13.31)	40.98 (12.41)	39.80 (12.41)	39.53 (12.77)
Head of Household Has Disabling Condition: No	33.82	21.36	44.82	82.82
Head of Household Has Disabling Condition: Yes	41.39	19.47	39.14	15.03
Head of Household Has Disabling Condition: Don't Know	38.97	29.11	31.92	2.12
Head of Household Has Disabling Condition: Refused	50.00	0.00	50.00	0.02
Monthly Income	871.76 (1111.02)	2008.84 (2753.28)	1487.54 (1528.84)	1382.31 (1800.82)
Number of Children	0.50 (1.05)	1.41 (1.49)	0.92 (1.25)	0.88 (1.28)
Wait Before Entry	253.89 (452.80)	141.48 (393.75)	249.51 (469.72)	228.10 (450.80)

Table 4.3: Summary statistics for the most influential features for determining which households will benefit from the optimal allocation (due to the large number of prior residence categories, those making up less than 5% of the population were omitted from the table)

Overall, these results suggest an ability to improve upon the allocation rules used by the homeless system. To note, although more than one optimal solutions could exist, we find evidence only for a single solution across runs. Interestingly, the efficiency gains are achieved primarily through "shuffling" households between emergency shelters (which is a uniformly poor service), prevention (which may be appropriate for more vulnerable households than previously believed), and transitional housing (an intense and expensive service with higher efficacy). There is clearly some household-level heterogeneity that could potentially be exploited to achieve gains.

4.2.3 Constraining Increased Probability of Reentry

Another important dimension of fairness raised in algorithmic decision-making pertains to the local costs of redistributing resources. Inefficiencies in the original allocation may be because decision-makers are prioritizing equity by assigning more vulnerable households to more intensive services (whether the measurement of vulnerability corresponds to the actual notion we care about is a separate question)[36]. Of course, this idea may be flawed in that some of these "more vulnerable" households may actually be equally well-served by less intensive services.

One way to potentially deal with fairness concerns like these is to make them explicit in the optimization. As an example, we consider what happens if we add a constraint that prevents any household from suffering too high a predicted cost, in terms of predicted increases in probability of reentry, from the change in allocation. For example:

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

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

To illustrate the results of the allocation when this constraint is added, Figure 4.4 shows the distribution of changes in the the new allocation when δ is set to 5 percentage points. The hard threshold prevents any negative changes of greater than 5 percentage points. When we include this constraint, the solution to the optimization problem yields an allocation with a predicted 2644 households (26.25%) reentering the system within two years. This is just a little bit higher than the optimized allocation without the constraint, but still a 4.36% decrease compared to the predicted reentry number for the original allocation. Looking again at individual households, 2619 (26.08%) are allocated into a service that lowers probability of reentry, 5746 (57.21%) are allocated into the original assignment, and 1678 (16.71%) are allocated into a service that increases probability of reentry. The majority of households who do worse suffer very small penalties.

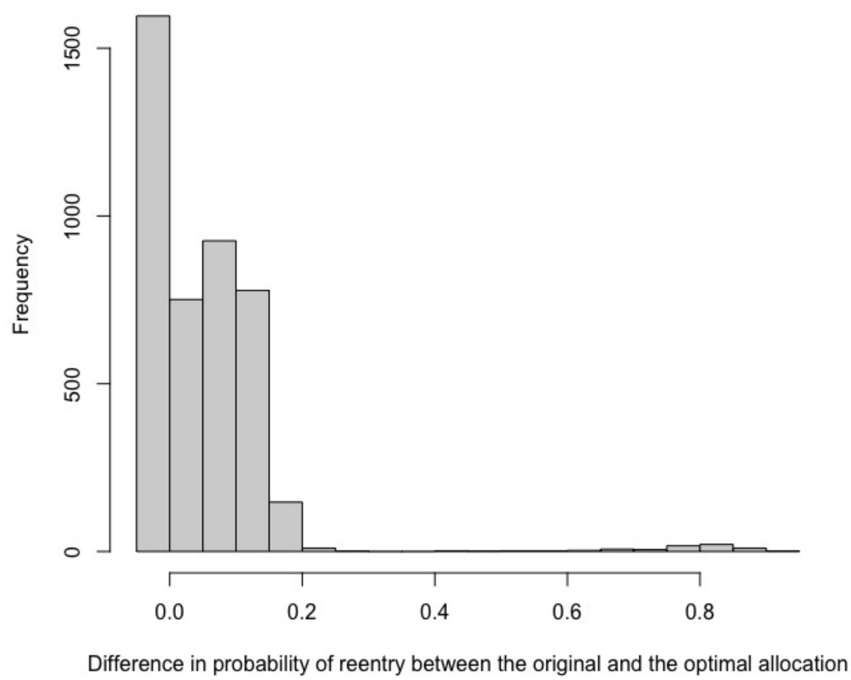

Figure 4.4: Histogram of improvement in reentry probability under the constrained optimized allocation (the 5762 individuals whose probability of reentry was unchanged are not included)

We empirically investigate the influence of imposing more and less restrictive fairness constraints on reentry rates. Figure 4.5 shows the percentage reduction in expected reentries as a function of δ (how much each household's predicted reentry probability is allowed to increase in the optimal allocation). That is, how much predicted cost households are allowed to incur from the change in allocation. An interesting result from this investigation is that, even when the constraint is set to 0.01 and barely allows any household's predicted reentry probability to increase, we achieve a modest reduction in expected reentries. Therefore, we can produce gains in efficiency even in the presence of strict fairness constraints.

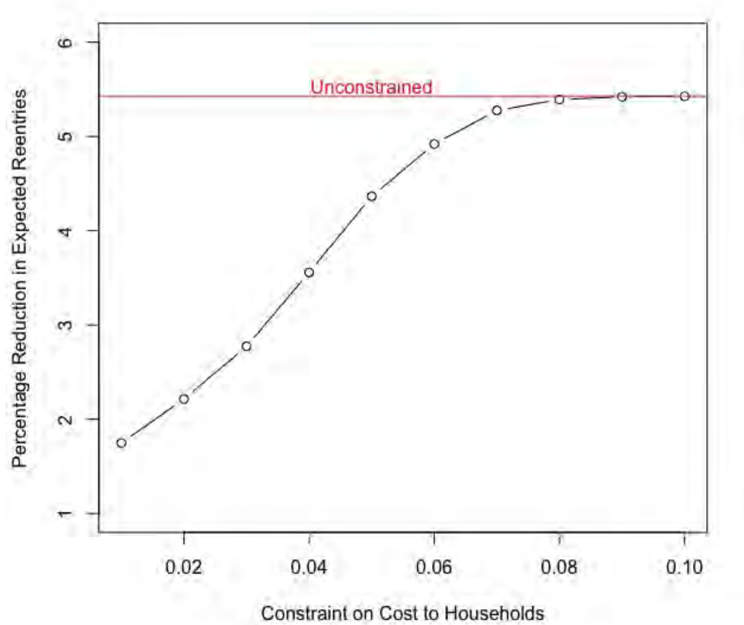

Figure 4.5: Graph showing percent decrease in expected number of reentries as a function of constraint on how much a household's predicted reentry probability is allowed increase in the optimal allocation

4.3 Discussion

Our work tests the feasibility of using data-driven counterfactual approaches to inform policies that guide homeless service provision. We analyze the potential for different allocation mechanisms to improve outcomes using counterfactual estimates of probability of reentry into the system. Our results suggest that optimal weekly assignments reduces system reentries. However, optimization of system-wide service delivery withholds useful services for one-third of households. Although the average harm to households is small in comparison to the benefits for other households, the results emphasize that optimal reallocation of services fails to improve the outcomes of all households in the homeless system. Assuming the original allocation to be fair, models explore the imposition of an approximate fairness constraint that avoid households from being reallocated to services that worsen the probability of reentry into the system compared to the original allocation. Results show smaller but meaningful reductions in reentries into the homeless system using fair data-driven allocations of services.

Our findings demonstrate the critical importance of fairness and justice considerations in the design of algorithmic allocations of homeless services delivery. The assumptions, implications, and potential unintended consequences must be thoroughly analyzed and addressed before implementing data-driven decision-making. One potential solution allows workers to override certain allocation decisions. The idea has previously been adopted as part of a homelessness prevention screening instrument used in New York City [72]. Shinn and colleagues note 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 could help to increase fairness and inform re-calibrations of models. It also makes the transition to an allocation program smoother by allowing homeless service workers to maintain control over allocations.

The results presented here must be considered in the context of limitations of this kind of study. It is difficult to rule out all potential confounds for treatment estimates. Our models leverage all available data from homeless services for predictions, and extensive sensitivity analyses provide some confidence in the results. However, the observational nature of the data constrains modeling for variables we were not aware of or to which we did not have access. If the estimated treatment effects are biased, this would inherently worsen efficiency gains by introducing unreliability.

Another key limitation concerns the potential for unobserved inequities in homeless service delivery. Administrative records only collect information on services provided; models remain vulnerable to service decisions that intentionally (i.e., explicit bias) or unintentionally (i.e., implicit bias) disadvantage specific groups. As illustrated in prior applications, algorithmic decision making risks perpetuating systematic inequalities captured in the data [24, 33, 65]. Surprisingly, initial tests in the present study suggest optimal allocation disproportionately advantages more vulnerable households. The unexpected findings potentially reveal counterproductive assumptions guiding service delivery. Currently, homeless policies prioritize scarce intensive services for more vulnerable households, whereas the data-driven allocation maximizes timely receipt of preventive services for first time entries into the homeless system [78]. These findings are consistent with a growing body of evidence on community-wide benefits of homelessness prevention [35]. Insights from the present study introduce new avenues for future work that informs data-driven homeless service delivery. Further investigation into heterogeneous effects of different homeless services offers opportunities to ask key policy questions of what works for whom. This is especially true for prevention services that unexpectedly show promise at first time entry. In addition, deeper investigation into winners and losers of data-driven allocation needs to test for potential disparities. Fairness considerations must extend to assess whether specific groups are being disproportionately reassigned to certain services (e.g. shelter versus prevention). Answering questions like this would help us learn how to decrease the number of households harmed by this type of service allocation.

In sum, this study demonstrates both the potential of, and the need for caution in, data-driven homeless service delivery. Although machine learning improves efficiency, fairness considerations arise that require careful implementation in practice. Data-driven insights also raise questions regarding policies that underlie service delivery – fitting an algorithms-in-the-loop process [40]. This study opens new lines of inquiry for designing and testing computational approaches that promote social good.

Chapter 5

Data-Driven Homeless Services Designed for Efficiency and Transparency

In the previous chapter, we showed that decreases in expected number of households reentering the homeless system can be achieved by optimizing allocations using counterfactual predictions. In this chapter, we expand on the work done previously to present measures of the efficacy of potential homeless services across the population of homeless households as well as for individual subpopulations of interest. These measures are used to develop simple, interpretable decision rules for allocation that are evidence-based yet do not require caseworkers to view and interpret predictive scores. Allocations made using these decision rules are compared to current allocations as well as random allocations in efficiency, financial cost, and fairness.

5.1 Analysis

5.1.1 Decision Rules

Using counterfactual predictions from BART models, decision rules for more efficient allocation of homeless services can be developed. Three methods are used to investigate these predictions and develop decision rules. Firstly, each household's "best" service can be found by calculating the minimum predicted probability of reentry between the predictions for each service for that household. Dividing the households into groups by their "best" services and calculating demographics for each group can glean information about characteristics of households that do well in each service, helping to determine what household characteristics should lead to an allocation to a particular service. Mean Decrease in Accuracy (MDA) can be used to determine which characteristics should be the focus of this investigation. The MDA measures how much the predictive accuracy of the model would decrease if that variable was left out and thus how "important" that variables is in distinguishing what intervention would give a household its lowest predicted probability of reentry. Secondly, average treatment effects (ATEs) for each pair of services can be calculated as follows: For each pair of potential service assignments, the difference in probabilities of reentry are calculated for each household and the mean and 2.5% and 97.5% quantiles (95% estimated credible interval) of these differences is recorded to estimate pairwise average treatment effects and variation. These ATEs can help determine which service outperform others. Lastly, conditional average treatment effects (CATEs) can be calculated by looking at ATEs for subpopulations of the data to determine which services outperform others for specific subgroups in the dataset. Using these methods, can lead to the development of decision rules for placement of households into homeless services.

5.1.2 Allocation Methods

Decision rules developed using this methodology can be implemented on the administrative dataset to determine the expected number of households that would reenter within two years given the new allocation as well as the estimated monetary cost of the new allocation. This allocation can be compared to the original allocation or a purely random allocation in both cost and number of reentries. As the ability to allocate services differently is constrained by the funding of individual services which impacts their availability, an integer linear program can be developed to optimally assign households to services while keeping to eligibility constraints and staying within a particular budget. This linear program is a simple modification of the Optimized Allocation proposed in [54]. The addition of this optimization as a comparison tests the feasibility of implementing new decision rules and gives insight into the reallocation of resources necessary to do so.

Using estimated costs explained in the next section, the integer linear program is as follows:

> Let x_{ij} be a binary variable representing whether or not household i is placed in intervention j. Then, the Integer Program is the following

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

where p_{ij} is the probability of household i reentering if they are placed in intervention j, d_j is the estimated, per household cost of intervention j, and B is the total budget for all households across all weeks.

This optimization minimizes the total expected number of reentries while allocating households to services such that each household can be assigned to only one service and the overall cost of allocations stays beneath a certain budget B. The IP framework along with optimization software were used to find an allocation within 10% of optimal for households who entered the system during each week, using the estimated cost of the original allocation as the total budget.

5.1.3 Cost Estimation

A literature search for estimated costs of homelessness provided approximate monthly costs of each service [49]. These monthly costs were adjusted for inflation to the 2022 USD and then used together with estimates of service length from the data to produce estimates of the per household total cost of each service as well as an estimate of the total cost of service provision for all households in the dataset shown in Table 5.1. The total cost spent on each service was calculated by multiplying the per household monthly cost of each service by the average months spent in the service and the number of households assigned to that service (multiplying across columns in Table 5.1). The total cost across services is the sum of the total cost for each service. These amounts were validated though comparison to reported HUD funding through the Emergency Solutions grant.

Service Type	Monthly Cost per Household	Average Months in Service	Number of Households	Total Cost
Emergency Shelter	$1,400	0.45	2997	$1,875,522.60
Transitional Housing	$2,150	0.61	1469	$1,926,593.50
Rapid Rehousing	$750	0.99	840	$623,700.00
Homelessness Prevention	$110	1.23	4737	$642,218.80
Total				$5,068,035

Table 5.1: Summary of service assignment and homeless system reentry within two years by type of service

5.1.4 Fairness

In order to understand fairness tradeoffs associated with each of these potential service assignments compared to the original assignment, two metrics for group fairness were calculated: Difference in Gain and Difference in Shortfall using definitions from Mashiat et al. (2022) [60]. Difference in Gain compares realized utility to minimum possible utility and Difference in Shortfall compares realized utility to maximum possible utility. For both metrics, a value of 0 indicates no bias toward either group. Therefore, the larger the number in absolute value the more bias in the allocation for that group.

5.2 Results

5.2.1 Pairwise Inference Using ATEs

Table 5.2 shows the (ATEs) across households for each test, where a positive treatment effect indicates the second listed service outperforming the first listed service. These treatment effects are calculated separately for those eligible for homelessness prevention and ineligible for homelessness prevention. Pairwise differences show that most pairs for which there seem to be meaningful treatment effect differences are those that included homelessness prevention.

Those assigned to homelessness prevention see a 7.80 percentage point decrease in probability of reentering the homeless system compared to having been assigned to emergency shelter, a 5.54 percentage point decrease compared to having been assigned to transitional housing, and a 6.04 percentage point decrease compare to having been assigned to rapid rehousing. For prevention ineligible households, the story is less clear, indicating that a single service may not be best for all ineligible households. Based on this analysis, it is clear that prevention is the best service on average for prevention eligible households. This leads to the first decision rule: **DR1: All prevention eligible households should be allocated to prevention**.

Comparison	Average Treatment Effect (pp)	95% Credible Interval
Prevention Eligible		
Prevention vs Emergency Shelter	-7.80	[-13.76, -3.34]
Prevention vs Transitional Housing	-5.54	[-11.57, -0.01]
Prevention vs Rapid Rehousing	-6.04	[-10.85, -1.38]
Emergency Shelter vs Transitional Housing	2.26	[-3.22, 9.11]
Emergency Shelter vs Rapid Rehousing	1.76	[-2.29, 6.52]
Transitional Housing vs Rapid Rehousing	-0.50	[-7.22, 5.31]
Prevention Ineligible		
Emergency Shelter vs Transitional Housing	4.65	[-3.91, 11.83]
Emergency Shelter vs Rapid Rehousing	2.27	[-4.30, 9.31]
Transitional Housing vs Rapid Rehousing	-2.38	[-11.95, 7.07]

Table 5.2: Table showing the average treatment effects for households for each comparison using BART estimates along with the associated credible intervals

5.2.2 Assessing Minimum Probability of Reentry for Prevention Ineligible Households

Each household's "best" possible allocation was determined to be the allocation that provides that household with the lowest predicted probability of reentry within two years. As **DR1** applies to all prevention eligible households, focus turns to prevention ineligible households. Out of the 5101 homelessness prevention ineligible households, 3349 (65.65%) were predicted to do best in transitional housing, 1537 (30.13%) in rapid rehousing, and 215 (4.21%) in emergency shelter.

Feature	**Transitional Housing** n = 3349 (% / M)	**Rapid Rehousing** n = 1537 (% / M)	**All Prevention Ineligible** n = 4942 (% / M)	**Prevention Eligible** n = 5101 (% / M)	**p-value**
Monthly Income	642.92(886.20)	953.91(976.33)	766.70(976.55)	2017.72(2193.52)	$< 2.2e\text{-}16$
Age	40.52(12.89)	33.04(12.24)	38.11(13.07)	40.99(12.28)	$< 2.2e\text{-}16$
Number of Children	0.13(0.54)	1.26(1.40)	0.54(1.08)	1.23(1.38)	$< 2.2e\text{-}16$
Number of Household Members	1.14(0.56)	2.29(1.43)	1.56(1.11)	2.41(1.49)	$< 2.2e\text{-}16$
Number of Calls Before Entry	3.94(6.04)	4.29(6.95)	4.28(6.83)	1.11(2.72)	0.09
Wait Before Entry (in days)	285.13(490.01)	241.79(447.79)	276.25(480.22)	178.40(412.47)	0.002
Female Head of Household	36.97	90.24	54.21	73.74	$< 2.2e\text{-}16$
African American Head of Household	79.67	75.01	78.18	89.82	0.0004
Disabled Head of Household	22.03	17.05	20.19	9.71	3.159e-05
Head of Household Age 18 to 24	14.30	27.59	18.07	7.08	$< 2.2e\text{-}16$

Table 5.3: Comparison of the characteristics of homelessness prevention ineligible households that are predicted to do best in rapid rehousing and transitional housing

Table 5.3 compares the homelessness prevention ineligible households predicted to do best in transitional housing to those predicted to do best in rapid rehousing across eligibility definitions. The characteristics presented in the table were chosen based on theoretical relevance as well as the mean decrease in accuracy (MDA) associated with that variable in building a model that distinguishes those who have the lowest probability of reentry if placed in transitional housing from those who have the lowest probability of reentry if placed in rapid rehousing. The final column in the table gives the p-value for a Student's t-test comparing values of that feature for those predicted to do best in transitional housing to those predicted to do best on rapid rehousing. A similar table containing this information for all features in the dataset is provided in Appendix B.

The trends in the table suggest that households with lower income, increased wait time, and disabilities should be allocated to transitional housing whereas youth (heads of household between the age of 18 and 24), families with children, and female headed-households should be allocated to rapid rehousing. As many families with children are also female-headed, we also compared female-headed families and single females. 72.92% of female-headed families were predicted to do best in rapid rehousing. In comparison, over 71.29% of single females were predicted to do best in transitional housing. Therefore, it seems the trend of females doing well in rapid rehousing is driven by females often heading families.

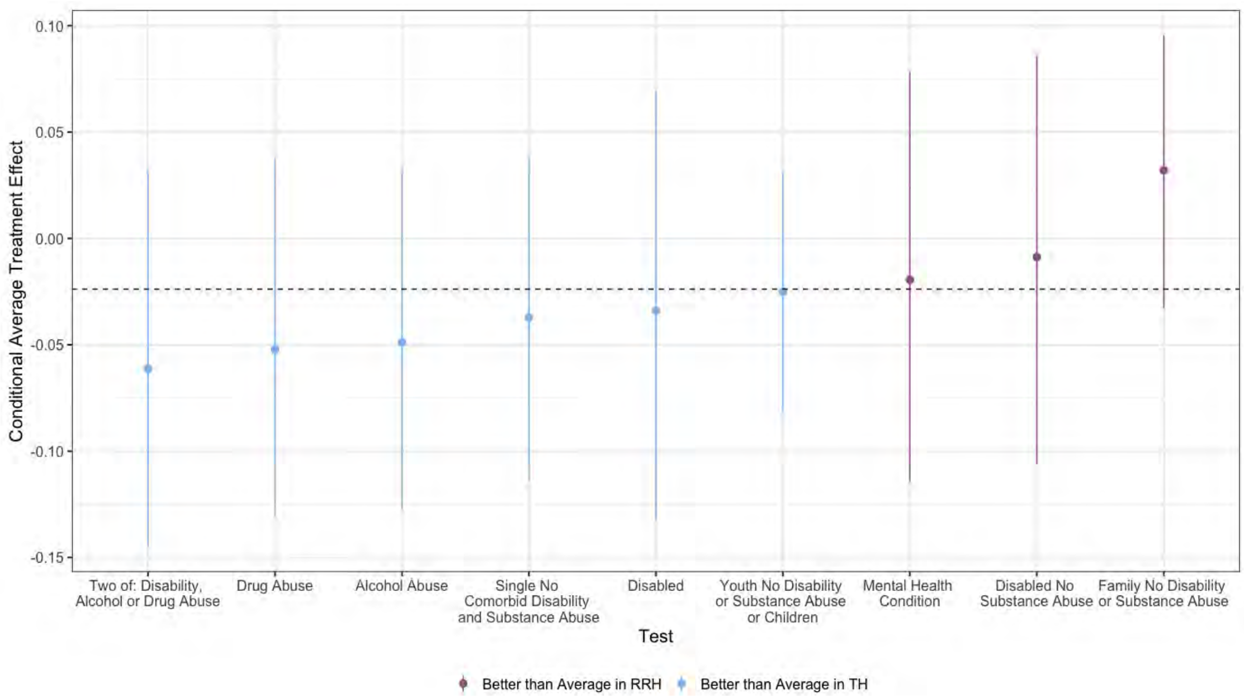

Figure 5.1: Plot showing CATEs of transitional housing vs rapid rehousing for subpopulations of theoretical interest. The dotted line represents the ATE of transitional housing vs rapid rehousing for the entire population. CATEs below 0 indicate transitional housing outperforming rapid rehousing. CATEs above 0 indicate rapid rehousing outperforming TH. CATEs above the dotted line indicate the population performing better than average in rapid rehousing compared to TH. CATEs below the dotted line indicate the population performing better than average in transitional housing

To further explore these differences, Figure 5.1 depicts the CATEs for several variables theorized to be associated with either better outcomes if provided rapid rehousing or if provided TH. Purple lines indicate better than average predicted outcomes in rapid rehousing compared to transitional housing. Values above 0 indicate rapid rehousing outperforming transitional housing for the subpopulation despite the ATE for transitional housing vs rapid rehousing indicating transitional housing outperforming rapid rehousing in the population as a whole. This analysis provides additional evidence that those with disabilities or substance abuse problems are predicted to perform best in transitional housing while families without disabilities or substance abuse problems perform best in rapid rehousing.

Thus, an optimal decision rule for prevention ineligible households would be prioritize those with substance abuse or disabilities for transitional housing and families for rapid rehousing. This leads to **DR2: Those with comorbid disabilities or substance abuse problems should be prioritized for transitional housing. DR3: Families without such comorbidities should be prioritized for rapid rehousing.** Figure 5.2 shows how allocations can be made using decision rules 1 through 3. The natural next question is where to place those that are left unassigned after all decision rules are implemented. These are single adults who are ineligible for prevention who do not have comorbid disabilities or substance abuse problems. Figure 5.1 shows transitional housing outperforms rapid rehousing for this group. Therefore, for the remaining households, we propose a lottery to determine which households receive spaces in transitional housing and which receive spaces in rapid rehousing. To determine the number of spaces used in the lottery, we used the budget determined in Subsection 5.1.3 to calculate the number of spaces in each required to meet but not exceed this budget. In the next section, this allocation is compared to the original allocation by the homeless system, an optimized allocation without constraints, and the optimized allocation constrained to keep costs of allocations under budget as described in Subsection 5.1.2.

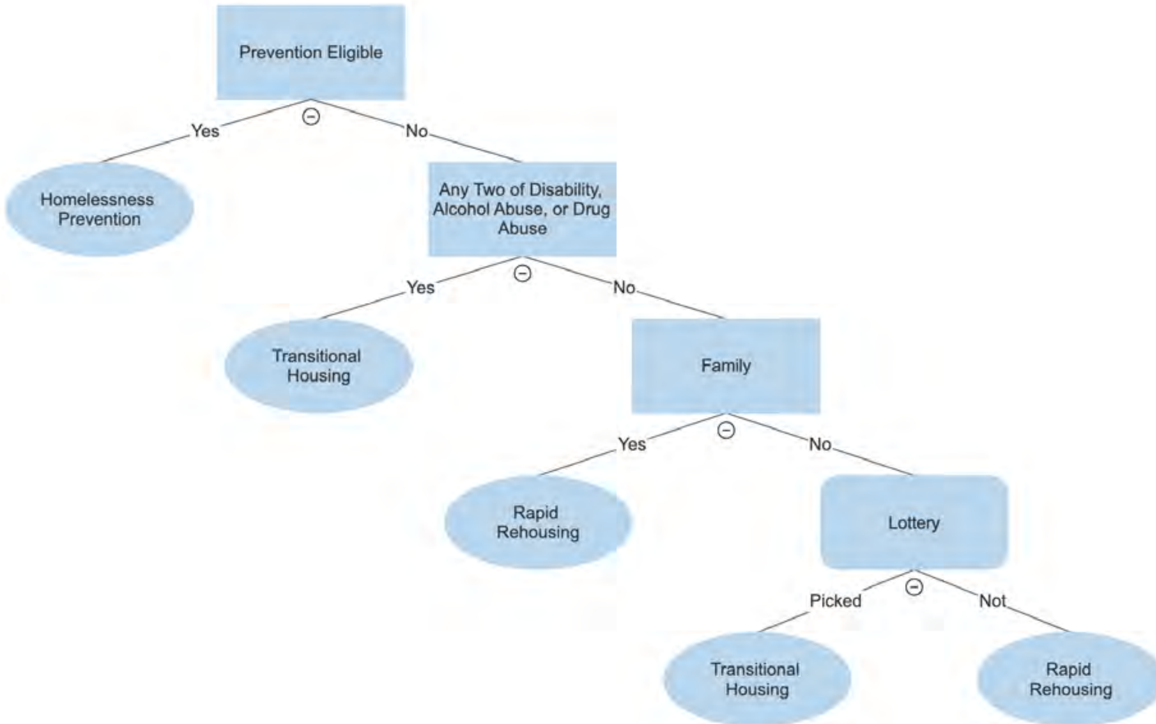

Figure 5.2: Flow chart showing the allocation based on the three developed decision rules

5.2.3 Comparison of Allocations

The results of an allocation based on the 3 proposed decision rules combined with the lottery are provided in Table 5.4 along with the results of the original allocation for comparison. Also included in the table are the results of the optimized allocation subject to budget constraints and an optimized allocation with no constraints which places each household in their "best" predicted service.

Use of the decision rules along with the lottery proposed in the previous section result in an estimated 3.7 percent decrease in probability of reentry compared to the original allocation. This is only a difference of 0.61 percentage points compared to an optimized assignment using a linear program. Though the best possible allocation results in an additional 1.17

Allocation	Estimated Cost (in US Dollars)	Estimated Savings (in US Dollars)	(Expected) Reentry Percentage
Original	5,068,035	-	27.82%
Decision Rules	5,067,472	563	26.78%
Optimized with Budget	5,067,988	47	26.17%
Unconstrained	6,337,995	-1,269,960	25.61%

Table 5.4: Expected cost and reentry percentages for the decision rule and budget allocations compared to the original and unconstrained allocations

percentage point reduction in probability of reentry, it also comes at an estimated cost of $1,269,960 more than what was spent in the original allocation.

5.2.4 Comparison of Group Fairness

The previous section compared allocations by estimated cost and expected percent reentry. In this section, allocations are compared using two group fairness metrics: difference in gain and difference in shortfall. Each of these are ratios comparing to the minimum possible utility and maximum possible utility respectively. The previously mentioned allocations are also compared to average results from 100 runs of a random allocation. The random allocation assigns all prevention eligible households to homelessness prevention and then randomizes allocations for all ineligible households such that the number of households receiving transitional housing and rapid rehousing is the same as that of the allocation based on our decision rules. Figure 5.3 shows the values of these fairness metrics across allocations for both groups affected by the proposed decision rules and other groups of theoretic interest.

The decision rule allocation is clearly favoring those with a comorbidity when considering both gain and shortfall. Interestingly, we see the other allocations favoring those without a comorbidity in terms of shortfall; an example of when different fairness metrics can show different biases/amounts of bias in an allocation. In the case of families without comorbidities,

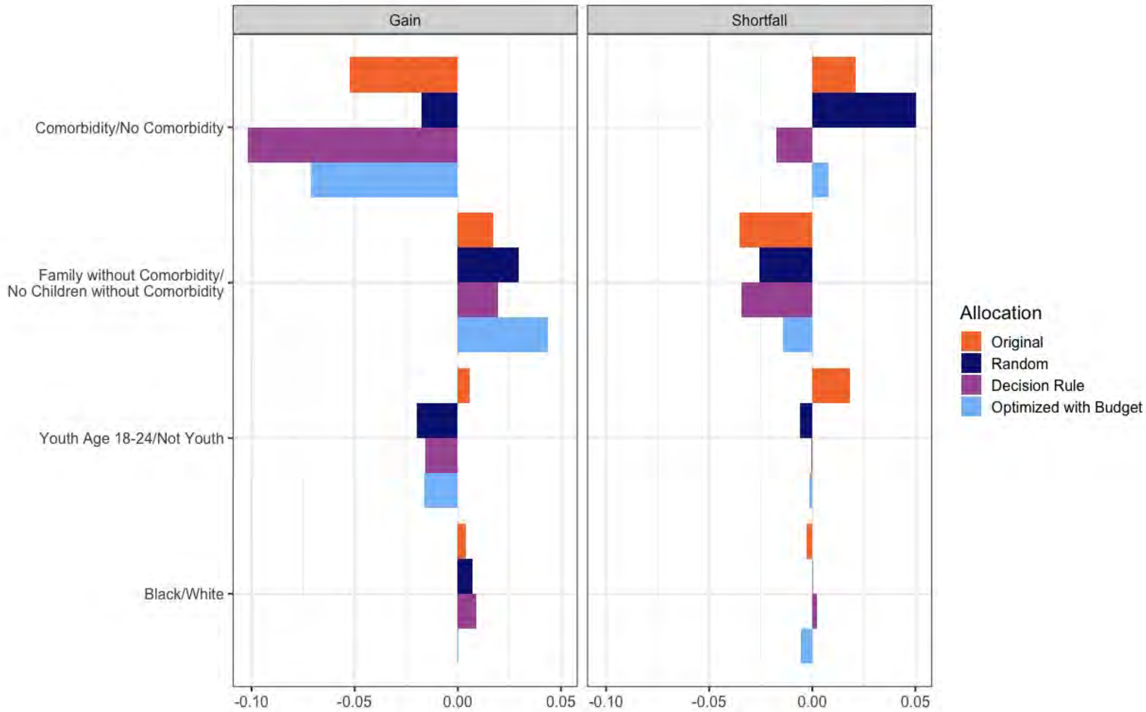

Figure 5.3: Bar chart comparing allocations on both difference in gain and difference in shortfall. Bars pointing to the left of the figure indicate bias toward Group A and bars pointing to the right of the figure indicate bias toward Group B where groups are listed as Group A/Group B

when comparing to the random allocation, the decision rule favors families more under both metrics. In addition, the decision rule allocation favors families more under both metrics than an optimized allocation and has similar values under both metrics to the original allocation. Therefore, as designed the decision rule allocation is favoring families and those with comorbidities.

Values of both gain and to a larger extend shortfall are small in absolute value for comparisons of youth and race. For race in particular, we see no disparate effects due to the small values of both metrics. We see that in terms of gain the decision rule allocation favors youth more so than the original allocation. Overall, we see no disparate effects of the decision rule allocation on the groups tested using either metric.

5.3 Discussion

This work uses machine learning and counterfactual estimation to better understand the efficacy of homeless services for different subpopulations of homeless households. It then builds on this information to develop three data-driven decision rules for placement of households into services. These decision rules improve upon current allocation mechanisms in terms of reductions in reentry into the homeless system while maintaining the same estimated cost as the current mechanism.

Another benefit of this allocation is that it is data-driven yet transparent. Rather than providing caseworkers with raw predictions from a machine learning algorithm or using linear programs to optimize assignment, these decision rules are explainable to stakeholders in the homeless system increasing the chance of implementation relative to "black-box" methods [56, etc].

As with any potential implementation, it is important to understand any bias present in the allocation as well as how its implementation might affect the system as a whole. The previous section shows that it can be difficult to interpret the fairness or the bias of an allocation. Firstly, different definitions of fairness or metrics used to calculate fairness produce different results and interpretations. In addition, one must decide what characteristics or variables they want to be fair or unbiased and whether fairness on that variables means complete lack of bias or bias toward those with protected attributes.

Is is also important to note that all analyses were done using data from one particular city during a span of five years. Any trends seen may differ between locations or may change over time. Therefore it is important to replicate and extend these results to more and larger datasets before implementation in any local homeless system.

In sum, the findings presented here have three main takeaways. The first is that there is opportunity for improvement to the current system of allocation by understanding trends in administrative data and using them to inform current policy. The second is that our evidence suggests households can be better served by allocation of prevention eligible households to homelessness prevention, ineligible comorbid households to transitional housing, and ineligible, non-comorbid families to rapid rehousing. Lastly, additional conversations between researchers and stakeholders are needed to better understand what is needed regarding fairness, transparency, morality, and justice in implementing this type of work.

Chapter 6

Just Resource Allocation? How Algorithmic Predictions and Human Notions of Justice Interact

In the previous chapters we have explored the potential for improving homeless allocation by using machine learning and counterfactual prediction. We have also discussed the fairness of these new allocations. In this chapter we take our dicussion of fairness and justice further by focusing on one central question: *Who should be prioritized for receipt of a scarce resource that is centrally controlled, funded, and allocated, in the absence of a market?*

This question arises in many contexts, ranging from organ donation [37], social service allocation [8], and military service [32], to entries to the New York City marathon. How institutions make these decisions has been studied under the moniker "local justice" in political philosophy, and it has become clear that different types of institutions use a range of different prioritization schemes, ranging from lotteries (military drafts) to prioritizing the most vulnerable (cadaveric organ transplantation) to prioritizing those predicted to benefit most from receipt of the resource (medical triage) [27, 32]. The advent of algorithmic decision making has brought with it the ability to make such prioritization decisions in a more automated manner, often through providing decision support to humans in the form of additional information about those seeking the scarce resource. While there has been considerable attention paid to the possible bias of such algorithmic predictions, a key question that has not been studied thus far is how such predictions interact with decision-maker conceptualizations of justice, and the possible impacts of such interactions on the overall goals of the institution allocating the resources.

Our goal in this chapter is to study exactly this issue using an experimental design for scarce resource allocation in the context of providing scarce homelessness resources. The scarce resource is one unit of *transitional housing*, an intensive and costly service that provides stable housing along with many other forms of support for an extended period. The baseline is *emergency shelter*, a less intensive and costly service that mainly provides space to stay for a more limited time. In the main task, participants are presented with information and asked which of two households they would prioritize for transitional housing.

We are interested in choices about whether to prioritize the *more vulnerable* household for transitional housing, or the household that would have a *better outcome* from receiving transitional housing. We conceptualize baseline vulnerability as the probability of returning to homelessness within the next two years if only given emergency shelter.

We have three major hypotheses.

H1: Decisions on scarce resource allocation primarily fall into two types – outcome-oriented prioritization versus vulnerability-oriented prioritization – reflecting common perceptions of justice. In our experiment, we will observe this as one group of decision-makers allocating transitional housing to households most likely to benefit from the service as characterized by lower probabilities of return to homelessness when receiving transitional housing, while another group prioritizes transitional housing for households with greater perceived need, as characterized by higher probabilities of return to homelessness when receiving only emergency shelter.

H2: Prior exposure to outcome predictions in any form introduces a goal-framing effect, leading to decision-makers becoming more likely to be outcome-oriented in future allocation decisions.

H3: In the absence of defined scarce resource allocation goals, the presentation of algorithmic predictions of outcomes reveals the prioritization types of decision-makers. Among outcome-oriented decision-makers, the group that sees algorithmic predictions of outcomes in addition to household information should make more allocations to transitional housing of those with better predicted outcomes from such allocations. Conversely, among vulnerability-oriented decision-makers, the group that sees algorithmic predictions of outcomes in addition to household information should make more allocations to transitional housing of those with worse predicted outcomes from emergency shelter.

In addition to enhancing our scientific understanding of how human notions of justice play out in scarce resource allocation, our results also have significant policy implications. Homelessness service caseworkers have discretion in their decision-making, and institutional guidelines (which often say to prioritize the most vulnerable in many contexts) often conflict with on-the-ground evaluation measures (where return to homelessness is a significant factor). While our experiments are on a lay audience rather than homelessness caseworkers, they highlight that decision support can have unexpected implications, for example through the framing channel we uncover (H2), as well as through the enhanced ability of decision-makers to make decisions concordant with their type (H3).

6.1 Experimental Design

In this section we discuss our methods for human subjects data collection. Using predictions described in previous chapters and our data on homeless households, we conduct a survey to study decision-making when resources are scarce.

6.1.1 Algorithmic Prediction

Our experimental study consisted of three tasks, which all subjects completed in the same order.[5] Probabilities of reentry from models described previously are binned into three buckets, high, medium, and low. The first task tests the effect of training on assessing vulnerability from vignettes (the **effect-of-training** task). The second task compares prioritizations with or without access to algorithmic predictions of vulnerability (the **effect-of-algo-predictions** task). The third task elicits decision-maker preferences when provided only vulnerability predictions without other information (the **type-elicitation** task). In the first task, subjects predict which probability bucket a household falls into. The second and third tasks ask which of two households to prioritize for transitional housing, the most expensive and intense service, here treated as the scarce resource. The full survey is provided in Appendix E.

Task One - The effect-of-training Task

In the **effect-of-training** task, participants were randomized such that half received training in the form of ten examples of households along with the predicted probability category if that household was placed into emergency shelter or transitional housing for each household. The other half saw the same ten example vignettes without indicating probability categories. Figure 6.1 shows what the no training group sees along with an example of how probability

[5] We note that all experimental protocols were approved by a relevant IRB.

category was indicated to the training group. All participants were then asked to categorize 10 new households as having low, medium, or high probability of needing services again within 2 years given they are placed in transitional housing based on these vignettes (see Figure 6.2).

Task Two - The effect-of-algo-predictions Task

In the **effect-of-algo-predictions** task, participants were presented with 10 pairs of vignettes and were asked to choose which household to prioritize for transitional housing. They were told that households who are not prioritized for transitional housing will receive emergency shelter. Half of the participants are randomized to see vignettes like those shown in Figure 6.3 panel (a). The other half see the same vignettes along with predictions of the probabilities that those households will need future services within 2 years if given transitional housing and if given emergency shelter as in Figure 6.3 panel (b). These predictions were presented as “Low”, “Medium”, or “High”.

Task Three - The type-elicitation Task

In the **type-elicitation** task, participants are again presented with 10 pairs of households and asked to decide which household to prioritize for transitional housing (see Figure 6.4). However, this time, participants are only shown the predictions of probability of needing future services within 2 years given placement in emergency shelter and placement in transitional housing. Participants are randomized into three groups. One third of participants are not told how to make these prioritizations. This is the group we focus on in the main results. Another third are told to make *Vulnerability-Oriented* decisions. The last third are told to make *Outcome-Oriented* decisions. These last two groups are intended to check that the task makes sense and participants can make decisions that are concordant with externally specified

This is the same table presented previously. Models predict that the households highlighted in red have a **high** probability of needing services again within 2 years if they are given **Transitional Housing**.

	Household 1	Household 2	Household 3	Household 4	Household 5	Household 6	Household 7	Household 8	Household 9	Household 10
Number of Household Members	1	1	1	1	1	1	1	2	1	1
Number of Children	0	0	0	0	0	0	0	1	0	0
Head of Household Gender	Female	Male	Male	Male	Female	Female	Male	Female	Female	Female
Head of Household Age	45	35	64	55	55	23	23	23	55	45
Head of Household Disabling Condition	No	Yes	Yes	No	Yes	No	No	No	No	No
Head of Household Receives Substance Abuse Services	No	No	No	No	No	No	Yes	No	No	No
Monthly Income	$800	$668	$1524	$400	$800	$367	$0	$1068	$800	$367
Number of Calls to Hotline	15	5	2	2	2	5	0	0	2	0
Prior Residence	Private Emergency Shelter	Private Emergency Shelter	Private Emergency Shelter	Private Emergency Shelter	Private Emergency Shelter	Private Transitional Housing	Private Emergency Shelter	Private Transitional Housing	Private Transitional Housing	Private Emergency Shelter

(a) An example of what participants who receive training are shown

Here is a table of 10 examples of households needing homeless services. Each column in the table represents a different household and each row represents a piece of information that homeless service providers have when making decisions on which households should receive each kind of service. Take a few minutes to familiarize yourself with the table and try to think of the likelihood of those households needing future services if given Transitional Housing or Emergency Shelter. Then, click "I Understand" when you are ready to proceed.

	Household 1	Household 2	Household 3	Household 4	Household 5	Household 6	Household 7	Household 8	Household 9	Household 10
Number of Household Members	1	1	1	1	1	1	1	2	1	1
Number of Children	0	0	0	0	0	0	0	1	0	0
Head of Household Gender	Female	Male	Male	Male	Female	Female	Male	Female	Female	Female
Head of Household Age	45	35	64	55	55	23	23	23	55	45
Head of Household Disabling Condition	No	Yes	Yes	No	Yes	No	No	No	No	No
Head of Household Receives Substance Abuse Services	No	No	No	No	No	No	Yes	No	No	No
Monthly Income	$800	$668	$1524	$400	$800	$367	$0	$1068	$800	$367
Number of Calls to Hotline	15	5	2	2	2	5	0	0	2	0
Prior Residence	Private Emergency Shelter	Private Emergency Shelter	Private Emergency Shelter	Private Emergency Shelter	Private Emergency Shelter	Private Transitional Housing	Private Emergency Shelter	Private Transitional Housing	Private Transitional Housing	Private Emergency Shelter

(b) An example of what participants who do not receive training are shown

Figure 6.1: A comparison of the training and no training conditions in Task 1.

Categorize the following households based on how likely you predict they are to need further services within 2 years if placed into Transitional Housing. Click on the household number label below the image to select your answer.

	Household 1	Household 2	Household 3	Household 4	Household 5	Household 6	Household 7	Household 8	Household 9	Household 10
Number of Household Members	1	1	1	1	1	1	1	1	1	1
Number of Children	0	0	0	0	0	0	0	0	0	0
Head of Household Gender	Male	Male	Female	Male	Female	Male	Male	Male	Male	Male
Head of Household Age	35	45	35	45	45	45	55	45	45	35
Head of Household Disabling Condition	Yes	Yes	No	No	No	Yes	No	Yes	No	No
Head of Household Receives Substance Abuse Services	Yes	No	No	No	Yes	Yes	No	Yes	Yes	No
Monthly Income	$400	$1524	$668	$200	$800	$668	$0	$0	$200	$0
Number of Calls to Hotline	5	1	15	25	1	3	25	5	2	25
Prior Residence	Private Emergency Shelter	Place not meant for habitation	Permanent housing for formerly homeless persons	Private Emergency Shelter	Place not meant for habitation	Private Emergency Shelter	Private Emergency Shelter	Place not meant for habitation	Private Emergency Shelter	Private Emergency Shelter

High probability of future need in Transitional Housing | Medium probability of future need in Transitional Housing | Low probability of future need in Transitional Housing

Household 1 Household 2 Household 3 Household 4 Household 5 Household 6 Household 7 Household 8 Household 9 Household 10

Figure 6.2: The question posed to participants in Task 1.

Which of the following households would you prioritize for Transitional Housing? Select your choice by clicking on the associated column in the image below.

	Household 1	Household 2
Number of Household Members	1	1
Number of Children	0	0
Head of Household Gender	Female	Female
Head of Household Age	35	45
Head of Household Disabling Condition	No	No
Head of Household Receives Substance Abuse Services	No	Yes
Monthly Income	$668	$800
Number of Calls to Hotline	15	1
Prior Residence	Permanent housing for formerly homeless persons	Place not meant for habitation

(a) An example question presented to the vignette only group

Which of the following households would you prioritize for Transitional Housing? Select your choice by clicking on the associated column in the image below.

	Household 1	Household 2
Number of Household Members	1	1
Number of Children	0	0
Head of Household Gender	Female	Male
Head of Household Age	55	35
Head of Household Disabling Condition	Yes	No
Head of Household Receives Substance Abuse Services	No	Yes
Monthly Income	$0	$0
Number of Calls to Hotline	5	2
Prior Residence	Private Emergency Shelter	Private Emergency Shelter
Predicted probability of needing future services within 2 years if given Transitional Housing	High	Low
Predicted probability of needing future services within 2 years if given Emergency Shelter	High	Medium

(b) An example question presented to the vignette and prediction group

Figure 6.3: A comparison of questions from Task 2 for each randomization group

goals. Participants then see two examples explaining how to make *Outcome-Oriented* versus *Vulnerability-Oriented* decisions. Lastly, they are told which goal to focus on and reminded of the definition of that goal before being presented with pairs of households to decide between.

Which of the following households would you prioritize for Transitional Housing? Select your choice by clicking on the associated column in the image below.

	Household 1	Household 2
Predicted probability of needing future services within 2 years if given Transitional Housing	Low	Medium
Predicted probability of needing future services within 2 years if given Emergency Shelter	High	High

Figure 6.4: An example question from Task 3

Survey Statistics

A total of 520 participants completed our survey, recruited through Amazon's Mechanical Turk platform. In both the **effect-of-algo-predictions** task and the **type-elicitation** task, two questions were duplicated as a reliability check. Any participant who answered inconsistently on both duplicate questions in either task was dropped from the study, resulting in 458 participants.[6] Responses were restricted to come from English speaking persons over age 18 in the United States. Of the 458 respondents, 38.81% identified as female. In a question where participants were asked to select their race within the ability to select multiple races, 86.57% identified exclusively as white. The average age of participants was 42.34 years (SD = 12.34).

6.2 Results

In this section, we provide evidence for our three hypotheses. First, we show there are two main "types" of decision-makers – labeled as *Vulnerability-Oriented* and *Outcome-Oriented* types – from the **type-elicitation** task. Note that we are not claiming that the type is

[6]We repeated all analyses using the full set of participants and results remained unchanged.

intrinsic and unchangeable, this is the type at the time of facing the decision-making task of choosing which household to allocate the scarce resource (transitional housing) to. Second, the type of a decision maker, as determined from the **type-elicitation** task, is affected by randomization in the **effect-of-training** and **effect-of-algo-predictions** tasks such that prior exposure to predictions increases the likelihood that a decision-maker is of the *Outcome-Oriented* type. Third, the **effect-of-algo-predictions** task demonstrates that a decision-maker's type determines the effect of providing vulnerability predictions in addition to vignettes about household characteristics. *Vulnerability-Oriented* types consistently make more vulnerability-oriented decisions when provided with predictions, while *Outcome-Oriented* types consistently make more outcome-oriented decisions when provided with predictions.

6.2.1 H1 - Type Revelation

Recall our first hypothesis,

H1: Decisions on scarce resource allocation primarily fall into two types – outcome-oriented prioritization versus vulnerability-oriented prioritization – reflecting common perceptions of justice. In our experiment, we will observe this as one group of decision-makers allocating transitional housing to households most likely to benefit from the service as characterized by lower probabilities of return to homelessness when receiving transitional housing, while another group prioritizes transitional housing for households with greater perceived need, as characterized by higher probabilities of return to homelessness when receiving only emergency shelter.

As a reminder, in the **type-elicitation** task, we elicited prioritization goals of participants by providing them with predictions of need for future services conditional on (1) receiving transitional housing support and (2) receiving space in an emergency shelter. Given a pair of households, participants were asked to decide which of the pair should be prioritized for transitional housing, given that the other would receive space in an emergency shelter (see Figure 6.4). Outcomes for transitional housing were always at least as good as those for shelter. Subjects were shown predictions of likelihood those households would need future services within 2 years. These predictions were presented as “Low”, “Medium”, or “High.” These predictions were based on a machine learning model trained on administrative data from actual households, as described in Section 6.1.1.

The framework of local justice [32] suggests that, given information of this kind, humans who take this information into account (as opposed to deciding randomly, for example), are likely to base their decisions on one of three possible criteria: prioritize (1) the household deemed to be most vulnerable prior to allocation; (2) the household that would be least vulnerable after allocation; or (3) the household whose vulnerability status would change the most due to the allocation. Since we use only three probability buckets to assess vulnerability to limit cognitive load, and are further limited by the constraint that transitional housing, as the scarce resource, is always at least as good as emergency shelter, it is difficult to design instances to cleanly disambiguate criterion 3 from the other two, so we focus on criteria (1) and (2).

To identify *Vulnerability-Oriented* and *Outcome-Oriented* types, we assign a score to each prioritization decision, which is a 0 if the decision is inconsistent with that prioritization type, and a 1 if it is consistent. For example, given the question posed in Figure 6.4, a participant who chose Household 1 made a *Outcome-Oriented* decision. As a result, we would give them a *Outcome-Oriented* score of 1 on this question. As there was no distinct *Vulnerability-Oriented* decision on this question, no *Vulnerability-Oriented* scores were assigned for this question. We sum up the scores for both criteria for each subject, and then scale the total scores to the range $[0, 10]$. We have three different randomized groups in this task – one group was told to prioritize according to the *Vulnerability-Oriented* criterion, a second group was told to prioritize according to the *Outcome-Oriented* criterion, and the third group was not given specific instructions on whom to prioritize. The first two groups serve as checks to make sure the task is properly designed (and the results are as expected – see Appendix C for details), while the third group is the one we use to assess types.

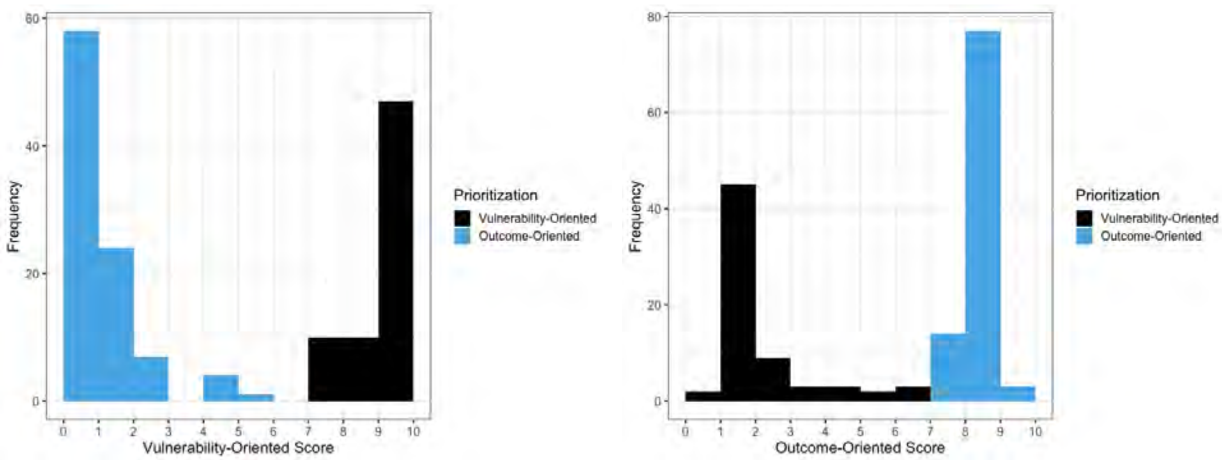

(a) Histogram of Vulnerability-Oriented Scores (b) Histogram of Outcome-Oriented Scores

Figure 6.5: Histograms showing the distributions of *Vulnerability-Oriented* (a) and *Outcome-Oriented* (b) scores elicited without a prioritization goal colored by revealed prioritization type (blue for *Outcome-Oriented* and black for *Vulnerability-Oriented*)

Figure 6.5 shows the distributions of the *Vulnerability-Oriented* and *Outcome-Oriented* scores for the group that was not asked to make a particular prioritization. We see a clear distributional difference and distinction in scores between the two types as colored by prioritization group. We define *Vulnerability-Oriented* types as those with *Vulnerability-Oriented* scores of 7 or above, and *Outcome-Oriented* types as those with *Outcome-Oriented* scores of 7 or above. Those with *Vulnerability-Oriented* scores of 7 or above are considered the *Vulnerability-Oriented* group. Of the 179 participants in the no goal group, 94 were in the the *Outcome-Oriented* group, 67 were in the *Vulnerability-Oriented* group, and 18 did not meet criteria for either group (both their *Outcome-Oriented* and *Vulnerability-Oriented* scores were below 7). Therefore, 90% of the participants were very consistent in their decision-making. H1 is clearly well-supported by the data.

However, as an alternative way of assigning types, *Vulnerability-Oriented* types can be defined as those whose *Vulnerability-Oriented* score is greater than their *Outcome-Oriented* score and similarly *Outcome-Oriented* types as those whose *Outcome-Oriented* score is greater than their *Vulnerability-Oriented* score. This definition results in almost the same type distribution with all 94 *Outcome-Oriented* types remaining *Outcome-Oriented* and 63 of 67 *Vulnerability-Oriented* types remaining *Vulnerability-Oriented* . Of the 18 who did not meet criteria for either type, 14 would become *Outcome-Oriented* and the remaining 4 *Vulnerability-Oriented* . Appendix D shows that our subsequent results are robust to this change in type definition.

6.2.2 H2 - Exposure to Predictions Primes Decision Makers to Become Outcome-Oriented

H2: Prior exposure to outcome predictions in any form introduces a goal-framing effect, leading to decision-makers becoming more likely to be outcome-oriented in future allocation decisions.

In order to test this hypothesis, we separated respondents who were in the "no goal" condition in the **effect-of-algo-predictions** task into four groups based on which randomization they received in the **effect-of-training** and **effect-of-algo-predictions** tasks. Of the four groups, only one had no prior exposure to outcome predictions (the "No Training + Vignette Only" group), whereas each of the other groups had seen outcome predictions in at least one of their prior tasks. Table 6.1 shows how many in each of the four groups ended up being identified as outcome-oriented versus vulnerability-oriented. The results are stark. By an almost 2:1 ratio, those with prior exposure to outcome predictions reveal themselves as outcome oriented, while those without reveal themselves as vulnerability oriented by almost the same ratio. The results are significant at the 0.01 level by a Fisher exact test (statistic value 0.0017).

Randomization Group	*Outcome-Oriented*	*Vulnerability-Oriented*
Training + Vignette Only Group	$N = 28$	$N = 12$
Training + Vignette and Predictions Group	$N = 27$	$N = 14$
No Training + Vignette Only Group	$N = 12$	$N = 23$
No Training + Vignette and Predictions Group	$N = 27$	$N = 18$

Table 6.1: Number of participants of each prioritization type based on the **type-elicitation** task for each possible combination of randomizations in the **effect-of-training** and **effect-of-algo-predictions** tasks.

This essentially indicates that the decision making population is divided into three major sets of individuals, in almost equal proportions. Those who are vulnerability oriented and will remain so regardless of exposure to predictions, those who are outcome oriented and will remain so regardless of exposure, and those who would be vulnerability oriented, but switch to being outcome oriented once exposed to information about outcomes.

6.2.3 H3 - The Effect of Algorithmic Predictions on Decision-Making

We now turn to evidence on the differential impacts of providing algorithmic risk predictions for *Vulnerability-Oriented* and *Outcome-Oriented* types.

H3: In the absence of defined scarce resource allocation goals, the presentation of algorithmic predictions of outcomes reveals the prioritization types of decision-makers. Among outcome-oriented decision-makers, the group that sees algorithmic predictions of outcomes in addition to household information should make more allocations to transitional housing of those with better predicted outcomes from such allocations. Conversely, among vulnerability-oriented decision-makers, the group that sees algorithmic predictions of outcomes in addition to household information should make more allocations to transitional housing of those with worse predicted outcomes from emergency shelter.

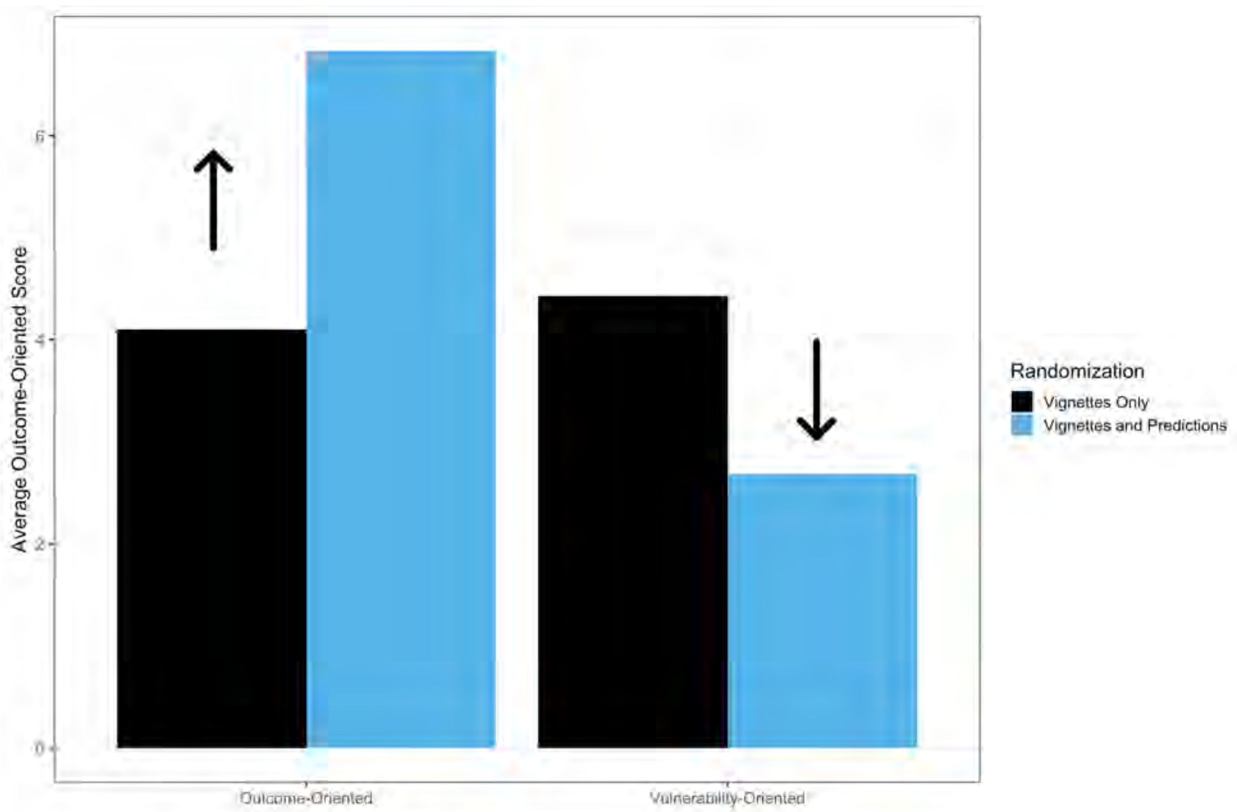

Figure 6.6: Average *Outcome-Oriented* score for each prioritization type in the **effect-of-algo-predictions** task across vignette only (black) and vignette-and-risk-prediction (blue) groups. Arrows show that, when shown predictions along with vignettes, those in the *Outcome-Oriented* group have higher *Outcome-Oriented* scores and those in the *Vulnerability-Oriented* group have lower *Outcome-Oriented* scores.

In the **effect-of-algo-predictions** task, subjects were shown ten pairs of households, and asked to prioritize one for transitional housing (see Figure 6.3). In each pair, one household always corresponded to the *Vulnerability-Oriented* prioritization and the other to the *Outcome-Oriented* prioritization. Subjects were randomized to see either just vignettes with information about the households, or the vignettes plus algorithmic risk prediction categories (i.e., low, medium, or high probability of re-entry conditional on receiving transitional housing or emergency shelter). We computed *Outcome-Oriented* scores for all participants (on a scale of 10 – in this task the *Vulnerability-Oriented* and *Outcome-Oriented* scores always sum to 10).

Figure 6.6 shows both *Vulnerability-Oriented* and *Outcome-Oriented* types have similar *Outcome-Oriented* scores (4.42(2.05) and 4.10(2.09) respectively, showing a slight lean towards prioritizing vulnerability) when shown just vignettes (no significant difference; p-value of 0.49). However, when shown both vignettes plus risk predictions, the *Vulnerability-Oriented* types see their scores decline to 2.67(2.19), showing that they become much more aligned with making *Vulnerability-Oriented* decisions. The *Outcome-Oriented* types, on the other hand, see a dramatic *increase* in their scores, to 6.83(2.80), showing that they become much more aligned with making *Outcome-Oriented* decisions. These differences are statistically significant (p-value = 5.04e-11) and clearly substantial in their effects. This is perhaps our most salient result. Additional information, in the form of algorithmic predictions, allows individuals to consistently make decisions that are either vulnerability- or outcome- oriented, with which one being determined by that individual. This also means that the behavior of individual decision-makers should be highly predictable after seeing a few examples (when they are presented with predictions) – they will quickly reveal whether they are vulnerability- or outcome- oriented.

We note that the fact that the **type-elicitation** task is chronologically last for each subject is so that being asked to perform this task does not affect subjects' performance of the **effect-of-algo-predictions** task (in particular, the group shown just the vignette in that task). Therefore, the results related to H2 are only related to **type-elicitation** in that **type-elicitation** determines the categorization of individuals into types.

6.2.4 Potential Channels for the Effect in Hypothesis 3,

We now turn to a discussion of the relative impacts of two potential channels, the *information* channel and the *framing* channel, for the differences in behavior when subjects see algorithmic risk predictions in addition to vignettes. If the effect were entirely or largely through the information channel, this would mean that subjects are attempting to make decisions that align with their prioritization types, but vignettes do not provide them sufficient information on vulnerability, therefore their decisions are noisier. The framing channel would instead imply that adding the information on algorithmic risk predictions makes subjects think about the implications of their choices differently, leading to different decisions.

As mentioned in Section 6.2.1, we have evidence that there are three main sets of individuals in the population. For simplicity, let us call them infungible outcome-oriented (IOO), infungible vulnerability-oriented (IVO), and vulnerability-to-outcome oriented (VOO). We assume that the effect we are seeing is a combination of providing the IOO and IVO individuals better information with which to make their decisions (the information effect), and of both framing and information for the VOO individuals. While disambiguating these channels is beyond the scope of the experimental design in this paper, we can shed some light on the relative information effect by considering the impact of training.

In theory, if training were perfectly successful, it would allow those who are trained to learn nothing additional from being presented predictions in addition to vignettes, and for them the information channel would be irrelevant. Then we would expect to see the training +

vignette groups performing very similarly to the groups that see vignettes + predictions, which is clearly not the case.

Analyzing data from the **effect-of-training** task allows us to directly ask if subjects can be trained to make more accurate risk predictions from vignettes. In this task, subjects are shown repeated examples of vignettes paired with the algorithmic risk predictions. We then test whether they are better able to assess the risk level predicted by the algorithm. The results show a significant improvement in classification accuracy (percentage correctly classified as low, medium, or high probability of re-entry) when participants receive training (Average Classification Score for Training Group = 4.75(1.80) N = 212, Average Classification Score for No Training Group = 3.69(1.52) N = 246, p-value = 4.73e-11). However, while statistically significant, the substantive impact on prediction accuracy here is small.

In combination, our results are supportive of the hypothesis that both channels play a role in the sharp differences we see in allocation decisions between those who receive vignettes + predictions versus those who receive only vignettes. The information channel is likely dominant for IVO and IOO individuals, while both channels will affect the behavior of VOO individuals.

6.3 Additional Results and Implications

In this section, we dig deeper into some of the results above. The results and analysis presented here are more speculative, but highlight interesting directions and connections to the existing literature.

6.3.1 The Interaction of Framing and Information

The cognitive bias literature suggests that presenting algorithmically derived predictions could introduce goal-framing by focusing attention on the outcomes of decisions [57, 76]. What would this suggest for the relative efficiency scores of different groups based on their types and randomization conditions? For the outcome-oriented types, we expect the ranking of efficiency scores would be:

1. Training + Predictions would have the highest outcome-oriented scores because of repeated frames for outcomes as a goal and the presence of predictions at task time.

2. No Training + Predictions would have the next highest outcome-oriented scores because there are no conflicting signals for making outcome-orientation a goal.

3. Training + No Predictions would have the third highest because of a combination of framing (prior exposure through training) and the information channel (being able to make somewhat better predictions because of training).

4. No Training + No Predictions would have the lowest efficiency among outcome-orientated decision-makers given the lack of additional outcome framing.

This is indeed the ranking we observe (Table 6.2), although some of the differences are small.

Expectations for the vulnerability-oriented decision-makers are less apparent. For these decision-makers, prior exposure to the outcome predictions did not change their orientation; however, the presence of algorithmically derived predictions could conflict with their vulnerability goal frame. It is interesting that Table 6.2 shows the lowest efficiency scores appear among vulnerability-oriented decision-makers who are presented predictions at task time; they may be reacting to the conflicting outcome frame and using the additional information from the predictions in confirming their vulnerability-orientation. The next lowest efficiency scores come from trained vulnerability-oriented decision-makers presented with predictions at task time; they may react to the framing but still attend to the additional information from predictions. The vulnerability-oriented vignette-only groups scores similarly to the outcome-oriented decision-makers from vignette-only groups.

We acknowledge that the small numbers in each category necessarily make this analysis speculative, but simultaneously, suggestive of an interesting avenue for future exploration. There could also be other framing effects in play, for example one that focuses decisions on specific attributes available in the vignettes, such as the presence of children, that might interact with goal framing [57].

Randomization Group	Prioritization Type	Efficiency Score M(SD)
Training + Vignette and Predictions Group	*Outcome-Oriented*	7.19(2.53), $N = 27$
No Training + Vignette and Predictions Group	*Outcome-Oriented*	6.48(3.06), $N = 27$
No Training + Vignette Only Group	*Vulnerability-Oriented*	4.70(2.01), $N = 23$
Training + Vignette Only Group	*Outcome-Oriented*	4.11(2.15), $N = 28$
No Training + Vignette Only Group	*Outcome-Oriented*	4.08(2.02), $N = 12$
Training + Vignette Only Group	*Vulnerability-Oriented*	3.92(2.11), $N = 12$
Training + Vignette and Predictions Group	*Vulnerability-Oriented*	3.07(2.13), $N = 14$
No Training + Vignette and Predictions Group	*Vulnerability-Oriented*	2.39(2.25), $N = 18$

Table 6.2: Means and standard deviations of efficiency scores on the **effect-of-algo-predictions** task for each possible combination of randomizations across prioritization types in descending order

6.3.2 Consistency of Decision-Making

Figure 6.7 shows histograms for the outcome-oriented score for each type, separated into those who saw only the vignettes and those who saw both vignettes and predictions. If individual decision-makers are consistent, we would expect to see many scores near the extremes, with

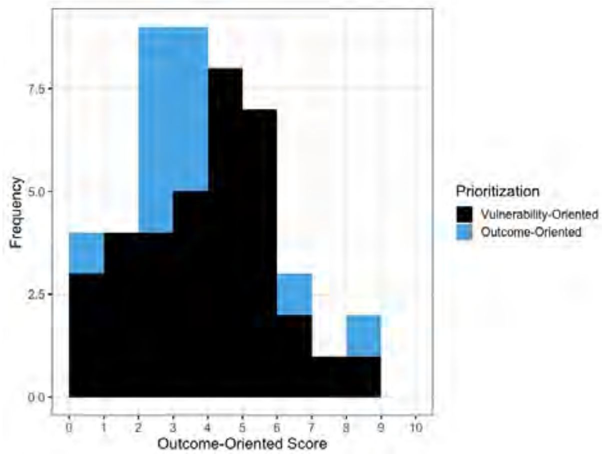

(a) Vignette Only

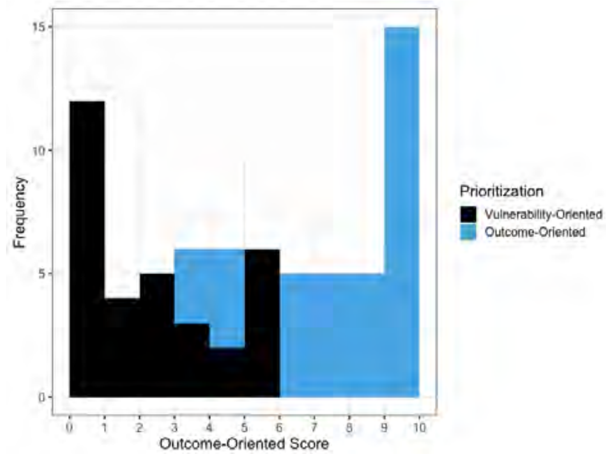

(b) Vignette and Prediction

Figure 6.7: Histograms showing the distribution of *Outcome-Oriented* scores on the **effect-of-algo-predictions** task for *Outcome-Oriented* types (blue) and *Vulnerability-Oriented* types (black)

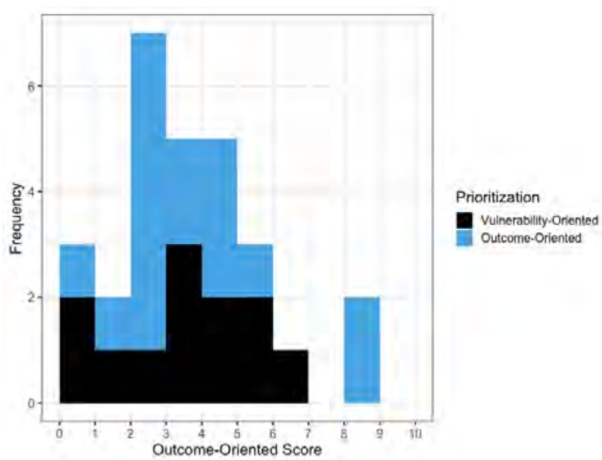

(a) Training Group

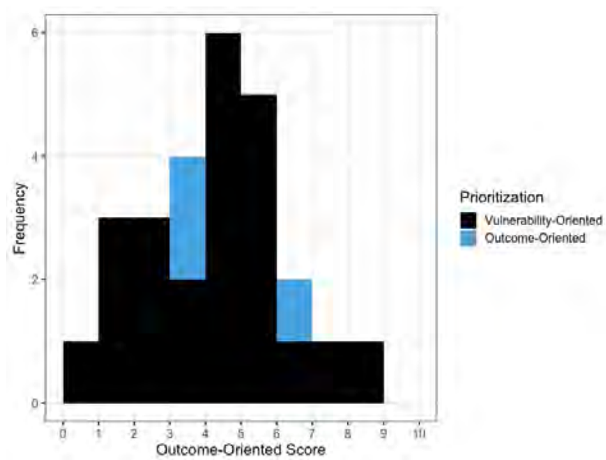

(b) No Training Group

Figure 6.8: Histograms showing the distribution of *Outcome-Oriented* scores on the **effect-of-algo-predictions** task for those in the vignette only group separated by *Outcome-Oriented* types (blue) and *Vulnerability-Oriented* types (black)

vulnerability-oriented decision-makers scoring low and outcome-oriented decision-makers scoring high. On the other hand, if individual decision-makers are inconsistent in their prioritization, that would likely manifest as a unimodal distribution with modes and means close to the middle for both types. Figure 6.7 shows clearly that decisions are much more inconsistent when decision-makers are provided only vignettes, and are consistent when they are also provided predictions. Looking at just the means of the distributions would be confounding in this case, because the mean may not change much since each type is driven to a different extreme of the distribution.

We can also examine whether this effect is because of provision of algorithmic predictions or of framing. Figure 6.8 shows that there are no major differences in the histograms by type between the training and no-training groups, demonstrating quite clearly that the effect is driven by the provision of predictions at the time of decision-making. These results could have important implications in practice, since they imply an increase in procedural fairness in such scarce resource allocation tasks when additional information is provided. Inconsistency in such decisions can make them seem arbitrary to those subject to the decisions, and can often lead to a loss of trust in the system.

6.4 Discussion and Policy Implications

It is important to note that there is no prioritization type that our work deems better or more useful than the other. Current homeless policy states that prioritization for housing should be based on risk [45]. This would coincide with a vulnerability-oriented prioritization. However, homeless services are often evaluated based on efficiency; as evidenced by the literature's focus on 2-year reentry as a metric for the usefulness of an intervention, which more closely corresponds to an outcome-oriented prioritization. This difference in priorities results in difficulties in the current system which could be exacerbated by the inclusion of an algorithmic decision-aid. Introducing additional information that can change the priorities of decision-makers should not be done without additional research. While the current study focuses on the decision-making of the general population as represented by

Mechanical Turk workers, it is important to study decision-making in those training to make these decisions. The replication of this work with homeless caseworkers would help to further understand prioritization preferences as well as how, if at all, the addition of information from an algorithmic decision aid might affect homeless service allocation.

In addition, we acknowledge that those who complete surveys on Mechanical Turk might not be representative of the U.S. population as a whole. For example, our respondents were overwhelmingly white and male. There are also many possible prioritization schemes aside from the two main schemes we focus on here; though these are most closely related with the current priorities of the homeless system. Further insight into how participants are making these decisions and what criteria they weigh most heavily would help to determine additional prioritization schemes or features of families that are deemed most important in making these decisions.

Overall, our findings suggest that additional information from algorithmic decision-aids might affect more than just the efficiency or fairness of decisions made in societal contexts. The priorities and focus of human-decision makers might become more polarized and thus might fall out of line with the priorities of the social system or society as a whole. Therefore, it is important to further understand, not merely the fairness of the tools' output, or the moral reasoning of the tool, but the morals of introducing these tools and the effects they could have on our society.

Chapter 7

Conclusion

This thesis investigates potential for improvement at the intersection of AI and homeless policy from model building using administrative data. We show that use of machine learning and optimization to leverage heterogeneity in effect of services improves efficiency. However, as there is no Pareto improvement, reduced overall reentry comes at the expense of increased reentry probability for some households. As batch optimization is a simplified version of the online allocation process that truly takes place, these results serve as a proof-of-concept and show that the potential for improvement exists and that fairness tradeoffs must be considered when putting such work into practice.

There are a multitude of ways that information from our model might be used to augment current allocation decisions. We explore the use of our models for understanding portions of the homeless population that are best served by particular homeless services and use this information to make simple and interpretable decision rules which we compare to the original and other benchmark allocation mechanisms on both cost and efficiency. We find that use of these decision rules leads to improved efficiency while maintaining the current cost of services.

We also discuss the fairness of this allocation mechanism compared to our benchmarks as well as tradeoffs between group fairness metrics. We find that [add here after analysis].

Another potentiality is that predictions from models like ours could be provided to homelessness caseworkers as additional information they can use to make their decisions. A key question here is how this additional information might affect decision-making. We find that inclusion of additional information from algorithmic decision-aids affects the priorities and focus of human-decision makers, causing them to become more polarized. These results emphasize the importance of studying efficiency, fairness/morality, and the potential moral or societal affects of implementing such tools.

Our aim is that this work leads to additional discussion for how to improve upon homeless allocation in both the scientific community and the community at large. We emphasize throughout our work the many factors that must be kept in mind throughout the research pipeline when studying scarce societal resources. The goal is to improve outcomes for homeless households, but some improvements come at the cost of outcomes for others. It is essential to attempt to remain unbiased or fair, but how and when fairness is determined is a topic of debate. We argue that there exist moral considerations at both the optimization-level and the implementation-level. This is not only true in our setting, but in many setting like ours where essential resources are scarce. We hope that our investigation prompts others to reason deliberately about the morals involved in introducing machine learning and AI into decision-making processes.

References

[1] Atila Abdulkadiroglu, Parag A Pathak, Alvin E Roth, and Tayfun Sonmez. "The Boston public school match." In: *American Economic Review* (2005), pp. 368–371.

[2] Mohammad Akbarpour, Shengwu Li, and Shayan Oveis Gharan. "Thickness and information in dynamic matching markets." In: *Journal of Political Economy* 128.3 (2020), pp. 783–815.

[3] Julia Angwin, Jeff Larson, Surya Mattu, and Lauren Kirchner. *Machine Bias*. 2016.

[4] 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*. 2013, pp. 60–66.

[5] Elliot Anshelevich and Sanmay Das. "Matching, Cardinal Utility, and Social Welfare." In: *ACM SIGECom Exchanges* 9.1 (2010).

[6] Elliot Anshelevich, Sanmay Das, and Yonatan Naamad. "Anarchy, stability, and utopia: Creating better matchings." In: *Autonomous Agents and Multi-Agent Systems* 26.1 (2013), pp. 120–140.

[7] Edmond Awad, Sohan Dsouza, Richard Kim, Jonathan Schulz, Joseph Henrich, Azim Shariff, Jean-François Bonnefon, and Iyad Rahwan. "The moral machine experiment." In: *Nature* 563.7729 (2018), pp. 59–64.

[8] 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: *International Conference on the Integration of Constraint Programming, Artificial Intelligence, and Operations Research*. Springer. 2018, pp. 35–51.

[9] Boris Babic, Anil Gaba, Ilia Tsetlin, and Robert Winkler. "Normativity, Epistemic Rationality, and Noisy Statistical Evidence." In: *British Journal for the Philosophy of Science* (2021).

[10] Kirk Bansak, Jeremy Ferwerda, Jens Hainmueller, Andrea Dillon, Dominik Hangartner, Duncan Lawrence, and Jeremy Weinstein. "Improving refugee integration through data-driven algorithmic assignment." In: *Science* 359.6373 (2018), pp. 325–329.

[11] Ruha Benjamin. “Assessing risk, automating racism.” In: *Science* 366.6464 (2019), pp. 421–422.

[12] Leo Breiman. “Random forests.” In: *Machine learning* 45.1 (2001), pp. 5–32.

[13] Anna Brown, Alexandra Chouldechova, Emily Putnam-Hornstein, Andrew Tobin, and Rhema Vaithianathan. “Toward Algorithmic Accountability in Public Services: A Qualitative Study of Affected Community Perspectives on Algorithmic Decision-Making in Child Welfare Services.” In: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM. 2019, 41:1–12.

[14] Molly Brown, Camilla Cummings, Jennifer Lyons, Andrés Carrión, and Dennis P Watson. “Reliability and validity of the Vulnerability Index-Service Prioritization Decision Assistance Tool (VI-SPDAT) in real-world implementation.” In: *Journal of Social Distress and the Homeless* 27.2 (2018), pp. 110–117.

[15] 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.” In: *Journal of Banking and Finance* 72 (Nov. 2016), pp. 218–239.

[16] Thomas Byrne, Ann Elizabeth Montgomery, and Jamison D Fargo. “Predictive modeling of housing instability and homelessness in the Veterans Health Administration.” In: *Health services research* 54.1 (2019), pp. 75–85.

[17] 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.” In: *2017 AAAI Fall Symposium Series*. 2017.

[18] Cheng Chen, Lan Zheng, Venkatesh Srinivasan, Alex Thomo, Kui Wu, and Anthony Sukow. “Conflict-aware weighted bipartite b-matching and its application to e-commerce.” In: *IEEE Transactions on Knowledge and Data Engineering* 28.6 (2016), pp. 1475–1488.

[19] Hugh A Chipman, Edward I George, Robert E McCulloch, et al. “BART: Bayesian additive regression trees.” In: *The Annals of Applied Statistics* 4.1 (2010), pp. 266–298.

[20] Hugh A. Chipman, Edward I. George, and Robert E. McCulloch. “Bayesian Ensemble Learning.” In: *Advances in Neural Information Processing Systems* 19 (2007), pp. 265–272.

[21] Alexandra Chouldechova, Diana Benavides-Prado, Oleksandr Fialko, and Rhema Vaithianathan. “A case study of algorithm-assisted decision making in child maltreatment hotline screening decisions.” In: *Proceedings of the Conference on Fairness, Accountability and Transparency*. 2018, pp. 134–148.

[22] *Coordinated Entry Evaluation City of St. Louis Continuum of Care - Final Report*. Dec. 2019.

[23] Sam Corbett-Davies and Sharad Goel. “The measure and mismeasure of fairness: A critical review of fair machine learning.” In: *arXiv preprint arXiv:1808.00023* (2018).

[24] 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*. ACM. 2017, pp. 797–806.

[25] Diane Coyle and Adrian Weller. "“Explaining” machine learning reveals policy challenges." In: *Science* 368.6498 (2020), pp. 1433–1434. DOI: 10.1126/science.aba9647.

[26] Dennis P Culhane, Jung Min Park, and Stephen Metraux. "The patterns and costs of services use among homeless families." In: *Journal of Community Psychology* 39.7 (2011), pp. 815–825.

[27] Sanmay Das. "Local Justice and the Algorithmic Allocation of Scarce Societal Resources." In: *Proceedings of the AAAI Conference on Artificial Intelligence*. 2022.

[28] David Delacrétaz, Scott Duke Kominers, and Alexander Teytelboym. *Refugee resettlement*. 2016.

[29] John P Dickerson, David F Manlove, Benjamin Plaut, Tuomas Sandholm, and James Trimble. "Position-indexed formulations for kidney exchange." In: *Proceedings of the 2016 ACM Conference on Economics and Computation*. ACM, 2016, pp. 25–42.

[30] Bistra Dilkina and Carla P Gomes. "Solving connected subgraph problems in wildlife conservation." In: *International Conference on Integration of Artificial Intelligence (AI) and Operations Research (OR) Techniques in Constraint Programming*. Springer. 2010, pp. 102–116.

[31] Julia Dressel and Hany Farid. "The accuracy, fairness, and limits of predicting recidivism." In: *Science Advances* 4.1 (2018), eaao5580.

[32] Jon Elster. *Local Justice: How Institutions Allocate Scarce Goods and Necessary Burdens*. New York, NY: Russell Sage Foundation, 1992.

[33] Danielle Ensign, Sorelle A Friedler, Scott Neville, Carlos Scheidegger, and Suresh Venkatasubramanian. "Runaway feedback loops in predictive policing." In: *Conference on Fairness, Accountability and Transparency*. PMLR. 2018, pp. 160–171.

[34] Virginia Eubanks. *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor*. New York, NY: St. Martin's Press, 2018.

[35] Patrick J Fowler, Kenneth Wright, Katherine E Marcal, Ellis Ballard, and Peter S Hovmand. "Capability Traps Impeding Homeless Services: A Community-Based System Dynamics Evaluation." In: *Journal of Social Service Research* 45.3 (2019), pp. 348–359.

[36] Patrick J. Fowler, Peter S. Hovmand, Katherine E. Marcal, and Sanmay Das. "Solving Homelessness from a Complex Systems Perspective: Insights for Prevention Responses." In: *Annual Review of Public Health* 40.1 (2019). PMID: 30601718, pp. 465–486. DOI: 10.1146/annurev-publhealth-040617-013553.

[37] Rachel Freedman, Jana Schaich Borg, Walter Sinnott-Armstrong, John P Dickerson, and Vincent Conitzer. "Adapting a kidney exchange algorithm to align with human values." In: *Artificial Intelligence* 283 (2020), p. 103261.

[38] Zachary Glendening and Marybeth Shinn. *Predicting repeated and persistent family homelessness: Do families' characteristics and experiences matter*. Tech. rep. OPRE Report, 2018.

[39] Zachary Glendening and Marybeth Shinn. "Risk models for returns to housing instability among families experiencing homelessness." In: *Cityscape (Washington, DC)* 19.3 (2017), p. 309.

[40] Ben Green and Yiling Chen. "Disparate Interactions: An Algorithm-in-the-Loop Analysis of Fairness in Risk Assessments." en. In: *Proceedings of the Conference on Fairness, Accountability, and Transparency - FAT* '19*. Atlanta, GA, USA: ACM Press, 2019, pp. 90–99. DOI: 10.1145/3287560.3287563.

[41] Daniel Gubits, Marybeth Shinn, Michelle Wood, Scott R. Brown, Samuel R. Dastrup, and Stephen H. Bell. "What Interventions Work Best for Families Who Experience Homelessness? Impact Estimates from the Family Options Study." en. In: *Journal of Policy Analysis and Management* 37.4 (2018), pp. 835–866. DOI: 10.1002/pam.22071.

[42] Jacqueline Hannan, Huei-Yen Winnie Chen, and Kenneth Joseph. "Who Gets What, According to Whom? An Analysis of Fairness Perceptions in Service Allocation." In: *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*. 2021, pp. 555–565.

[43] Meghan Henry, Anna Mahathey, Tyler Morrill, Anna Robinson, Azim Shivji, and Rian Watt. *The 2018 Annual Homeless Assessment Report (AHAR) to Congress: Part 1: Point-in-Time Estimates of Homelessness*. Tech. rep. Washington D.C.: United States Department of Housing and Urban Development, 2018, p. 100.

[44] Jennifer L Hill. "Bayesian nonparametric modeling for causal inference." In: *Journal of Computational and Graphical Statistics* 20.1 (2011), pp. 217–240.

[45] HUD. *Coordinated Entry Core Elements*. 2018.

[46] HUD. *Introductory Guide to the Continuum of Care (CoC) Program: Understanding the CoC Program and the requirements of the CoC Program Interim Rule*. 2012.

[47] *HUD Exchange Website*.

[48] Fredrik Johansson, Uri Shalit, and David Sontag. "Learning representations for counterfactual inference." In: *International Conference on Machine Learning*. 2016, pp. 3020–3029.

[49] 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.

[50] Jon Kleinberg, Himabindu Lakkaraju, Jure Leskovec, Jens Ludwig, and Sendhil Mullainathan. "Human decisions and machine predictions." In: *The Quarterly Journal of Economics* 133.1 (2018), pp. 237–293.

[51] Jon Kleinberg, Jens Ludwig, Sendhil Mullainathan, and Ziad Obermeyer. "Prediction policy problems." In: *The American Economic Review* 105.5 (2015), pp. 491–495.

[52] Jon Kleinberg, Sendhil Mullainathan, and Manish Raghavan. "Inherent trade-offs in the fair determination of risk scores." In: *arXiv preprint arXiv:1609.05807* (2016).

[53] Scott Duke Kominers, Alexander Teytelboym, and Vincent P Crawford. "An invitation to market design." In: *Oxford Review of Economic Policy* 33.4 (2017), pp. 541–571.

[54] Amanda Kube, Sanmay Das, and Patrick J. Fowler. "Allocating Interventions Based on Predicted Outcomes: A Case Study on Homelessness Services." en. In: vol. 33. July 2019, pp. 622–629. DOI: `10.1609/aaai.v33i01.3301622`.

[55] H. W. Kuhn. "The Hungarian method for the assignment problem." In: *Naval Research Logistics Quarterly* 2.1-2 (1955), pp. 83–97.

[56] Min Kyung Lee et al. "WeBuildAI: Participatory framework for algorithmic governance." In: *Proceedings of the ACM on Human-Computer Interaction* 3.CSCW (2019), pp. 1–35.

[57] Irwin Levin, Sandra Schneider, and Gary Gaeth. "All Frames Are Not Created Equal: A Typology and Critical Analysis of Framing Effects," in: *Organizational Behavior and Human Decision Processes* 76 (Feb. 1998), pp. 149–188. DOI: `10.1006/obhd.1998.2804`.

[58] Zhuoshu Li, Kelsey Lieberman, William Macke, Sofia Carrillo, Chien-Ju Ho, Jason Wellen, and Sanmay Das. "Incorporating Compatible Pairs in Kidney Exchange: A Dynamic Weighted Matching Model." In: *Proceedings of the ACM Conference on Economics and Computation*. 2019, pp. 349–367.

[59] Zhiyuan "Jerry" Lin, Jongbin Jung, Sharad Goel, and Jennifer Skeem. "The limits of human predictions of recidivism." In: *Science Advances* 6.7 (2020), eaaz0652.

[60] Tasfia Mashiat, Xavier Gitiaux, Huzefa Rangwala, Patrick J Fowler, and Sanmay Das. "Trade-offs between Group Fairness Metrics in Societal Resource Allocation." In: Upcoming.

[61] Steven Melendez. *How big data could help make it easier for landlords to evict you*. Sept. 2019.

[62] Stephen Metraux, Nicholas Eng, Jay Bainbridge, and Dennis P Culhane. "The impact of shelter use and housing placement on mortality hazard for unaccompanied adults and adults in family households entering New York City shelters: 1990–2002." In: *Journal of Urban Health* 88.6 (2011), pp. 1091–1104.

[63] Cathy O'Neil. *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. New York, NY: Broadway Books, 2016.

[64] Ziad Obermeyer. "Race representation matters in cancer care." In: *The Lancet Digital Health* 3 (2020), e408.

[65] Geoff Pleiss, Manish Raghavan, Felix Wu, Jon Kleinberg, and Kilian Q Weinberger. "On Fairness and Calibration." In: (2017), pp. 5684–5693.

[66] Anna Reosti. ""We Go Totally Subjective": Discretion, Discrimination, and Tenant Screening in a Landlord's Market." In: *Law and Social Inquiry* 45.3 (Aug. 2020), pp. 618–657.

[67] Jason M Rodriguez and Marybeth Shinn. "Intersections of family homelessness, CPS involvement, and race in Alameda County, California." In: *Child abuse & neglect* 57 (2016), pp. 41–52.

[68] Alvin E Roth, Tayfun Sonmez, and M Utku Unver. "Pairwise kidney exchange." In: *Journal of Economic Theory* 125.2 (2005), pp. 151–188.

[69] Alvin E. Roth and Marilda Sotomayor. *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*. Econometric Society Monograph Series. Cambridge UK: Cambridge University Press, 1990.

[70] Valerie Schneider. "Locked out by Big Data: How Big Data Algorithms and Machine Learning May Undermine Housing Justice." In: *Columbia Human Rights Law Review* 52 (2020), p. 251.

[71] Marybeth Shinn, Scott R. Brown, Brooke E. Spellman, Michelle Wood, Daniel Gubits, and Jill Khadduri. "Mismatch Between Homeless Families and the Homelessness Service System." In: *Cityscape (Washington, D.C.)* 19.3 (2017), pp. 293–307.

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

[73] Sarah Tan, Julius Adebayo, Kori Inkpen, and Ece Kamar. "Investigating human+ machine complementarity for recidivism predictions." In: *arXiv preprint arXiv:1808.09123* (2018).

[74] Andrew C Trapp, Alexander Teytelboym, Alessandro Martinello, Tommy Andersson, and Narges Ahani. "Placement optimization in refugee resettlement." In: *Operations Research* (2021). To appear.

[75] Dan Treglia, Eileen L Johns, Maryanne Schretzman, Jacob Berman, Dennis P Culhane, David C Lee, and Kelly M Doran. "When crises converge: hospital visits before and after shelter use among homeless New Yorkers." In: *Health Affairs* 38.9 (2019), pp. 1458–1467.

[76] Amos Tversky and Daniel Kahneman. "Judgment under Uncertainty: Heuristics and Biases: Biases in judgments reveal some heuristics of thinking under uncertainty." In: *Science* 185.4157 (1974), pp. 1124–1131.

[77] Amos Tversky and Daniel Kahneman. "The Framing of Decisions and the Psychology of Choice." In: *Science* 211.4481 (1981), pp. 453–458. DOI: 10.1126/science.7455683.

[78] United States Congress. *Homeless Emergency Assistance and Rapid Transition to Housing Act*. 2009.

[79] Michelle Vaccaro and Jim Waldo. "The effects of mixing machine learning and human judgment." In: *Communications of the Association for Computing Machinery* 62.11 (2019), pp. 104–110.

[80] Phebe Vayanos, Duncan McElfresh, Yingxiao Ye, John Dickerson, and Eric Rice. "Active preference elicitation via adjustable robust optimization." In: *arXiv preprint arXiv:2003.01899* (2020).

[81] Till Von Wachter, Marianne Bertrand, Harold Pollack, Janey Rountree, and Brian Blackwell. "Predicting and preventing homelessness in Los Angeles." In: *California Policy Lab and University of Chicago Poverty Lab* (2019).

[82] *Vulnerability Index - Service Prioritization Decision Assistance Tool (VI-SPDAT): Prescreen Triage Tool for Families*. 2015.

[83] David Weinberger. "Don't Make Artificial Intelligence Artificially Stupid in the Name of Transparency." In: *Wired* (Jan. 2018).

[84] 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 International Conference on Autonomous Agents and Multiagent Systems*. 2016, pp. 740–748.

Appendix A

Descriptive Statistics

Table A.1 provides descriptive statistics for all features in the dataset.

Feature	**Emergency Shelter** n=2997	**Transitional Housing** n=1469	**Rapid Rehousing** n=840	**Homelessness Prevention** n=4737	**Total** n=10043
	% / M	% / M	% / M	% / M	% / M
Household Characteristics					
Number of Household Members	1.75	1.15	1.67	2.44	1.98
Spouse Present	1.90	0.20	4.52	10.05	5.72
Number of Children	0.73	0.14	0.60	1.25	0.88
Number of Children Ages 0 to 2	0.26	0.08	0.09	0.15	0.17
Number of Children Ages 3 to 5	0.17	0.03	0.11	0.18	0.15
Number of Children Ages 6 to 10	0.17	0.02	0.15	0.31	0.21
Number of Children Ages 11 to 14	0.09	0.01	0.11	0.24	0.15
Number of Children Ages 15 to 17	0.03	0.00	0.07	0.17	0.10
Number of Unrelated Adults	0.00	0.00	0.03	0.07	0.04
Number of Unrelated Children	0.01	0.00	0.03	0.11	0.06
Number of Calls Before Entry	4.52	4.14	3.13	1.07	2.72
Wait Before Entry (in days)	267.79	276.30	291.83	176.74	228.10
Head of Household Characteristics					
Female	75.54	16.20	52.14	82.75	68.61
Age (years)	36.38	38.55	43.43	41.13	39.53
White	18.18	24.78	17.62	8.76	14.66
African American	79.85	73.66	81.43	90.10	83.91
Hispanic or Latino Ethnicity	1.60	1.36	0.71	0.76	1.10
Veteran	3.14	8.51	7.02	2.48	3.92
Disabling Condition	17.25	24.30	19.88	9.90	15.04
Physical Disability	20.49	15.04	26.43	18.85	19.42
Received Physical Disability Services	7.17	6.54	8.93	8.09	7.66
Developmental Disability	3.47	2.25	4.76	1.90	2.66
Received Developmental Disability Services	0.40	0.27	1.67	0.42	0.50
Chronic Health Condition	33.47	29.07	35.12	36.90	34.58
Received Chronic Health Services	17.65	14.43	21.43	22.04	19.57
HIV/AIDS	0.70	0.48	0.95	0.53	0.61
Received HIV/AIDS Services	0.37	0.20	0.83	0.13	0.27
Mental Health Problem	34.57	25.05	37.62	26.89	29.81
Received Mental Health Services	12.85	11.71	15.71	10.83	11.97
Alcohol Abuse Problem	6.14	11.37	5.48	3.95	5.81
Drug Abuse Problem	13.71	24.10	11.55	10.39	13.48
Both Alcohol and Drug Abuse Problem	9.64	15.32	8.81	4.71	8.08
Received Substance Abuse Services	10.08	32.40	9.52	10.32	13.41
Domestic Violence Survivor	1.23	0.48	1.07	0.61	0.82
Chronically Homeless	0.27	2.86	17.98	1.03	2.49

Table A.1: Summary of the dataset by service type

Appendix B

Extended Comparison of Ineligible Households

Tables B.1 is an extended version of Table 5.3 containing all features in the dataset.

Feature	Transitional Housing n=3349	Rapid Rehousing n=1537	All Prevention Ineligible n=4942	Prevention Eligible n=5101
	% / M	% / M	% / M	% / M
Household Characteristics				
Number of Household Members	1.14(0.56)	2.29(1.43)	1.56(1.11)	2.41(1.49)
Spouse Present	0.87	29.80	1.73	9.83
Number of Children	0.13(0.54)	1.26(1.40)	0.54(1.08)	1.23(1.38)
Number of Children Ages 0 to 2	0.05(0.25)	0.45(0.64)	0.18(0.47)	0.15(0.40)
Number of Children Ages 3 to 5	0.03(0.18)	0.29(0.57)	0.12(0.39)	0.18(0.46)
Number of Children Ages 6 to 10	0.03(0.20)	0.29(0.63)	0.13(0.43)	0.30(0.63)
Number of Children Ages 11 to 14	0.01(0.14)	0.15(0.47)	0.07(0.32)	0.23(0.52)
Number of Children Ages 15 to 17	0.01(0.14)	0.05(0.23)	0.03(0.18)	0.17(0.43)
Number of Unrelated Adults	0.01(0.09)	0.01(0.10)	0.01(0.11)	0.07(0.29)
Number of Unrelated Children	0.00(0.08)	0.02(0.22)	0.01(0.16)	0.11(0.47)
Number of Calls Before Entry	3.94(6.04)	4.29(6.95)	4.28(6.83)	1.11(2.72)
Wait Before Entry (in days)	285.13(490.01)	241.79(447.13)	276.25(480.22)	178.40(412.47)
Monthly Income (in US Dollars)	642.92(886.20)	953.91(976.33)	766.70(976.55)	2017.72(2193.52)
Head of Household Characteristics				
Female	36.97	90.24	54.21	82.86
Age (years)	40.52(12.89)	33.04(12.24)	38.11(13.07)	40.99(0.07)
Head of Household Between 18 and 24 years	14.30	27.59	18.07	7.08
White	18.90	22.51	20.07	9.07
African American	79.67	75.02	78.18	89.82
Hispanic or Latino Ethnicity	1.07	2.15	1.39	0.79
Veteran	7.29	1.89	5.47	2.33
Disabling Condition	22.04	17.05	20.19	9.71
Physical Disability	21.44	15.48	19.47	19.36
Received Physical Disability Services	10.21	1.95	7.39	7.93
Developmental Disability	3.34	4.03	3.43	1.86
Received Developmental Disability Services	0.57	0.65	0.57	0.42
Chronic Health Condition	31.74	32.53	33.13	36.08
Received Chronic Health Services	18.04	14.83	17.74	21.45
HIV/AIDS	0.90	0.33	0.71	0.51
Received HIV/AIDS Services	0.48	0.26	0.41	0.12
Mental Health Problem	31.95	34.42	32.74	26.79
Received Mental Health Services	14.12	11.26	13.35	10.54
Alcohol Abuse Problem	9.82	3.12	7.72	3.84
Drug Abuse Problem	21.08	7.42	16.60	10.26
Both Alcohol and Drug Abuse Problem	15.02	3.90	11.39	4.65
Received Substance Abuse Services	23.29	3.51	16.60	10.12
Domestic Violence Survivor	0.84	1.24	1.02	0.61
Chronically Homeless	5.26	1.11	3.86	1.07

Table B.1: Comparison of the characteristics of homelessness prevention ineligible households that are predicted to do best in rapid rehousing and transitional housing

Appendix C

Score Distributions For the Type-Elicitation Task

Figure C.1 provides histograms of *Vulnerability-Oriented* and *Outcome-Oriented* scores for participants randomized to make *Vulnerability-Oriented* decisions in the **type-elicitation** task. These scores show that most participants understood the task they were given and were able to make a high proportion of *Vulnerability-Oriented* decisions.

Figure C.2 provides the same histograms for participants randomized to make *Outcome-Oriented* decisions in the **type-elicitation** task. Here we see participants were able to make a high proportions of *Outcome-Oriented* decisions.

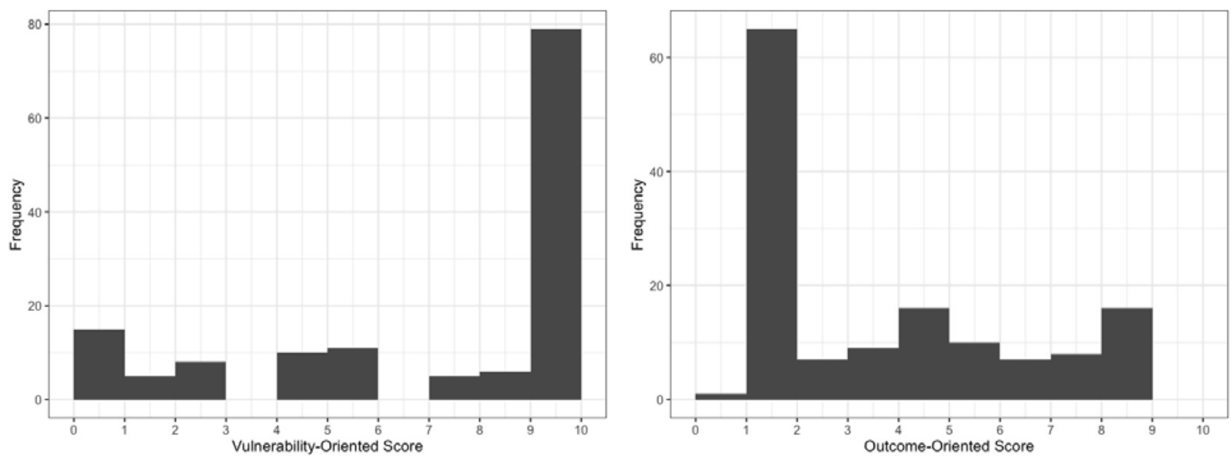

(a) Histogram of *Vulnerability-Oriented* Scores (b) Histogram of *Outcome-Oriented* Scores

Figure C.1: Histograms showing the distributions of *Vulnerability-Oriented* and *Outcome-Oriented* scores for the group told to make neediest-first decisions

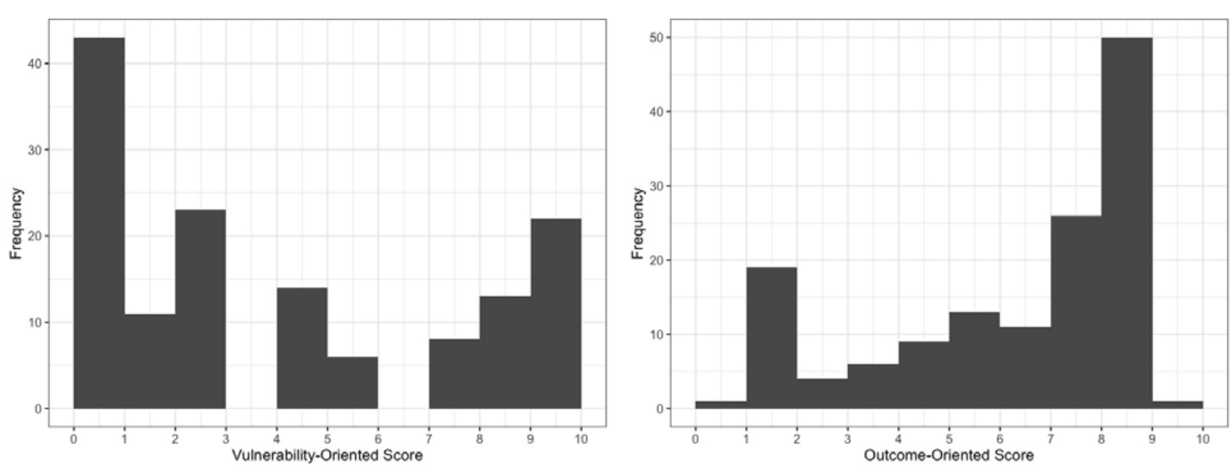

(a) Histogram of *Vulnerability-Oriented* Scores (b) Histogram of *Outcome-Oriented* Scores

Figure C.2: Histograms showing the distributions of *Vulnerability-Oriented* and *Outcome-Oriented* scores for the group told to make *Outcome-Oriented* decisions

Appendix D

Robustness Check

Figure D.1 is an alternate version of Figure 6.6 where type is defined as *Vulnerability-Oriented* if the participant had a higher *Vulnerability-Oriented* score than *Outcome-Oriented* score and *Outcome-Oriented* otherwise. In this case, no participant is considered to not have a type. We see much the same main result with this type definition as the type definition described in Section 3.1.

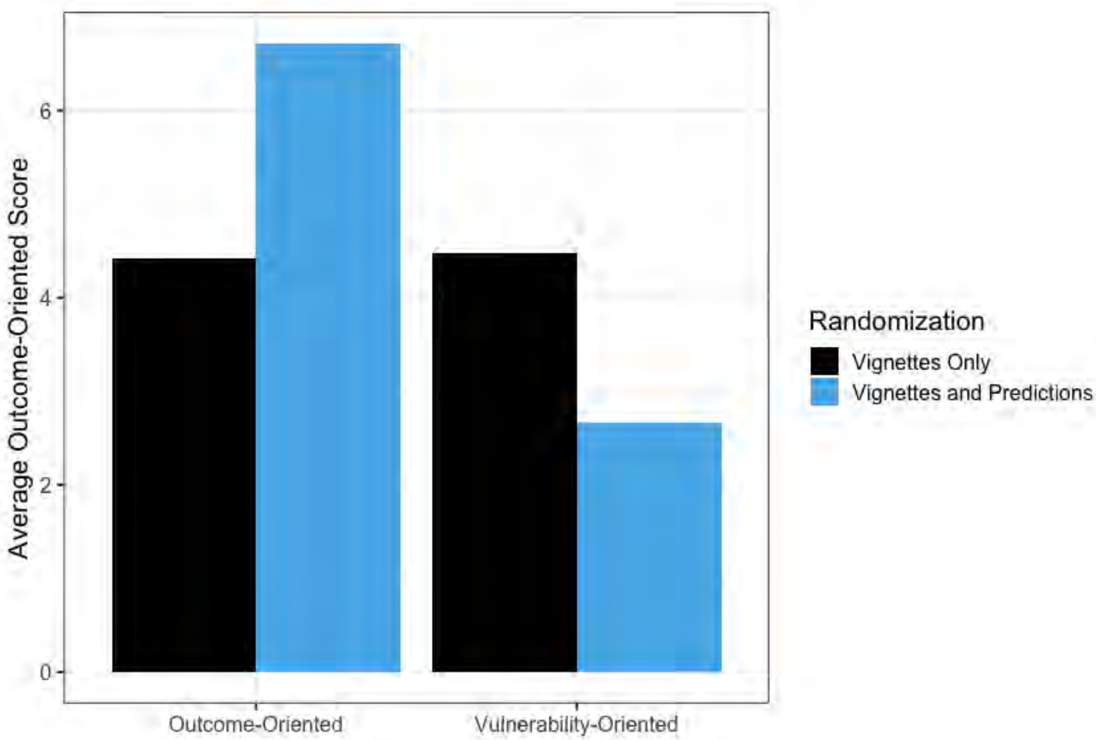

Figure D.1: Barplot comparing the average *Outcome-Oriented* score for each prioritization group across vignette only and vignette and prediction groups where types are defined by which criterion the participant scored highest on

Appendix E

Data Collection Instrument

The following pages contain a copy of instructions and tasks participants saw in our online survey.

Consent

Overview

Thank you for participating in research conducted by investigators from Washington University in St. Louis. The survey investigates the decision-making processes involved when allocating scarce social resources. We set the context in homeless service delivery that provide limited housing resources for households in precarious living accommodations. You will be asked to review descriptions of households seeking supports, and then, assign them to one of two potential homeless services. The survey takes about 20 minutes and is funded by the National Science Foundation award number 1939677 and Amazon.

Goal

The main goal of the survey is to better understand the ways information is used in making decisions on how to allocate scarce social resources. By providing descriptions of households seeking homeless services and asking for an assignment, we observe what decisions are made with different information. Survey data allow us to compare human decisions to those made by a computer. The study aims to improve decision-making on scarce resources for homeless service delivery. Your participation contributes to advancing our understanding.

Procedures

By continuing to the survey, you are volunteering to participate in the study. You will be introduced to the context of homeless services and the tasks involved, and then, will be asked to make a series of decisions to assign households to services. The survey is voluntary; you may stop at anytime by closing the browser. If you decide not to take part in the study or if you stop participating at any time, you won't be penalized or lose any benefits for which you otherwise qualify.

Risks and costs

We will keep the information you provide confidential. All responses will remain anonymous, and reports of study findings will not include information that identifies you.

Benefits

Although you will not gain personally, we hope that others may benefit in the future from what we learn as a result of this study. You will receive compensation for your time through the Mechanical Turk compensation platform. You will receive $3 for your participation.

Confidentiality

Again, we will keep the information you provide confidential. All responses will remain anonymous, and reports of study findings will not include information that identifies you.

Participant Certification

I have read and understand the study description. I understand the purpose of the research project and what I will be asked to do. I may stop my participation in this research study at any time and that I can refuse to answer any question(s). I hereby give my informed and free consent to be a participant in this study.

Instructions

Communities across the US provide homeless services that respond to household requests for assistance in securing stable housing. Households call a hotline to request assistance and provide basic demographic information, including household size, monthly income, whether anyone receives disability supports, and their last residence. Homeless service providers must decide what services to offer households based on the need and availability of resources.

Two key services include the following:
1) **Emergency Shelter** provides an immediate response to homelessness;
2) **Transitional Housing** provides long-term housing as well as individual case management which can include treatment for disabilities or health conditions.

Emergency Shelter consumes fewer resources to provide. Transitional Housing requires additional supports, and thus, is more scarce. Budget constraints do not allow all households to receive Transitional Housing. Those who cannot have access to Transitional Housing often receive Emergency Shelter or stay in an Emergency Shelter until space becomes available in Transitional Housing or another service.

One key outcome of interest to homeless service providers and researchers is whether receiving services reduces a household's need for services again in the future. Using the demographic information provided during hotline calls, researchers have developed models to predict the probability that a household will need services again within 2 years of being allocated a service. All households are assumed to have lower or equal probabilities of needing future services if given Transitional Housing than if given Emergency Shelter.

During the following activities you will be asked to see the same information that service providers receive and make decisions about how to allocate homeless services based on that information.

Training Block

Homeless service providers make decisions on which services to allocate to which households. These decisions are complex and based on complicated patterns of information about each household. Throughout this survey, you will be given information similar to what homeless service providers see when allocating services, and you will be asked to make decisions based on that information.

We will start with an example similar to what you will see in future tasks. You will see a total of 10 example households presented together in a table. As you proceed through this section, different households will be highlighted to show an estimate based on past administrative data of the likelihood of those households needing future services if given Transitional Housing or Emergency Shelter. Please read through these examples and estimates and spend some time looking for patterns in the data. Click "I Understand" when you are ready to proceed.

Here is a table of 10 examples of households needing homeless services. Each column in the table represents a different household and each row represents a piece of information that homeless service providers have when making decisions on which households should receive each kind of service. The next few screens will show the same table with the same ten households but will highlight the columns associated with households that have a certain predicted probability of needing future services if given either Transitional Housing or Emergency Shelter. Take a few minutes to familiarize yourself with each table and look for patterns in the information to help you categorize similar households later in the survey. Then, click "I Understand" when you are ready to proceed.

	Household 1	Household 2	Household 3	Household 4	Household 5	Household 6	Household 7	Household 8	Household 9	Household 10
Number of Household Members	1	1	1	1	1	1	1	2	1	1
Number of Children	0	0	0	0	0	0	0	1	0	0
Head of Household Gender	Female	Male	Male	Male	Female	Female	Male	Female	Female	Female
Head of Household Age	45	35	64	55	55	23	23	23	55	45
Head of Household Disabling Condition	No	Yes	Yes	No	Yes	No	No	No	No	No
Head of Household Receives Substance Abuse Services	No	No	No	No	No	No	Yes	No	No	No
Monthly Income	$800	$668	$1524	$400	$800	$367	$0	$1068	$800	$367
Number of Calls to Hotline	15	5	2	2	2	5	0	0	2	0
Prior Residence	Private Emergency Shelter	Private Emergency Shelter	Private Emergency Shelter	Private Emergency Shelter	Private Emergency Shelter	Private Transitional Housing	Private Emergency Shelter	Private Transitional Housing	Private Transitional Housing	Private Emergency Shelter

This is the same table presented previously. Models predict that the households highlighted in red have a **high** probability of needing services again within 2 years if they are given **Transitional Housing**.

	Household 1	Household 2	Household 3	Household 4	Household 5	Household 6	Household 7	Household 8	Household 9	Household 10
Number of Household Members	1	1	1	1	1	1	1	2	1	1
Number of Children	0	0	0	0	0	0	0	1	0	0
Head of Household Gender	Female	Male	Male	Male	Female	Female	Male	Female	Female	Female
Head of Household Age	45	35	64	55	55	23	23	23	55	45
Head of Household Disabling Condition	No	Yes	Yes	No	Yes	No	No	No	No	No
Head of Household Receives Substance Abuse Services	No	No	No	No	No	No	Yes	No	No	No
Monthly Income	$800	$668	$1524	$400	$800	$367	$0	$1068	$800	$367
Number of Calls to Hotline	15	5	2	2	2	5	0	0	2	0
Prior Residence	Private Emergency Shelter	Private Emergency Shelter	Private Emergency Shelter	Private Emergency Shelter	Private Emergency Shelter	Private Transitional Housing	Private Emergency Shelter	Private Transitional Housing	Private Transitional Housing	Private Emergency Shelter

This is the same table presented previously. Models predict that the households highlighted in yellow have a **medium** probability of needing services again within 2 years if they are given **Transitional Housing**.

	Household 1	Household 2	Household 3	Household 4	Household 5	Household 6	Household 7	Household 8	Household 9	Household 10
Number of Household Members	1	1	1	1	1	1	1	2	1	1
Number of Children	0	0	0	0	0	0	0	1	0	0
Head of Household Gender	Female	Male	Male	Male	Female	Female	Male	Female	Female	Female
Head of Household Age	45	35	64	55	55	23	23	23	55	45
Head of Household Disabling Condition	No	Yes	Yes	No	Yes	No	No	No	No	No
Head of Household Receives Substance Abuse Services	No	No	No	No	No	No	Yes	No	No	No
Monthly Income	$800	$668	$1524	$400	$800	$367	$0	$1068	$800	$367
Number of Calls to Hotline	15	5	2	2	2	5	0	0	2	0
Prior Residence	Private Emergency Shelter	Private Emergency Shelter	Private Emergency Shelter	Private Emergency Shelter	Private Emergency Shelter	Private Transitional Housing	Private Emergency Shelter	Private Transitional Housing	Private Transitional Housing	Private Emergency Shelter

This is the same table presented previously. Models predict that the households highlighted in green have a **low** probability of needing services again within 2 years if they are given **Transitional Housing**.

	Household 1	Household 2	Household 3	Household 4	Household 5	Household 6	Household 7	Household 8	Household 9	Household 10
Number of Household Members	1	1	1	1	1	1	1	2	1	1
Number of Children	0	0	0	0	0	0	0	1	0	0
Head of Household Gender	Female	Male	Male	Male	Female	Female	Male	Female	Female	Female
Head of Household Age	45	35	64	55	55	23	23	23	55	45
Head of Household Disabling Condition	No	Yes	Yes	No	Yes	No	No	No	No	No
Head of Household Receives Substance Abuse Services	No	No	No	No	No	No	Yes	No	No	No
Monthly Income	$800	$668	$1524	$400	$800	$367	$0	$1068	$800	$367
Number of Calls to Hotline	15	5	2	2	2	5	0	0	2	0
Prior Residence	Private Emergency Shelter	Private Emergency Shelter	Private Emergency Shelter	Private Emergency Shelter	Private Emergency Shelter	Private Transitional Housing	Private Emergency Shelter	Private Transitional Housing	Private Transitional Housing	Private Emergency Shelter

This is the same table presented previously. Models predict that the households highlighted in red have a **high** probability of needing services again within 2 years if they are given **Emergency Shelter**.

	Household 1	Household 2	Household 3	Household 4	Household 5	Household 6	Household 7	Household 8	Household 9	Household 10
Number of Household Members	1	1	1	1	1	1	1	2	1	1
Number of Children	0	0	0	0	0	0	0	1	0	0
Head of Household Gender	Female	Male	Male	Male	Female	Female	Male	Female	Female	Female
Head of Household Age	45	35	64	55	55	23	23	23	55	45
Head of Household Disabling Condition	No	Yes	Yes	No	Yes	No	No	No	No	No
Head of Household Receives Substance Abuse Services	No	No	No	No	No	No	Yes	No	No	No
Monthly Income	$800	$668	$1524	$400	$800	$367	$0	$1068	$800	$367
Number of Calls to Hotline	15	5	2	2	2	5	0	0	2	0
Prior Residence	Private Emergency Shelter	Private Emergency Shelter	Private Emergency Shelter	Private Emergency Shelter	Private Emergency Shelter	Private Transitional Housing	Private Emergency Shelter	Private Transitional Housing	Private Transitional Housing	Private Emergency Shelter

This is the same table presented previously. Models predict that the households highlighted in yellow are predicted to have a **medium** probability of needing services again within 2 years if they are given **Emergency Shelter**.

	Household 1	Household 2	Household 3	Household 4	Household 5	Household 6	Household 7	Household 8	Household 9	Household 10
Number of Household Members	1	1	1	1	1	1	1	2	1	1
Number of Children	0	0	0	0	0	0	0	1	0	0
Head of Household Gender	Female	Male	Male	Male	Female	Female	Male	Female	Female	Female
Head of Household Age	45	35	64	55	55	23	23	23	55	45
Head of Household Disabling Condition	No	Yes	Yes	No	Yes	No	No	No	No	No
Head of Household Receives Substance Abuse Services	No	No	No	No	No	No	Yes	No	No	No
Monthly Income	$800	$668	$1524	$400	$800	$367	$0	$1068	$800	$367
Number of Calls to Hotline	15	5	2	2	2	5	0	0	2	0
Prior Residence	Private Emergency Shelter	Private Emergency Shelter	Private Emergency Shelter	Private Emergency Shelter	Private Emergency Shelter	Private Transitional Housing	Private Emergency Shelter	Private Transitional Housing	Private Transitional Housing	Private Emergency Shelter

This is the same table presented previously. Models predict that the households highlighted in green have a **low** probability of needing services again within 2 years if they are given **Emergency Shelter**.

	Household 1	Household 2	Household 3	Household 4	Household 5	Household 6	Household 7	Household 8	Household 9	Household 10
Number of Household Members	1	1	1	1	1	1	1	2	1	1
Number of Children	0	0	0	0	0	0	0	1	0	0
Head of Household Gender	Female	Male	Male	Male	Female	Female	Male	Female	Female	Female
Head of Household Age	45	35	64	55	55	23	23	23	55	45
Head of Household Disabling Condition	No	Yes	Yes	No	Yes	No	No	No	No	No
Head of Household Receives Substance Abuse Services	No	No	No	No	No	No	Yes	No	No	No
Monthly Income	$800	$668	$1524	$400	$800	$367	$0	$1068	$800	$367
Number of Calls to Hotline	15	5	2	2	2	5	0	0	2	0
Prior Residence	Private Emergency Shelter	Private Emergency Shelter	Private Emergency Shelter	Private Emergency Shelter	Private Emergency Shelter	Private Transitional Housing	Private Emergency Shelter	Private Transitional Housing	Private Transitional Housing	Private Emergency Shelter

No Training Block

Homeless service providers make decisions on which services to allocate to which households. These decisions are complex and based on complicated patterns of information about each household. Throughout this survey, you will be given information similar to what homeless service providers see when allocating services, and you will be asked to make decisions based on that information.

We will start with examples similar to what you will see in future tasks. You will see a total of 10 examples presented together in a table. Read through the examples and try to think of the likelihood of those households needing future services if given Transitional Housing or Emergency Shelter. Click "I Understand" when you are ready to proceed.

Here is a table of 10 examples of households needing homeless services. Each column in the table represents a different household and each row represents a piece of information that homeless service providers have when making decisions on which households should receive each kind of service. Take a few minutes to familiarize yourself with the table and try to think of the likelihood of those households needing future services if given Transitional Housing or Emergency Shelter. Then, click "I Understand" when you are ready to proceed.

	Household 1	Household 2	Household 3	Household 4	Household 5	Household 6	Household 7	Household 8	Household 9	Household 10
Number of Household Members	1	1	1	1	1	1	1	2	1	1
Number of Children	0	0	0	0	0	0	0	1	0	0
Head of Household Gender	Female	Male	Male	Male	Female	Female	Male	Female	Female	Female
Head of Household Age	45	35	64	55	55	23	23	23	55	45
Head of Household Disabling Condition	No	Yes	Yes	No	Yes	No	No	No	No	No
Head of Household Receives Substance Abuse Services	No	No	No	No	No	No	Yes	No	No	No
Monthly Income	$800	$668	$1524	$400	$800	$367	$0	$1068	$800	$367
Number of Calls to Hotline	15	5	2	2	2	5	0	0	2	0
Prior Residence	Private Emergency Shelter	Private Emergency Shelter	Private Emergency Shelter	Private Emergency Shelter	Private Emergency Shelter	Private Transitional Housing	Private Emergency Shelter	Private Transitional Housing	Private Transitional Housing	Private Emergency Shelter

Instructions Task 1

Next, we will present tables with descriptions of households seeking homeless services. Please use these descriptions to predict the probability that a household will need further services within 2 years if placed in Transitional Housing.
You will be presented with 10 households to sort into 3 groups: Low predicted probability, Medium predicted probability, and High predicted probability.

Task 1

Categorize the following households based on how likely you predict they are to need further services within 2 years if placed into Transitional Housing. Click on the household number label below the image to select your answer.

	Household 1	Household 2	Household 3	Household 4	Household 5	Household 6	Household 7	Household 8	Household 9	Household 10
Number of Household Members	1	1	1	1	1	1	1	1	1	1
Number of Children	0	0	0	0	0	0	0	0	0	0
Head of Household Gender	Male	Male	Female	Male	Female	Male	Male	Male	Male	Male
Head of Household Age	35	45	35	45	45	45	55	45	45	35
Head of Household Disabling Condition	Yes	Yes	No	No	No	Yes	No	Yes	No	No
Head of Household Receives Substance Abuse Services	Yes	No	No	No	Yes	Yes	No	Yes	Yes	No
Monthly Income	$400	$1524	$668	$200	$800	$668	$0	$0	$200	$0
Number of Calls to Hotline	5	1	15	25	1	3	25	5	2	25
Prior Residence	Private Emergency Shelter	Place not meant for habitation	Permanent housing for formerly homeless persons	Private Emergency Shelter	Place not meant for habitation	Private Emergency Shelter	Private Emergency Shelter	Place not meant for habitation	Private Emergency Shelter	Private Emergency Shelter

Household 1 Household 2 Household 3 Household 4 Household 5
Household 6 Household 7 Household 8 Household 9 Household 10

Instructions Task 2

Now that you understand, we will present descriptions of households seeking homeless services in pairs. Please decide which household to prioritize for **Transitional Housing**.
You will be presented with 12 pairs of households.

Task 2 - Vignette Only Group

Which of the following households would you prioritize for Transitional Housing? Select your choice by clicking on the associated column in the image below.

	Household 1	Household 2
Number of Household Members	1	1
Number of Children	0	0
Head of Household Gender	Female	Female
Head of Household Age	35	45
Head of Household Disabling Condition	No	No
Head of Household Receives Substance Abuse Services	No	Yes
Monthly Income	$668	$800
Number of Calls to Hotline	15	1
Prior Residence	Permanent housing for formerly homeless persons	Place not meant for habitation

Which of the following households would you prioritize for Transitional Housing? Select your choice by clicking on the associated column in the image below.

	Household 1	Household 2
Number of Household Members	1	1
Number of Children	0	0
Head of Household Gender	Female	Female
Head of Household Age	35	45
Head of Household Disabling Condition	No	No
Head of Household Receives Substance Abuse Services	No	Yes
Monthly Income	$668	$800
Number of Calls to Hotline	15	1
Prior Residence	Permanent housing for formerly homeless persons	Place not meant for habitation

Which of the following households would you prioritize for Transitional Housing? Select your choice by clicking on the associated column in the image below.

	Household 1	Household 2
Number of Household Members	1	1
Number of Children	0	0
Head of Household Gender	Male	Male
Head of Household Age	55	45
Head of Household Disabling Condition	No	Yes
Head of Household Receives Substance Abuse Services	No	Yes
Monthly Income	$0	$668
Number of Calls to Hotline	25	3
Prior Residence	Private Emergency Shelter	Private Emergency Shelter

Which of the following households would you prioritize for Transitional Housing? Select your choice by clicking on the associated column in the image below.

	Household 1	Household 2
Number of Household Members	1	1
Number of Children	0	0
Head of Household Gender	Male	Male
Head of Household Age	45	45
Head of Household Disabling Condition	No	Yes
Head of Household Receives Substance Abuse Services	No	Yes
Monthly Income	$200	$0
Number of Calls to Hotline	25	5
Prior Residence	Private Emergency Shelter	Place not meant for habitation

Which of the following households would you prioritize for Transitional Housing? Select your choice by clicking on the associated column in the image below.

	Household 1	Household 2
Number of Household Members	1	1
Number of Children	0	0
Head of Household Gender	Male	Male
Head of Household Age	45	45
Head of Household Disabling Condition	No	Yes
Head of Household Receives Substance Abuse Services	No	Yes
Monthly Income	$200	$0
Number of Calls to Hotline	25	5
Prior Residence	Private Emergency Shelter	Place not meant for habitation

Which of the following households would you prioritize for Transitional Housing? Select your choice by clicking on the associated column in the image below.

	Household 1	Household 2
Number of Household Members	1	1
Number of Children	0	0
Head of Household Gender	Male	Male
Head of Household Age	45	35
Head of Household Disabling Condition	No	No
Head of Household Receives Substance Abuse Services	Yes	No
Monthly Income	$200	$0
Number of Calls to Hotline	2	25
Prior Residence	Private Emergency Shelter	Private Emergency Shelter

Which of the following households would you prioritize for Transitional Housing? Select your choice by clicking on the associated column in the image below.

	Household 1	Household 2
Number of Household Members	1	3
Number of Children	0	2
Head of Household Gender	Female	Female
Head of Household Age	45	35
Head of Household Disabling Condition	No	Yes
Head of Household Receives Substance Abuse Services	No	Yes
Monthly Income	$400	$1524
Number of Calls to Hotline	0	25
Prior Residence	Place not meant for habitation	Unknown

Which of the following households would you prioritize for Transitional Housing? Select your choice by clicking on the associated column in the image below.

	Household 1	Household 2
Number of Household Members	3	1
Number of Children	2	0
Head of Household Gender	Female	Female
Head of Household Age	23	35
Head of Household Disabling Condition	No	No
Head of Household Receives Substance Abuse Services	No	No
Monthly Income	$200	$0
Number of Calls to Hotline	5	0
Prior Residence	Private Emergency Shelter	Unknown

Which of the following households would you prioritize for Transitional Housing? Select your choice by clicking on the associated column in the image below.

	Household 1	Household 2
Number of Household Members	3	1
Number of Children	2	0
Head of Household Gender	Female	Male
Head of Household Age	23	55
Head of Household Disabling Condition	No	Yes
Head of Household Receives Substance Abuse Services	No	Yes
Monthly Income	$1524	$0
Number of Calls to Hotline	3	2
Prior Residence	Private Emergency Shelter	Private Emergency Shelter

Which of the following households would you prioritize for Transitional Housing? Select your choice by clicking on the associated column in the image below.

	Household 1	Household 2
Number of Household Members	2	1
Number of Children	1	0
Head of Household Gender	Female	Female
Head of Household Age	23	45
Head of Household Disabling Condition	No	No
Head of Household Receives Substance Abuse Services	No	Yes
Monthly Income	$3256	$200
Number of Calls to Hotline	25	0
Prior Residence	Unknown	Private Emergency Shelter

Which of the following households would you prioritize for Transitional Housing? Select your choice by clicking on the associated column in the image below.

	Household 1	Household 2
Number of Household Members	1	1
Number of Children	0	0
Head of Household Gender	Female	Male
Head of Household Age	55	35
Head of Household Disabling Condition	Yes	No
Head of Household Receives Substance Abuse Services	No	Yes
Monthly Income	$0	$0
Number of Calls to Hotline	5	2
Prior Residence	Private Emergency Shelter	Private Emergency Shelter

Which of the following households would you prioritize for Transitional Housing? Select your choice by clicking on the associated column in the image below.

	Household 1	Household 2
Number of Household Members	5	1
Number of Children	4	0
Head of Household Gender	Female	Female
Head of Household Age	23	23
Head of Household Disabling Condition	No	No
Head of Household Receives Substance Abuse Services	No	No
Monthly Income	$668	$32
Number of Calls to Hotline	25	0
Prior Residence	Private Emergency Shelter	Private Transitional Housing

Task 2 - Vignette and Predictions Group

Which of the following households would you prioritize for Transitional Housing? Select your choice by clicking on the associated column in the image below.

	Household 1	Household 2
Number of Household Members	1	1
Number of Children	0	0
Head of Household Gender	Female	Female
Head of Household Age	35	45
Head of Household Disabling Condition	No	No
Head of Household Receives Substance Abuse Services	No	Yes
Monthly Income	$668	$800
Number of Calls to Hotline	15	1
Prior Residence	Permanent housing for formerly homeless persons	Place not meant for habitation
Predicted probability of needing future services within 2 years if given Transitional Housing	High	Low
Predicted probability of needing future services within 2 years if given Emergency Shelter	High	Medium

Which of the following households would you prioritize for Transitional Housing? Select your choice by clicking on the associated column in the image below.

	Household 1	Household 2
Number of Household Members	1	1
Number of Children	0	0
Head of Household Gender	Female	Female
Head of Household Age	35	45
Head of Household Disabling Condition	No	No
Head of Household Receives Substance Abuse Services	No	Yes
Monthly Income	$668	$800
Number of Calls to Hotline	15	1
Prior Residence	Permanent housing for formerly homeless persons	Place not meant for habitation
Predicted probability of needing future services within 2 years if given Transitional Housing	High	Low
Predicted probability of needing future services within 2 years if given Emergency Shelter	High	Medium

Which of the following households would you prioritize for Transitional Housing? Select your choice by clicking on the associated column in the image below.

	Household 1	Household 2
Number of Household Members	1	1
Number of Children	0	0
Head of Household Gender	Male	Male
Head of Household Age	55	45
Head of Household Disabling Condition	No	Yes
Head of Household Receives Substance Abuse Services	No	Yes
Monthly Income	$0	$668
Number of Calls to Hotline	25	3
Prior Residence	Private Emergency Shelter	Private Emergency Shelter
Predicted probability of needing future services within 2 years if given Transitional Housing	High	Low
Predicted probability of needing future services within 2 years if given Emergency Shelter	High	Medium

Which of the following households would you prioritize for Transitional Housing? Select your choice by clicking on the associated column in the image below.

	Household 1	Household 2
Number of Household Members	1	1
Number of Children	0	0
Head of Household Gender	Male	Male
Head of Household Age	45	45
Head of Household Disabling Condition	No	Yes
Head of Household Receives Substance Abuse Services	No	Yes
Monthly Income	$200	$0
Number of Calls to Hotline	25	5
Prior Residence	Private Emergency Shelter	Place not meant for habitation
Predicted probability of needing future services within 2 years if given Transitional Housing	High	Low
Predicted probability of needing future services within 2 years if given Emergency Shelter	High	Medium

Which of the following households would you prioritize for Transitional Housing? Select your choice by clicking on the associated column in the image below.

	Household 1	Household 2
Number of Household Members	1	1
Number of Children	0	0
Head of Household Gender	Male	Male
Head of Household Age	45	45
Head of Household Disabling Condition	No	Yes
Head of Household Receives Substance Abuse Services	No	Yes
Monthly Income	$200	$0
Number of Calls to Hotline	25	5
Prior Residence	Private Emergency Shelter	Place not meant for habitation
Predicted probability of needing future services within 2 years if given Transitional Housing	High	Low
Predicted probability of needing future services within 2 years if given Emergency Shelter	High	Medium

Which of the following households would you prioritize for Transitional Housing? Select your choice by clicking on the associated column in the image below.

	Household 1	**Household 2**
Number of Household Members	1	1
Number of Children	0	0
Head of Household Gender	Male	Male
Head of Household Age	45	35
Head of Household Disabling Condition	No	No
Head of Household Receives Substance Abuse Services	Yes	No
Monthly Income	$200	$0
Number of Calls to Hotline	2	25
Prior Residence	Private Emergency Shelter	Private Emergency Shelter
Predicted probability of needing future services within 2 years if given Transitional Housing	Low	High
Predicted probability of needing future services within 2 years if given Emergency Shelter	Medium	High

Which of the following households would you prioritize for Transitional Housing? Select your choice by clicking on the associated column in the image below.

	Household 1	Household 2
Number of Household Members	1	3
Number of Children	0	2
Head of Household Gender	Female	Female
Head of Household Age	45	35
Head of Household Disabling Condition	No	Yes
Head of Household Receives Substance Abuse Services	No	Yes
Monthly Income	$400	$1524
Number of Calls to Hotline	0	25
Prior Residence	Place not meant for habitation	Unknown
Predicted probability of needing future services within 2 years if given Transitional Housing	Low	High
Predicted probability of needing future services within 2 years if given Emergency Shelter	Medium	High

Which of the following households would you prioritize for Transitional Housing? Select your choice by clicking on the associated column in the image below.

	Household 1	Household 2
Number of Household Members	3	1
Number of Children	2	0
Head of Household Gender	Female	Female
Head of Household Age	23	35
Head of Household Disabling Condition	No	No
Head of Household Receives Substance Abuse Services	No	No
Monthly Income	$200	$0
Number of Calls to Hotline	5	0
Prior Residence	Private Emergency Shelter	Unknown
Predicted probability of needing future services within 2 years if given Transitional Housing	High	Low
Predicted probability of needing future services within 2 years if given Emergency Shelter	High	Medium

Which of the following households would you prioritize for Transitional Housing? Select your choice by clicking on the associated column in the image below.

	Household 1	Household 2
Number of Household Members	3	1
Number of Children	2	0
Head of Household Gender	Female	Male
Head of Household Age	23	55
Head of Household Disabling Condition	No	Yes
Head of Household Receives Substance Abuse Services	No	Yes
Monthly Income	$1524	$0
Number of Calls to Hotline	3	2
Prior Residence	Private Emergency Shelter	Private Emergency Shelter
Predicted probability of needing future services within 2 years if given Transitional Housing	High	Low
Predicted probability of needing future services within 2 years if given Emergency Shelter	High	Medium

Which of the following households would you prioritize for Transitional Housing? Select your choice by clicking on the associated column in the image below.

	Household 1	Household 2
Number of Household Members	2	1
Number of Children	1	0
Head of Household Gender	Female	Female
Head of Household Age	23	45
Head of Household Disabling Condition	No	No
Head of Household Receives Substance Abuse Services	No	Yes
Monthly Income	$3256	$200
Number of Calls to Hotline	25	0
Prior Residence	Unknown	Private Emergency Shelter
Predicted probability of needing future services within 2 years if given Transitional Housing	High	Low
Predicted probability of needing future services within 2 years if given Emergency Shelter	High	Medium

Which of the following households would you prioritize for Transitional Housing? Select your choice by clicking on the associated column in the image below.

	Household 1	Household 2
Number of Household Members	1	1
Number of Children	0	0
Head of Household Gender	Female	Male
Head of Household Age	55	35
Head of Household Disabling Condition	Yes	No
Head of Household Receives Substance Abuse Services	No	Yes
Monthly Income	$0	$0
Number of Calls to Hotline	5	2
Prior Residence	Private Emergency Shelter	Private Emergency Shelter
Predicted probability of needing future services within 2 years if given Transitional Housing	High	Low
Predicted probability of needing future services within 2 years if given Emergency Shelter	High	Medium

Which of the following households would you prioritize for Transitional Housing? Select your choice by clicking on the associated column in the image below.

	Household 1	Household 2
Number of Household Members	5	1
Number of Children	4	0
Head of Household Gender	Female	Female
Head of Household Age	23	23
Head of Household Disabling Condition	No	No
Head of Household Receives Substance Abuse Services	No	No
Monthly Income	$668	$32
Number of Calls to Hotline	25	0
Prior Residence	Private Emergency Shelter	Private Transitional Housing
Predicted probability of needing future services within 2 years if given Transitional Housing	High	Low
Predicted probability of needing future services within 2 years if given Emergency Shelter	High	Medium

Task 3 - Instructions for Efficient Group

Deciding who should receive a scarce resource like Transitional Housing is a complex task and there are different ways of prioritizing households.

Two possible prioritizations are:
1) **Efficient** - Give Transitional Housing to the households that would have the biggest benefit from being in Transitional Housing compared to Emergency Shelter
2) **Neediest First** - Give Transitional Housing to the households that would do worst in Emergency Shelter

For example:
Let Household 1 have a probability of reentry if given Transitional Housing of 70% and a probability of reentry if given Emergency Shelter of 90%.

Let Household 2 have a probability of reentry if given Transitional Housing of 30% and a probability of reentry if given Emergency Shelter of 70%

The **Neediest First** prioritization would give Household 1 Transitional Housing because with a probability of reentry of 90%, Household 1 is predicted to do worse in Emergency Shelter than Household 2

The **Efficient prioritization** would give Household 2 Transitional Housing because Household 2 gets a benefit of 40 percentage points (70 - 30) by moving from Emergency Shelter to Transitional Housing. Whereas Household 1 only gets a benefit of 20 percentage points (90 - 70).

Another example:
Let Household 1 have a probability of reentry if given Transitional Housing of 20% and a probability of reentry if given Emergency Shelter of 30%.
Let Household 2 have a probability of reentry if given Transitional Housing of 60% and a probability of reentry if given Emergency Shelter of 90%

The **Neediest First** prioritization would give Household 2 Transitional Housing because with a probability of reentry of 90%, Household 2 is predicted to do worse in Emergency Shelter than Household 1

The **Efficient prioritization** would also give Household 2 Transitional Housing because Household 2 gets a benefit of 30 percentage points (90 - 60) by moving from Emergency Shelter to Transitional Housing. Whereas Household 1 only gets a benefit of 10 percentage points (30 - 20).

Next, instead of descriptions of households, we will present only predictions based on past administrative data about whether households will need future services. These will again be presented pairs. Please decide which household to prioritize for Transitional Housing. You will be presented with 12 pairs of households.

Your goal is to make the most **efficient** assignment. As a reminder, an efficient assignment gives Transitional Housing to households that would have the biggest benefit from being in Transitional Housing compared to Emergency Shelter.

Task 3 - Instruction for Neediest Group

Deciding who should receive a scarce resource like Transitional Housing is a complex task and there are different ways of prioritizing households.

Two possible prioritizations are:
1) **Efficient** - Give Transitional Housing to the households that would have the biggest benefit from being in Transitional Housing compared to Emergency Shelter
2) **Neediest First** - Give Transitional Housing to the households that would do worst in Emergency Shelter

For example:
Let Household 1 have a probability of reentry if given Transitional Housing of 70% and a probability of reentry if given Emergency Shelter of 90%.
Let Household 2 have a probability of reentry if given Transitional Housing of 30% and a probability of reentry if given Emergency Shelter of 70%

The **Neediest First** prioritization would give Household 1 Transitional

Housing because with a probability of reentry of 90%, Household 1 is predicted to do worse in Emergency Shelter than Household 2

The **Efficient prioritization** would give Household 2 Transitional Housing because Household 2 gets a benefit of 40 percentage points (70 - 30) by moving from Emergency Shelter to Transitional Housing. Whereas Household 1 only gets a benefit of 20 percentage points (90 - 70).

Another example:
Let Household 1 have a probability of reentry if given Transitional Housing of 20% and a probability of reentry if given Emergency Shelter of 30%.
Let Household 2 have a probability of reentry if given Transitional Housing of 60% and a probability of reentry if given Emergency Shelter of 90%

The **Neediest First** prioritization would give Household 2 Transitional Housing because with a probability of reentry of 90%, Household 2 is predicted to do worse in Emergency Shelter than Household 1

The **Efficient prioritization** would also give Household 2 Transitional Housing because Household 2 gets a benefit of 30 percentage points (90 - 60) by moving from Emergency Shelter to Transitional Housing. Whereas Household 1 only gets a benefit of 10 percentage points (30 - 20).

Next, instead of descriptions of households, we will present only predictions based on past administrative data about whether households will need future services. These will again be presented pairs. Please decide which household to prioritize for Transitional Housing. You will be presented with 12 pairs of households.

Your goal is to make an assignment that gives the **neediest** households Transitional Housing. As a reminder, the neediest households are those who are predicted to do worst in Emergency Shelter.

Task 3 - Instruction for No Goal Group

Next, instead of descriptions of households, we will present only predictions based on past administrative data about whether households will need future services. These will again be presented pairs. Please decide which household to prioritize for **Transitional Housing**.
You will be presented with 12 pairs of households.

Task 3 - Questions

Which of the following households would you prioritize for Transitional Housing? Select your choice by clicking on the associated column in the image below.

	Household 1	**Household 2**
Predicted probability of needing future services within 2 years if given Transitional Housing	High	Low
Predicted probability of needing future services within 2 years if given Emergency Shelter	High	Low

Which of the following households would you prioritize for Transitional Housing? Select your choice by clicking on the associated column in the image below.

	Household 1	**Household 2**
Predicted probability of needing future services within 2 years if given Transitional Housing	High	Medium
Predicted probability of needing future services within 2 years if given Emergency Shelter	High	Medium

Which of the following households would you prioritize for Transitional Housing? Select your choice by clicking on the associated column in the image below.

	Household 1	Household 2
Predicted probability of needing future services within 2 years if given Transitional Housing	Low	Medium
Predicted probability of needing future services within 2 years if given Emergency Shelter	Low	Medium

Which of the following households would you prioritize for Transitional Housing? Select your choice by clicking on the associated column in the image below.

	Household 1	Household 2
Predicted probability of needing future services within 2 years if given Transitional Housing	High	Low
Predicted probability of needing future services within 2 years if given Emergency Shelter	High	Medium

Which of the following households would you prioritize for Transitional Housing? Select your choice by clicking on the associated column in the image below.

	Household 1	Household 2
Predicted probability of needing future services within 2 years if given Transitional Housing	High	Low
Predicted probability of needing future services within 2 years if given Emergency Shelter	High	Medium

Which of the following households would you prioritize for Transitional Housing? Select your choice by clicking on the associated column in the image below.

	Household 1	**Household 2**
Predicted probability of needing future services within 2 years if given Transitional Housing	High	Low
Predicted probability of needing future services within 2 years if given Emergency Shelter	High	High

Which of the following households would you prioritize for Transitional Housing? Select your choice by clicking on the associated column in the image below.

	Household 1	**Household 2**
Predicted probability of needing future services within 2 years if given Transitional Housing	Low	Medium
Predicted probability of needing future services within 2 years if given Emergency Shelter	Low	High

Which of the following households would you prioritize for Transitional Housing? Select your choice by clicking on the associated column in the image below.

	Household 1	**Household 2**
Predicted probability of needing future services within 2 years if given Transitional Housing	Medium	Low
Predicted probability of needing future services within 2 years if given Emergency Shelter	Medium	Medium

Which of the following households would you prioritize for Transitional Housing? Select your choice by clicking on the associated column in the image below.

	Household 1	Household 2
Predicted probability of needing future services within 2 years if given Transitional Housing	Low	Low
Predicted probability of needing future services within 2 years if given Emergency Shelter	High	Medium

Which of the following households would you prioritize for Transitional Housing? Select your choice by clicking on the associated column in the image below.

	Household 1	Household 2
Predicted probability of needing future services within 2 years if given Transitional Housing	Low	Low
Predicted probability of needing future services within 2 years if given Emergency Shelter	High	Medium

Which of the following households would you prioritize for Transitional Housing? Select your choice by clicking on the associated column in the image below.

	Household 1	Household 2
Predicted probability of needing future services within 2 years if given Transitional Housing	Low	Medium
Predicted probability of needing future services within 2 years if given Emergency Shelter	High	High

Which of the following households would you prioritize for Transitional Housing? Select your choice by clicking on the associated column in the image below.

	Household 1	Household 2
Predicted probability of needing future services within 2 years if given Transitional Housing	Low	Medium
Predicted probability of needing future services within 2 years if given Emergency Shelter	Medium	High

Demographics

What is your age?

What is the highest level of school you have completed or the highest degree you have received?

- Less than high school degree
- High school graduate (high school diploma or equivalent including GED)
- Some college but no degree
- Associate degree in college (2-year)
- Bachelor's degree in college (4-year)
- Master's degree
- Doctoral degree
- Professional degree (JD, MD)

Are you Spanish, Hispanic, or Latino or none of these?

- Yes
- None of these

Are you Spanish, Hispanic, or Latino?

- Spanish
- Hispanic
- Latino

Choose one or more races that you consider yourself to be:

- ☐ White
- ☐ Black or African American
- ☐ American Indian or Alaska Native
- ☐ Asian
- ☐ Native Hawaiian or Pacific Islander
- ☐ Other ________

With which gender do you most identify?

- ○ Male
- ○ Female
- ○ Transgender Male
- ○ Transgender Female
- ○ Gender Variant/Nonconforming
- ○ ________ Other
- ○ Prefer Not to Answer

Prior to your involvement in this study, how would you rate your familiarity with homelessness or homeless services?

- ○ Not at all Familiar
- ○ Slightly Familiar
- ○ Somewhat Familiar
- ○ Moderately Familiar
- ○ Extremely Familiar