

Who Pays the RENT? Implications of Spatial Inequality for Prediction-Based Allocation Policies

Tasfia Mashiat^{1,*}, Patrick J. Fowler², Sanmay Das³

¹University of Iowa

²Washington University in St. Louis

³Virginia Tech

tasfia-mashiat@uiowa.edu, pjfowler@wustl.edu, sanmay@vt.edu

Abstract

AI-powered scarce resource allocation policies rely on predictions to target either specific individuals (e.g., high-risk) or settings (e.g., neighborhoods). Recent research on individual-level targeting demonstrates conflicting results; some models show that targeting is not useful when inequality is high, while other work demonstrates potential benefits. To study and reconcile this apparent discrepancy, we develop a stylized framework based on the Mallows model to understand how the spatial distribution of inequality affects the effectiveness of door-to-door outreach policies. We introduce the RENT (Relative Efficiency of Non-Targeting) metric, which we use to assess the effectiveness of targeting approaches compared with neighborhood-based approaches in preventing tenant eviction when high-risk households are more versus less spatially concentrated. We then calibrate the model parameters to eviction court records collected in a medium-sized city in the USA. Results demonstrate considerable gains in the number of high-risk households canvassed through individually targeted policies, even in a highly segregated metro area with concentrated risks of eviction. We conclude that apparent discrepancies in the prior literature can be reconciled by considering 1) the source of deployment costs and 2) the observed versus modeled concentrations of risk. Our results inform the deployment of AI-based solutions in social service provision that account for particular applications and geographies.

Introduction

Prediction-based resource allocation systems calibrate available task-relevant information to make targeted decisions for individuals in different societal service provisions, such as homelessness, education, medical care, child welfare, etc. (Chouldechova et al. 2018; Kube, Das, and Fowler 2019; Bruce et al. 2011; Ayer et al. 2019). Generally, the prediction produces a score for each individual that indicates the need for intervention. Scoring informs decision-making in the context of perceived needs among the overall population. In limited resource allocation settings, scores provide a rank for targeting interventions until the budget runs out.

Alternatively, communities often rely on non-targeting allocation methods that proportion resource provision based on

the aggregated welfare of a larger unit containing multiple individuals. A motivating example for this paper involves door-to-door canvassing in neighborhoods with high rates of tenant eviction. Canvassers approach all residents in an area to offer information on available resources, regardless of housing precarity or the need for services. Such outreach strategies are easier and faster to coordinate than prediction-based allocations (Moon and Guha 2024; Johnson and Zhang 2022); however, for relatively low base-rate events, like eviction filing, the lack of targeting proves inefficient and potentially wasteful.

Although the selection process of units of targeted versus non-targeted allocation systems differs, the underlying goal remains the same – *Within a given budget, serve as many needy individuals as possible.*

Recent work of Shirali, Abebe, and Hardt (2024) raises interesting questions regarding the relative utility of prediction-based versus non-targeting allocation methods. In their proposed mathematical model, they show that prediction-based targeting policies outperform non-targeting policies for allocating interventions only when the across-unit inequality of welfare is low and the allocation budget is high. Otherwise, non-targeted allocation outperforms prediction-based individualized allocation in all other social welfare and budget contexts.

We note two key model assumptions that introduce potential alternatives that limit generalizations made by Shirali et al. First, the authors consider two categories of allocation costs in their conceptualization: the cost-to-target and the cost-to-treat for each individual. In their work, the modeling of non-targeting policies only includes cost-to-treatment, while targeted policies incorporate both costs. Thus, the model cannot easily capture the costs of service delivery to different locations in situations where the inequality is spatially heterogeneous. Second, Shirali et al. assume that resource allocators possess prior knowledge regarding which individuals exhibit the highest levels of welfare and refrain from allocating interventions to such individuals, even under conditions of non-targeted allocation. Together, the assumptions impose a form of targeting upon specific individuals, despite the intention of non-targeted allocation across units. The assumptions fit certain resource allocation scenarios, for example, allocation of resources to schools in order to treat specific students who may need additional assistance (Perdomo et al.

*Work was performed while the author was a PhD student at George Mason University.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

2025). However, they miss essential features of other scenarios, such as door-to-door outreach, where the marginal cost of treating an additional unit depends on spatial location relative to other at-risk units (Mashiat et al. 2024).

In this paper, we focus on the door-to-door outreach setting. Individuals (or households, we use the terms interchangeably in this paper) reside within specific spatial locations, and costs are associated with traveling from one household to another. As described, our motivating example is the allocation of outreach case workers to canvass residential properties located in neighborhoods to prevent tenant eviction (Mashiat et al. 2024). Each of these properties has an eviction risk score and a geospatial location. We assess the success of allocation policies based on the distribution of High-Risk properties across neighborhoods. We introduce a novel metric, RENT (Relative Efficiency of Non-Targeting), and a stylized framework based on the classic Mallows model of ranking to evaluate the effectiveness of outreach allocation policies for preventing evictions.

Our results demonstrate the connection between the spatial distribution of risk and the overall efficiency of targeting. In particular, accounting for the cost of outreach, targeting policies substantially outperform non-targeting policies in reaching more at-risk households when inequality is spatially dispersed. This raises the question of what level of dispersion is necessary to demonstrate the effectiveness of targeting policies. We introduce a method for calibrating our stylized framework to real-world data and analyze eviction records from St. Louis (one of the most segregated cities in the United States on both racial and socioeconomic measures) (Cooperman 2014). Surprisingly, we find that the Mallows model dispersion parameters corresponding to the data from St. Louis are quite high, indicating that even when segregation is high by standard measures (and thus, inequality of risk is concentrated by a lay definition), targeting adds value for achieving the best use of limited social service resources, canvassing a higher number of risky properties thus reducing evictions than non-targeted measures.

Our Contributions

- We introduce a stylized framework based on the classic Mallows model of ranking for generating spatial distributions of inequality with different concentrations of risk across neighborhoods.
- We use the framework to study the relative effectiveness of non-targeting methods versus prediction-based targeting methods, which we call the RENT metric.
- Leveraging a community partnership, we calibrate the stylized model and evaluate the performance of targeting and non-targeting policies in St. Louis, MO. Surprisingly, even in a highly segregated city, the dispersion parameter that best models the risk distribution is high, showing that targeting still has significant value and utility.

Literature Review

Risk-based Intervention Allocation

Individual targeting-based allocation policies are widely applied in various resource allocation domains. Targeting poli-

cies often rely on predictive models that derive risk scores based on how well individuals would do contingent on receiving the intervention (Kube, Das, and Fowler 2019), or on balancing the expected favorable outcome of an intervention against its associated cost in dynamic settings (Tang et al. 2023).

However, targeting can also be based on the vulnerability of individuals receiving the intervention; the risk scores indicate vulnerability (Pokharel, Das, and Fowler 2024; Elster 1992). The vulnerability scores are then used to form a ranking or assign a category (e.g., high-risk, low-risk). Based on a preference metric/ prioritization scheme set by intervention providers, individuals are selected to receive the intervention. Another conceptualization, which is particularly relevant in tenant eviction, considers risk as the probability of experiencing a bad event (e.g., eviction, emergency room visits, bankruptcy, etc.) in a future time frame (Ayer et al. 2019; Deo et al. 2013; Bastani et al. 2021; Lee, Lavieri, and Volk 2019).

In the operations research community, one popular paradigm is “predict-then-optimize”, referring to first utilizing a predictive model and ensuring its accuracy, then optimizing the outcome or the allocation within a budget (Vanderschueren et al. 2022). In allocating homeless services, Kube, Das, and Fowler (2019) propose a similar approach in which they allocate counterfactual services to each household and identify the service for which the probability of returning to homelessness is the lowest. Work on decision-focused learning adds the allocation decisions into modeling prediction, such that the predictive model is aware of the need for allocation (Ren, Byun, and Wilder 2024; Cortes-Gomez et al. 2025; Elmachtoub, Liang, and McNellis 2020; Shah et al. 2022; Mukhopadhyay et al. 2017; Wilder, Dilkina, and Tambe 2019).

In a different setting, a risk prediction tool, the Dropout Early Warning System (DEWS) (Bruce et al. 2011), ranks students according to their school dropout likelihood score. Based on this ranking, the system is designed to choose a threshold and categorize students with a probability of graduation below the threshold as needing interventions, compared to those with scores above the threshold (Perdomo et al. 2025). However, Perdomo and colleagues show that modeling the dropout risk individually does not improve the outcome compared to targeting a broader scale. Interestingly, Mashiat et al. (2024) find that a risk-based outreach routing policy outperforms several baselines that target tenant eviction prevention at a broader geographic scale.

Ranking Generation

Ranking based on measured scores denotes the relative order of individuals and is often used as a prioritization approach for treatment allocation. Several methods have been proposed to aggregate ranking information, such as the Bradley-Terry, Plackett-Luce, and Mallows models (Bradley and Terry 1952; Plackett 1975; Mallows 1957). Of these, the Mallows model has been applied most in ranking contexts (Vitelli et al. 2018).

The Mallows model (Mallows 1957) is a fundamental and widely applied approach for the generation and interpretation

of ranking data. It is frequently compared to the normal distribution, but is limited to permutations. The model has two key elements: the central order, which represents a chosen ranking of available options, and the dispersion parameter $\phi \in [0, 1]$. The value of ϕ influences whether the rankings are more closely concentrated around the central order (i.e., ϕ closer to 0) versus more uniformly dispersed across all possible permutations (i.e., ϕ closer to 1).

The Mallows model is popular for generating realistic synthetic data by controlling the parameters. Due to its flexibility in producing arbitrarily large synthetic ranking data, the model can be applied to several problem settings, such as voting (Betzler, Bredereck, and Niedermeier 2014), clustering (Busse, Orbanz, and Buhmann 2007), matching (Iruozki, Calvo, and Lozano 2019), and analyzing distributions of ranking data (Lu and Boutilier 2014; Collas and Iruozki 2021).

In our work, we apply the Mallows model to explore the spatial distribution of rankings in different groups of interest. The group membership of individuals is a representation of settings where individuals are part of a larger unit. For example, students in a school or schools in a school district. We focus on tenant eviction, where residential properties are located in different neighborhoods.

Tenant eviction is a major crisis in the US and has serious negative implications, especially for low-income households. Evicted households experience economic fallout and emotional stress, including depression, anxiety, and substance abuse (Himmelstein and Desmond 2021; Acharya, Bhatta, and Dhakal 2022; Bushman and Mehdipanih 2022; Kim 2021). Although governmental and non-governmental agencies coordinate community responses to eviction, efforts often struggle to identify and assist households at risk of eviction with limited and poorly targeted resources (Boen et al. 2023; Donnelly and Farina 2021; Leifheit et al. 2021).

Despite the current limitations, the machine learning community is engaging in efforts to address critical challenges. Initiatives include forecasting future evictions by employing machine learning and deep learning frameworks to direct outreach case workers to residential properties that are at high risk of eviction (Tabar et al. 2022; Sarkar et al. 2024; Mashiat et al. 2024). However, the effectiveness of these prediction-based targeting policies as a function of the distribution of inequality remains an open question. This problem motivates our work.

Spatial Inequality in Resource Allocation

In societal resource allocation settings, inequality is defined in terms of a discrepancy in the allocation of social goods, such as homeless services, child welfare, and aggregated social welfare (Kube et al. 2022; Azizi et al. 2018; Saxena and Guha 2024; Shirali, Abebe, and Hardt 2024) among different populations or groups. In this work, we focus on spatial inequality, precisely in the context of eviction prevention.

Previous research has reported that factors such as economic disparities, political power dynamics, systemic discrimination, geographic location, etc., can drive spatial inequality (Gligorić et al. 2023). Such inequalities are observed in access to proper healthcare (McCrum et al. 2022), transportation (Sanchez, Stolz, and Ma 2003), education (Zahl-

Thanem and Rye 2024), etc. For example, reports show discrepancies in the allocation of COVID-19 vaccines between states in the USA based on the existence of a vulnerable population (Bilal et al. 2022). We also derive a similar notion of spatial inequality between neighborhoods in terms of the distribution of eviction risk and analyze the implications on the effectiveness of outreach.

Reconciling Existing Results

The two most directly related papers to this one are that of Shirali et al., discussed at length above, and the work of Mashiat et al. (2024) on using risk scores to inform eviction outreach. Our work is motivated by the setting of Mashiat et al., in that we model the specific problem of door-to-door outreach for eviction prevention. At the same time, it directly interrogates the key question raised by Shirali et al. – how does the distribution of inequality (in our case, the inequality is manifested as eviction risk) across space affect the value of targeting? Mashiat et al. report significant value in targeting, while Shirali et al. describe a model where targeting has limited value when inequality is high. We attempt to reconcile these two results and uncover two major factors in the apparent discrepancy. First, the cost of outreach in eviction prevention is such that the marginal cost of outreach is deeply dependent on the spatial distribution of risk, a feature that is not present in the model of Shirali et al. Second, what we might intuitively consider a high inequality regime for eviction risk in practice (the highly segregated neighborhoods of St. Louis) may not correspond with what would mathematically be high inequality in a formal model like the one we present here, based on the Mallows model.

A General Framework

We now introduce our stylized framework for generating spatial distributions of risk. We use this framework specifically to analyze eviction outreach policies, but it could also be adapted for other types of resource allocation tasks.

Spatial Distribution of Risk or Inequality

Our model assumes that there are N distinct neighborhoods, each with a fixed number (n) of properties (the restriction to each neighborhood having the same number of properties is not necessary or limiting, but simplifies the exposition). Each property has a specific *risk score*, and ordering by risk scores yields a ranking of all properties. A cutoff based on the percentile or risk score can be used to differentiate between *High-Risk* and *Low-Risk* properties.

First, consider the case of complete concentration of inequality by neighborhood. In this case, Neighborhood 1 (numbering by risk, without loss of generality) has the n riskiest properties, Neighborhood 2 the next n riskiest, etc. Let us call this the *homogeneous risk ranking*, σ_0 . In order to vary how inequality/risk is distributed across neighborhoods, we use the Mallows model (Mallows 1957), which takes two parameters, a central ranking σ and a dispersion parameter $\phi \in [0, 1)$, and is defined as follows,

$$P(r) = P(r \mid \sigma, \phi) = \frac{1}{Z} \phi^{d(r, \sigma)} \quad (1)$$




Figure 1: Gini index across neighborhoods increases with the increase in the dispersion parameter ϕ

where r is a ranking, d is a distance function between rankings, and Z is a normalization constant.

For $\phi \rightarrow 1$, the rankings become equiprobable, while $\phi = 0$ concentrates all the weight on the central ranking σ . The most commonly used distance functions are the Kendall Tau and Spearman Footrule distances (Vitelli et al. 2018). The Kendall Tau distance counts the number of pairwise disagreements between two rankings, while the Spearman Footrule distance sums the deviations in ranking for each item. We experiment with both in this work and find very similar results, so we report only results with the Kendall Tau distance.

By using $\sigma = \sigma_0$ and varying the dispersion parameter, we can effectively control how spread out the High-Risk properties are over neighborhoods; we assume that the sampled permutation maintains property identities (and, therefore, risk score/status), while “moving” them into different neighborhoods. For a simple example, suppose there are 6 properties in total, spread across 3 neighborhoods, and the ranking (5, 1, 3, 2, 6, 4) is sampled. This means that the properties ranked fifth and first in terms of risk are in Neighborhood 1, the properties ranked third and second are in Neighborhood 2, and the properties ranked sixth and fourth are in Neighborhood 3. In the homogeneous risk ranking, the two riskiest properties are in Neighborhood 1, the next two in Neighborhood 2, and the two least risky in Neighborhood 3.

We note the model above could in theory lead to some permutations that are similar in risk characteristics but with neighborhood “identities” changed while appearing to be distant from the baseline rankings. In order to ensure that this does not create an unintended effect, we conducted a simulated experiment with varying values of ϕ to ensure that increasing ϕ does affect the overall distribution of risk across neighborhoods as intended. The simulation involved 600 properties divided into 30 neighborhoods. Keeping the center order (σ_0) homogeneous, we varied ϕ in the range (0,1). We then estimated the Gini index across neighborhoods. Figure 1 shows that as the value of the dispersion parameter increases, the Gini index also increases. Thus, increasing ϕ does indeed induce a more dispersed distribution of High-Risk properties and reduces homogeneity across neighborhoods.

Canvassing Costs and Policies

The resource allocation task we specifically consider is canvassing high(er)-risk properties to provide information about resources for those who may be facing eviction – a task that has received prior attention in the literature (Mashiat et al. 2024; Tabar et al. 2022). The importance of the neighborhood structure is that the marginal cost of visiting additional properties in a given neighborhood is lower than the cost of going to a different neighborhood. To model this, while keeping the analysis tractable, we fix the cost of traveling between properties within a neighborhood as 1 and allow the cost of traveling between neighborhoods, α , to vary.¹

Non-targeting policies conduct outreach at the neighborhood level. The structure is simple; policies rank neighborhoods according to the number of High-Risk properties and allocate outreach canvassers to all properties within the neighborhood, continuing down the list of neighborhoods until the budget runs out.

Prediction-based **targeting** policies use individual risk assessments to decide which properties to canvass. We consider two types of such policies. *High-Risk Property Targeting (HPT)* is similar to non-targeting policies with the difference that within a neighborhood, the policy only canvasses High-Risk properties. Thus, HPT ranks neighborhoods in terms of the number of High-Risk properties and then sequentially visits all High-Risk properties in each neighborhood, in that order, until exhaustion of the budget.

In some situations, fairness or regulatory constraints may require that higher-risk properties are always canvassed preferentially over lower-risk ones. In these cases, the actual risk score or rank matters, not just whether a property is High- or Low-Risk. We define the *Top-k Property Targeting (TPT)* policy as the one that visits the k highest risk properties. Given a value of k , the optimal path through these properties is easy to compute since the properties can be partitioned into their neighborhoods. It is also efficient to perform a binary search to find the highest value of k that can be achieved within a given budget.

Evaluation

In order to evaluate the effectiveness of targeting, we must consider the accuracy of predictions, as well as the budget of the agency implementing the policy. First, we need to consider the relative frequencies of false positives and false negatives in prediction. Let

$$\begin{aligned} Pr(\text{Eviction}|\text{High-Risk}) &= p \\ Pr(\text{Eviction}|\text{Low-Risk}) &= q \end{aligned}$$

Suppose the agency has a budget of c . We define a metric, the *Relative Efficiency of Non-Targeting (RENT)* to evaluate the effectiveness of outreach allocation policies. *RENT* at budget c is defined as follows:

$$RENT_c = \frac{|S_{B,c}|}{|S_{T,c}|} \quad (2)$$

¹Of course, in the real-world context both intra- and inter-neighborhood travel costs will vary, but we find this a useful and reasonable approximation in experiments with actual data.

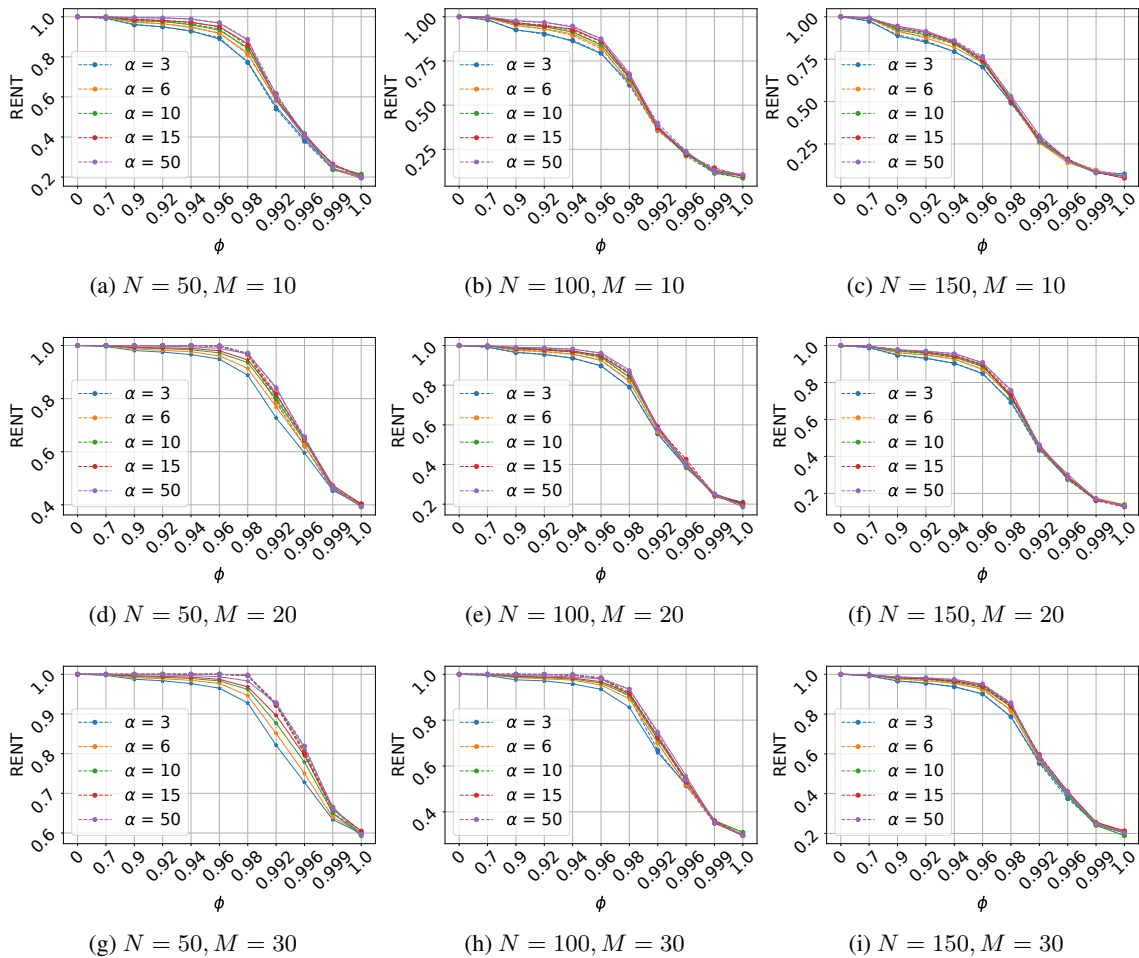


Figure 2: Results of simulation for 600 properties located in N neighborhoods. The total cost is estimated for canvassing the top M neighborhoods ordered by the central ranking. ϕ , the dispersion parameter, is on the X-axis, while the Y-axis shows the $RENT$ values. For both **targeting** policy HPT (solid) and TPT (dashed), as the distribution of inequality across neighborhoods becomes less homogeneous and divorced from the central order, $RENT$ decreases. As a result, the effectiveness of **targeting** policies increases.

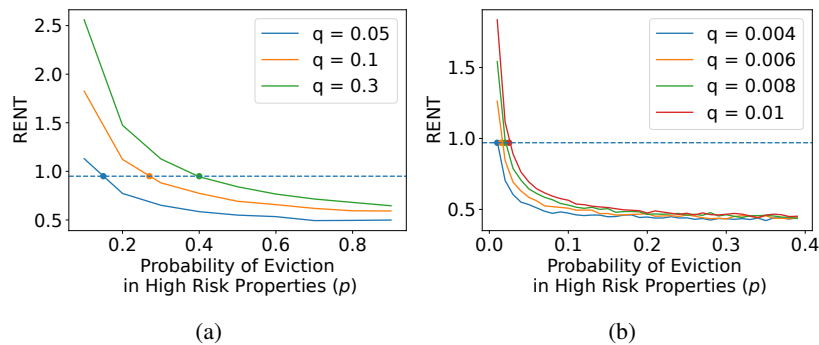


Figure 3: $RENT$ as a function of the precision of prediction (p) and the false negative rate (q). We see that $RENT$ is well below 1, indicating a high relative value of targeting for most reasonable values of precision and false negative rate. Figure (a) shows a broader set of possible values, while Figure (b) focuses on values of p and q that are more plausible for eviction prediction, where the base rate of actual evictions is quite low, and therefore precision is unlikely to exceed 0.4, while the false negative rate will usually be below 0.01.

Where, $S_{B,c}$ denotes the set of properties that experienced eviction canvassed by the **non-targeting** policy and $S_{T,c}$ denotes the set of properties that experienced eviction canvassed by a **targeting** policy for budget c . The lower the RENT, the more effective targeting is.

Simulation Results

We now turn to understanding how the dispersion parameter ϕ affects RENT. We do so by simulating a city with different numbers of neighborhoods. Suppose that there are 600 properties in a city. As case studies, we divide these properties into $N = \{50, 100, 150\}$ neighborhoods.

We consider values of α (the travel cost between neighborhoods) of $\{3, 6, 10, 15, 50\}$. For each value of the dispersion parameter ϕ examined, we sample rankings from a Mallows model with parameters σ_0 and ϕ .

We determine a budget c by calculating the cost to canvass all properties within M neighborhoods using a non-targeting policy. Subsequently, we identify the properties canvassed by HPT and TPT at budget c . We assume that in HPT, the top 20% of all properties are categorized as High-Risk, and in TPT the value of k is set to the maximum number of highest ranked properties that can be visited within budget c . For our initial experiments, we assume that $p = 1$ and $q = 0$, so that the risk predictions are essentially perfect. Figure 2 shows how RENT varies with ϕ for different settings of N , M , and α .

There are two main observations from these results. First, the RENT metric is close to 1 for a wide range of ϕ values and then drops rapidly as ϕ approaches 1. This finding implies that targeting is much more valuable when High-Risk properties are widely dispersed between neighborhoods, as expected. It raises the question of how the ϕ parameter corresponds to what we might expect to see in the real world, which we address in the next section with data from St. Louis. The second observation is that the relative value of targeting is higher when there are more neighborhoods and lower travel costs between neighborhoods. Although intuitive, our framework provides a way of capturing, formalizing, and qualifying the results.

The next question concerns how the accuracy of predictions affects RENT. As p declines and q increases, the relative value of targeting must decrease. It is easy to compute values of S_T and S_B because S_T in expectation is just p times the number of properties that the targeting method visits, while S_B is p times the number of high-risk properties visited by the non-targeting method plus q times the number of low-risk properties visited.

Figure 3(a) shows how RENT varies for $p = [0.1 - 0.9]$, $q = [0.05, 0.1, 0.3]$, and $\phi = 0.99$ (close to the ϕ values for real eviction data, later discussed). Targeting performs better as long as the precision p is reasonably higher than the false negative rate q (for example, once $p = 0.3$ for $q = 0.1$ and about $p = 0.15$ for $q = 0.05$). Since evictions are relatively low baseline probability events, we actually typically see very low values of both p (less than 0.4) and q (less than 0.01) in eviction prediction, while still being able to maintain a predictive edge. Figure 3(b) shows the RENT metric in this regime, and it is interesting to see that RENT is quite low,

implying significant benefits of targeting when p is 0.05 or above.

Application to Real-World Data

In order to understand the applicability of our framework, we turn to a detailed analysis of data on eviction risk from St. Louis, MO. We use geospatial data on properties, predictions of eviction risk, and actual court eviction filings from recent work of Mashiat et al. (2024). St. Louis is a useful example because it is generally considered highly segregated. St. Louis City is ranked 10th in segregation among US cities according to 2020 Census data by the Berkeley Othing and Belonging Institute, and the region is infamous for the so-called “Delmar Divide” where areas north and south of Delmar Boulevard exhibit significantly different racial and socioeconomic makeups and corresponding disparities in life expectancy. We consider census-defined neighborhoods of St. Louis City and municipalities of St. Louis County but refer to these units simply as neighborhoods in what follows for the sake of simplicity.

Dataset Description and Background

The eviction dataset provided by Mashiat et al. has eviction records and risk scores for 26770 residential properties. Among these 19232 are located in St. Louis City and 7543 in St. Louis County (Mashiat et al. 2024).

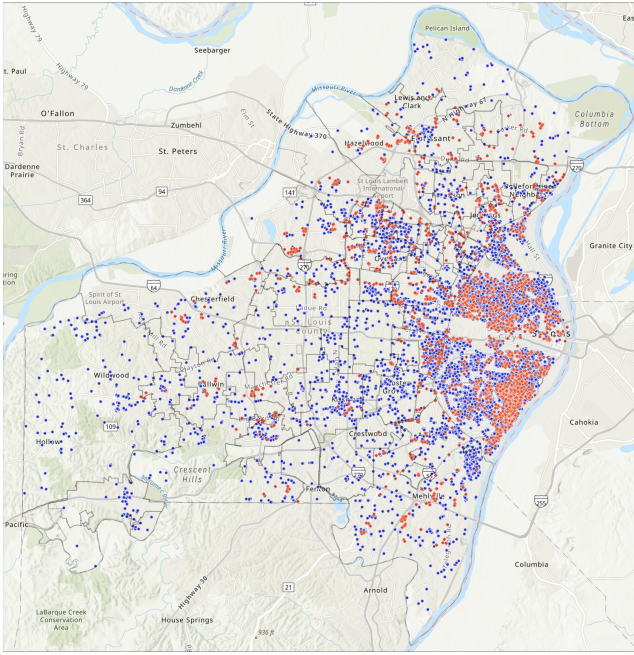
St. Louis City has 79 neighborhoods, while the County has 98 neighborhoods. To remove outliers and capture distributional trends, we only consider neighborhoods with more than 30 residential properties with two or more units.

Mashiat et. al. categorized 8.5% of the properties as High-Risk in their best-performing eviction prediction model, where 1434 and 864 properties are located in the City and the County, respectively. We compute the ratio of High-Risk properties to the total number of properties within each neighborhood and call this the *High-Risk Rate*. Figure 4b and 5 show the distribution of the High-Risk Rate across neighborhoods. The High-Risk properties are also marked as red in Figure 4a and are stratified across the City and the County. We can see that the distribution of the High-Risk Rate across neighborhoods is quite different between the City and the County, with the County displaying much more uniformity in the rate of High-Risk properties across neighborhoods (Figure 5).

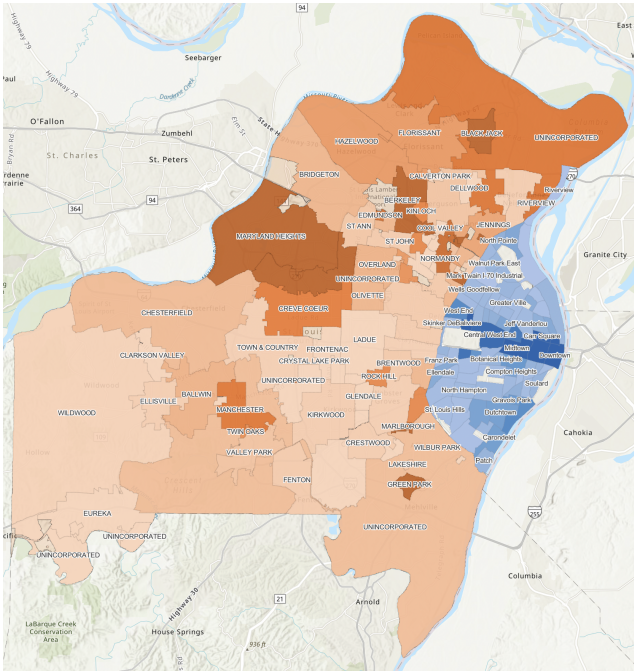
We now turn to calibrating the model developed above to the data from St. Louis City and County in order to find values of ϕ consistent with the observed distribution of High-Risk properties.

Calibrating the Mallows Model

Since our interest is in modeling the heterogeneity of the rates of High-Risk properties across neighborhoods, we conduct two types of calibration processes. For both of them, we simplify the process by assuming that all neighborhoods have $n = 100$ properties each. After removing outliers and neighborhoods with fewer than 30 properties, the remaining number of neighborhoods (N) is 62 for the City and 43 for the County.



(a) Residential properties (properties with higher eviction risk are marked red and low risk are marked blue).



(b) Distribution of High-Risk Rate

Figure 4: (a) Residential properties in St. Louis city and county; the red dots represent the High-Risk properties, and the blue dots represent Low-Risk properties (b) Distribution of High-Risk Rate in St. Louis City (blue) and County (orange) neighborhoods; the color scale represents the value of the rate, and darker shades indicate higher values.




Figure 5: Distribution of High-Risk Rate in St. Louis City (blue) and County (orange) neighborhoods.




Figure 6: The computed Gini index (G) for neighborhoods in St. Louis city and county aligns with the Gini index corresponding to the dispersion parameter (ϕ), when $\phi_{city} = 0.9983$ and $\phi_{county} = 0.998$.

In our first process, we set the High-Risk rates for the City to 0.07 and the County to 0.11, based on real-world data. We initialize our homogeneous central ranking σ_0 such that the top 7% and 11% of properties are categorized as High-Risk for the City and the County, respectively. We perform 1000 simulations for a range of dispersion parameters ϕ . In each simulation, we rank the neighborhoods by their risk rates. We record the minimum and maximum High-Risk rates for each neighborhood by rank. This gives a plausible distribution for what the High-Risk Rate could be in the neighborhood ranked r for each value of r . We then compare the observed distributions of the High-Risk rates by rank with the ranges generated in simulation for different values of ϕ . We find that the risk rate of each neighborhood falls within the risk ranges of the generated rankings produced by the framework when $\phi = \{0.999 \text{ (City)}, 0.9867 \text{ (County)}\}$.

In our second process, we estimate the Gini index (G) across the same range of dispersion parameters (ϕ). Figure 6 shows ϕ and the corresponding G values. We calculate G based on the distribution of High-Risk properties across neighborhoods in the city ($G_{city} = 0.9583$) and the county ($G_{county} = 0.9453$). We observe from Figure 6,

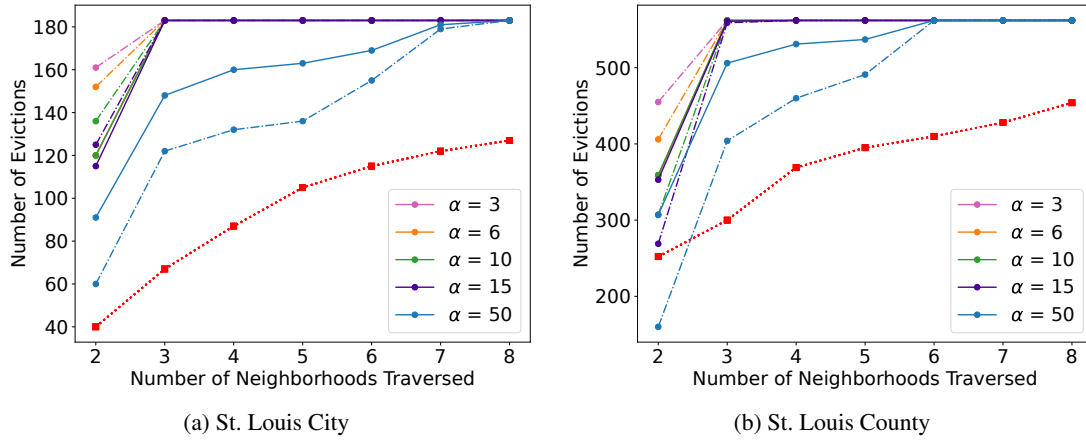


Figure 7: Comparative analysis of HPT (solid) and TPT (dashed) and non-targeting (red line) policies based on the number of discovered evictions in November, December 2022 - January 2023 for St. Louis City and County. The X-axis represents the number of neighborhoods traversed by the non-targeting policy, and the Y-axis represents the number of discovered evictions.

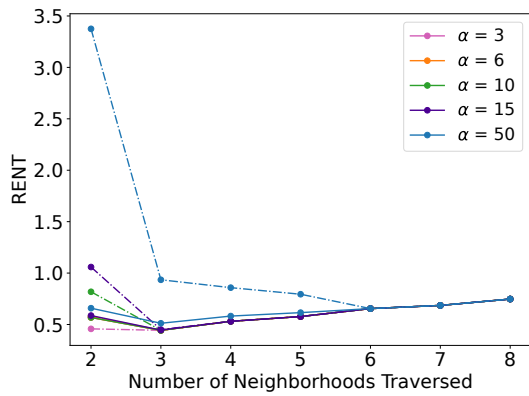
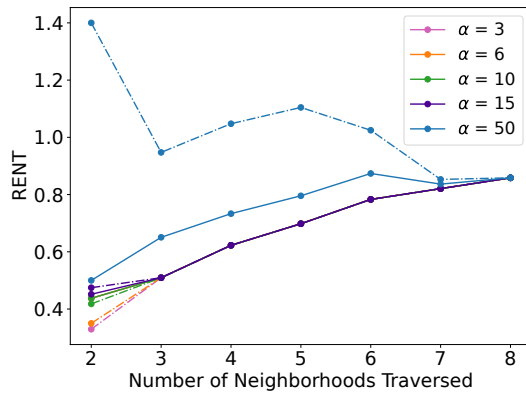


Figure 8: Comparative analysis of RENT values for HPT (solid) and TPT (dashed) relative to the non-targeting policy for St. Louis City and St. Louis County. The X-axis represents the number of neighborhoods traversed by the non-targeting policy, and the Y-axis represents the RENT metric.

that the corresponding ϕ values for these Gini indices, are $\phi = \{0.9983 \text{ (City)}, 0.998 \text{ (County)}\}$

The calculated ϕ values for both city and county are close to 0.99 in both analyses. This shows that even for an area considered emblematic of high levels of segregation and inequality, the dispersion parameters in the Mallows model need to be very high to match the characteristics of the real data. These parameter values correspond to regimes where the RENT metric is very low (Figure 2), indicating that targeting can potentially have significant value.

The Value of Targeting

Next, we turn to comparing the proposed targeting policies (HPT and TPT) relative to the non-targeting policy. As the City and the County are separate jurisdictions, we assume that the budget allocated for outreach in these two areas is determined independently. Suppose that the budgets for canvassing City and County properties are set to c_1 and c_2 , respectively.

We remove the constraint that mandates an equal number of properties in all neighborhoods and consider the actual number of properties for each neighborhood. We use the risk scores computed for predicting evictions in November, December 2022 - January 2023. We note that the computed values of p and q for the City are 0.078 and 0.008, respectively, while those for the County are 0.247 and 0.005.

Using prior eviction data, we rank the neighborhoods based on the number of observed evictions in the previous three months (August - October 2022). The cost of traveling between properties is set to 1, while the cost of traveling between neighborhoods α varies, similar to the simulations above.

Figure 7 shows the comparative analysis between **targeting** and **non-targeting** allocation policies based on the number of discovered evictions. The only situation where the targeting policies do not substantially dominate the non-targeting policy occurs in St. Louis County when the overall budget is low and the travel cost between neighborhoods is




Figure 9: Calibration Curve for eviction risk scores where bin size is 30.

extremely high (Figure 7b, $\alpha = 50$) in comparison to the travel cost within a neighborhood, and the targeting policy is constrained to have to visit higher-risk properties first (TPT, rather than HPT). Figure 8 shows the same results in terms of the RENT metric.

The Utility of Targeting


We now turn to the problem of understanding the value of targeting in terms of eviction reduction under different assumptions about the effectiveness of the intervention that canvassers offer. In order to do this, we first turn to addressing the question of risk score reliability.

Calibrating Risk Scores To understand the reliability of generated risk scores, we calibrate the risk scores for all properties using both Platt Scaling and Isotonic Regression. Figure 9 shows the calibration curves for the uncalibrated risk scores and the calibrated risk scores. The calibrated risk scores from Isotonic Regression are closer to the line representing perfect calibration. We use the eviction probabilities obtained by calibrating risk scores using Isotonic Regression to evaluate and compare the effectiveness of targeting versus non-targeting canvassing policies.


We assume that the intervention reduces the probability of an eviction at a canvassed property by some percentage. We examine four scenarios: no intervention (where risk scores remain unchanged and no reduction in evictions is assumed), and reduction in eviction probability by 30%, 50%, and 70%. Our utility then becomes the expected decrease in the number of evictions through the canvassing intervention.

Similar to our earlier experiment, we consider the High-Risk and Low-Risk properties of neighborhoods from both the city and the county. We set the cost based on the number of neighborhoods canvassed. For simplicity, we set the cost of traveling between properties to 1, within a neighborhood, while the cost of traveling between neighborhoods, $\alpha = 3$.

As the allocated total cost increases, canvassers can visit more properties. Both types of policies lead to better outcomes with higher reductions in the expected number of evictions. However, targeting policies (solid lines) result in



(a) St. Louis City



(b) St. Louis County

Figure 10: We assess utility — defined as the reduction in the expected number of evictions — across various levels of reduction in eviction probability and numbers of canvassed neighborhoods. In the plots, the X-axis shows the number of neighborhoods canvassed for the non-targeting policy (or equivalent cost travel for the targeting policy), while the Y-axis indicates the reduction in expected evictions for targeting (solid lines) and non-targeting (dashed lines) policies. In both St. Louis City (a) and County (b), the targeting approach consistently yielded greater reductions in evictions.

larger reductions in expected evictions compared with non-targeting policies (dashed lines) (Figure 10).

For example, when the probability of eviction is reduced by 30% and the budget is set to the cost of visiting all properties in up to 3 neighborhoods, the targeting policy leads to an expected reduction of 67 evictions in St. Louis County, while the non-targeting policy only leads to an expected decrease of 12 evictions (Figure 10b). As the figure shows, the benefits of targeting become even more noticeable as the reduction in the probability of eviction increases. We also observe similar results with the data on St. Louis City (Figure 10a).

Discussion

This work investigates resource allocation in the context of spatial inequalities of risk in geographic units. We propose

a general framework based on the classic Mallows model to replicate such inequalities in simulation. To control the degree of inequality, we utilize the dispersion parameter of the Mallows model. We analyze and compare both prediction-based targeting and non-targeting allocation policies in the presence and absence of inequality. Our in-depth simulation settings and experiments show that the targeting policy outperforms the non-targeting policy across different resource allocation scenarios. The conditions include, but are not limited to, low-capacity, high-capacity, low-cost, high-cost, etc.

Finally, we apply our general framework to St. Louis, MO, known for its racial and socioeconomic segregation. We show that using our framework can reasonably replicate the distribution of risk within St. Louis City and County. Thus, the comparative results demonstrated in simulated settings hold for St. Louis by association.

Although we present significant findings on spatial inequality and its impact on allocation policies, the results must consider key limitations. Our framework is based on the classic Mallows model of ranking; while the ability to calibrate these models to at least two real, albeit neighboring, urban areas (St. Louis City and County) is promising, it will be important to understand what happens with alternative ranking models as well. Testing the framework for other urban areas will be an interesting extension of the work. Additionally, our framework models inequality in terms of eviction risk; yet, inequality can be defined based on many other features, such as welfare, aggregated income, median rent, etc. Replicating other societal resource allocation settings with limited resources that incorporate a risk-score-based preference metric using our framework represents an exciting application and provides opportunities for further testing in the future. We are confident that policymakers can evaluate both targeting and non-targeting context-aware policies using our framework.

Ethical Statement

Our study is based on an initial model designed to compare allocation policies to help policymakers evaluate AI-powered targeting and non-targeting policies, in the context of social service provision. As a case study, we demonstrate the application of our model to tenant eviction prevention through the allocation of outreach case workers. We are not specifically advocating for the use of any particular policy in practice. The main objective of our work is to assess and evaluate policies in terms of how well they perform in assisting at-risk tenants in the presence of inequality. In the real-world context, the specific features of different approaches (for example, leaving out certain neighborhoods to reach more individuals) would need to be carefully considered by stakeholders.

Acknowledgments

We are grateful for support from NSF Award 2533162. We also acknowledge the St. Louis Regional Response Team and members of the collaborative network who offer ongoing assistance in the research design, including our civic technology partner who assembled the administrative records.

References

- Acharya, B.; Bhatta, D.; and Dhakal, C. 2022. The risk of eviction and the mental health outcomes among the US adults. *Preventive Medicine Reports*, 29: 101981.
- Ayer, T.; Zhang, C.; Bonifonte, A.; Spaulding, A. C.; and Chhatwal, J. 2019. Prioritizing hepatitis C treatment in US prisons. *Operations Research*, 67(3): 853–873.
- Azizi, M. J.; Vayanos, P.; Wilder, B.; Rice, E.; and Tambe, M. 2018. Designing fair, efficient, and interpretable policies for prioritizing homeless youth for housing resources. In van Hoeve, W.-J., ed., *Integration of Constraint Programming, Artificial Intelligence, and Operations Research*, 35–51. Cham: Springer International Publishing. ISBN 978-3-319-93031-2.
- Bastani, H.; Drakopoulos, K.; Gupta, V.; Vlachogiannis, I.; Hadjichristodoulou, C.; Lagiou, P.; Magiorkinis, G.; Paraskevis, D.; and Tsiodras, S. 2021. Efficient and targeted COVID-19 border testing via reinforcement learning. *Nature*, 599(7883): 108–113.
- Betzler, N.; Bredereck, R.; and Niedermeier, R. 2014. Theoretical and empirical evaluation of data reduction for exact Kemeny rank aggregation. *Autonomous Agents and Multi-Agent Systems*, 28: 721–748.
- Bilal, U.; Mullachery, P. H.; Schnake-Mahl, A.; Rollins, H.; McCulley, E.; Kolker, J.; Barber, S.; and Diez Roux, A. V. 2022. Heterogeneity in spatial inequities in COVID-19 vaccination across 16 large US cities. *American Journal of Epidemiology*, 191(9): 1546–1556.
- Boen, C. E.; Keister, L. A.; Gibson-Davis, C. M.; and Luck, A. 2023. The buffering effect of state eviction and foreclosure policies for mental health during the COVID-19 Pandemic in the United States. *Journal of Health and Social Behavior*, 00221465231175939.
- Bradley, R. A.; and Terry, M. E. 1952. Rank analysis of incomplete block designs: I. The method of paired comparisons. *Biometrika*, 39(3/4): 324–345.
- Bruce, M.; Bridgeland, J. M.; Fox, J. H.; and Balfanz, R. 2011. On Track for Success: The Use of Early Warning Indicator and Intervention Systems to Build a Grad Nation. *Civic Enterprises*.
- Bushman, G.; and Mehdipanah, R. 2022. Housing and health inequities during COVID-19: Findings from the National Household Pulse Survey. *J Epidemiol Community Health*, 76(2): 121–127.
- Busse, L. M.; Orbanz, P.; and Buhmann, J. M. 2007. Cluster analysis of heterogeneous rank data. In *Proceedings of the 24th International Conference on Machine Learning, ICML '07*, 113–120. New York, NY, USA: Association for Computing Machinery. ISBN 9781595937933.
- Chouldechova, A.; Benavides-Prado, D.; Fialko, O.; and Vaithianathan, R. 2018. A case study of algorithm-assisted decision making in child maltreatment hotline screening decisions. In Friedler, S. A.; and Wilson, C., eds., *Proceedings of the 1st Conference on Fairness, Accountability and Transparency*, volume 81 of *Proceedings of Machine Learning Research*, 134–148. PMLR.

- Collas, F.; and Irurozki, E. 2021. Concentric mixtures of Mallows models for top- k rankings: sampling and identifiability. In *International Conference on Machine Learning*, 2079–2088. PMLR.
- Cooperman, J. 2014. The story of segregation in St. Louis. *St. Louis Magazine*, 17.
- Cortes-Gomez, S.; Patiño, C. M.; Byun, Y.; Wu, S.; Horvitz, E.; and Wilder, B. 2025. Decision-Focused Uncertainty Quantification. In *The Thirteenth International Conference on Learning Representations*.
- Deo, S.; Iravani, S.; Jiang, T.; Smilowitz, K.; and Samuelson, S. 2013. Improving health outcomes through better capacity allocation in a community-based chronic care model. *Operations Research*, 61(6): 1277–1294.
- Donnelly, R.; and Farina, M. P. 2021. How do state policies shape experiences of household income shocks and mental health during the COVID-19 pandemic? *Social Science & Medicine*, 269: 113557.
- Elmachtoub, A. N.; Liang, J. C. N.; and McNellis, R. 2020. Decision trees for decision-making under the predict-then-optimize framework. In *International Conference on Machine Learning*, 2858–2867. PMLR.
- Elster, J. 1992. *Local justice: How institutions allocate scarce goods and necessary burdens*. Russell Sage Foundation.
- Gligorić, K.; Kamath, C.; Weiss, D. J.; Bavadekar, S.; Liu, Y.; Shekel, T.; Schulman, K.; and Gabrilovich, E. 2023. Revealed versus potential spatial accessibility of healthcare and changing patterns during the COVID-19 pandemic. *Communications Medicine*, 3(1): 157.
- Himmelstein, G.; and Desmond, M. 2021. Eviction and health: A vicious cycle exacerbated by a pandemic. *Health Affairs Health Policy Brief*, 10.
- Irurozki, E.; Calvo, B.; and Lozano, J. 2019. Mallows and generalized Mallows model for matchings. *Bernoulli*, 25: 1160–1188.
- Johnson, R. A.; and Zhang, S. 2022. What is the Bureaucratic Counterfactual? Categorical versus Algorithmic Prioritization in U.S. Social Policy. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '22, 1671–1682. New York, NY, USA. ISBN 9781450393522.
- Kim, D. 2021. Financial hardship and social assistance as determinants of mental health and food and housing insecurity during the COVID-19 pandemic in the United States. *SSM-population health*, 16: 100862.
- Kube, A.; Das, S.; and Fowler, P. J. 2019. Allocating interventions based on predicted outcomes: A case study on homelessness services. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01): 622–629.
- Kube, A.; Das, S.; Fowler, P. J.; and Vorobeychik, Y. 2022. Just resource allocation? How algorithmic predictions and human notions of justice interact. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, EC '22, 1184–1242. New York, NY, USA. ISBN 9781450391504.
- Lee, E.; Lavieri, M. S.; and Volk, M. 2019. Optimal screening for hepatocellular carcinoma: A restless bandit model. *Manufacturing & Service Operations Management*, 21(1): 198–212.
- Leifheit, K. M.; Pollack, C. E.; Raifman, J.; Schwartz, G. L.; Koehler, R. D.; Bronico, J. V. R.; Benfer, E. A.; Zimmerman, F. J.; and Linton, S. L. 2021. Variation in state-level eviction moratorium protections and mental health among US adults during the COVID-19 pandemic. *JAMA Network Open*, 4(12): e2139585–e2139585.
- Lu, T.; and Boutilier, C. 2014. Effective sampling and learning for mallows models with pairwise-preference data. *J. Mach. Learn. Res.*, 15(1): 3783–3829.
- Mallows, C. L. 1957. Non-null ranking models. I. *Biometrika*, 44(1/2): 114–130.
- Mashiat, T.; DiChristofano, A.; Fowler, P. J.; and Das, S. 2024. Beyond Eviction Prediction: Leveraging Local Spatiotemporal Public Records to Inform Action. In *The 2024 ACM Conference on Fairness, Accountability, and Transparency*, 1383–1394.
- McCrum, M. L.; Wan, N.; Han, J.; Lizotte, S. L.; and Horns, J. J. 2022. Disparities in spatial access to emergency surgical services in the US. In *JAMA health forum*, volume 3, e223633–e223633. American Medical Association.
- Moon, E. S.-Y.; and Guha, S. 2024. A Human-Centered Review of Algorithms in Homelessness Research. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, CHI '24. New York, NY, USA: Association for Computing Machinery. ISBN 9798400703300.
- Mukhopadhyay, A.; Vorobeychik, Y.; Dubey, A.; and Biswas, G. 2017. Prioritized Allocation of Emergency Responders based on a Continuous-Time Incident Prediction Model. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*, AAMAS '17, 168–177. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.
- Perdomo, J. C.; Britton, T.; Hardt, M.; and Abebe, R. 2025. Difficult Lessons on Social Prediction from Wisconsin Public Schools. In *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '25, 2682–2704. New York, NY, USA: Association for Computing Machinery. ISBN 9798400714825.
- Plackett, R. L. 1975. The analysis of permutations. *Journal of the Royal Statistical Society Series C: Applied Statistics*, 24(2): 193–202.
- Pokharel, G.; Das, S.; and Fowler, P. 2024. Discretionary trees: understanding street-level bureaucracy via machine learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 22303–22312.
- Ren, K.; Byun, Y.; and Wilder, B. 2024. Decision-Focused Evaluation of Worst-Case Distribution Shift. In *The 40th Conference on Uncertainty in Artificial Intelligence*.
- Sanchez, T. W.; Stolz, R.; and Ma, J. S. 2003. Moving to Equity: Addressing Inequitable Effects of Transportation Policies on Minorities. Technical report, UCLA: The Civil Rights Project / Proyecto Derechos Civiles.
- Sarkar, A.; DiChristofano, A.; Das, S.; Fowler, P. J.; Jacobs, N.; and Vorobeychik, Y. 2024. Geospatial Active Search for

Preventing Evictions. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems*, 2456–2458.

Saxena, D.; and Guha, S. 2024. Algorithmic harms in child welfare: Uncertainties in practice, organization, and street-level decision-making. *ACM Journal on Responsible Computing*, 1(1): 1–32.

Shah, S.; Wang, K.; Wilder, B.; Perrault, A.; and Tambe, M. 2022. Decision-focused learning without decision-making: Learning locally optimized decision losses. *Advances in Neural Information Processing Systems*, 35: 1320–1332.

Shirali, A.; Abebe, R.; and Hardt, M. 2024. Allocation Requires Prediction Only if Inequality Is Low. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, 45114–45153. PMLR.

Tabar, M.; Jung, W.; Yadav, A.; Wilson Chavez, O.; Flores, A.; and Lee, D. 2022. Forecasting the Number of Tenants At-Risk of Formal Eviction: A Machine Learning Approach to Inform Public Policy. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22*, 5178–5184. AI for Good.

Tang, B.; Koçyiğit, Ç.; Rice, E.; and Vayanos, P. 2023. Learning optimal and fair policies for online allocation of scarce societal resources from data collected in deployment. *arXiv preprint arXiv:2311.13765*.

Vanderschueren, T.; Verdonck, T.; Baesens, B.; and Verbeke, W. 2022. Predict-then-optimize or predict-and-optimize? An empirical evaluation of cost-sensitive learning strategies. *Information Sciences*, 594: 400–415.

Vitelli, V.; Sørensen, Ø.; Crispino, M.; Frigessi, A.; and Arjas, E. 2018. Probabilistic preference learning with the Mallows rank model. *Journal of Machine Learning Research*, 18(158): 1–49.

Wilder, B.; Dilkina, B.; and Tambe, M. 2019. Melding the data-decisions pipeline: Decision-focused learning for combinatorial optimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 1658–1665.

Zahl-Thanem, A.; and Rye, J. F. 2024. Spatial inequality in higher education: a growing urban–rural educational gap? *European Sociological Review*.