Choice Errors and Choice Restrictions in Public Housing Allocation

Canishk Naik
UC Berkeley

Neil Thakral
Brown University

March 2026*

Abstract

Public housing allocation requires balancing two objectives: prioritizing vulnerable households and preventing congestion from applicants holding out for desirable units. Housing authorities manage this trade-off by restricting applicants' choices, but the effectiveness of these restrictions depends on the errors applicants make when choosing. We combine administrative vacancy data from a large London council with a choice experiment that separately identifies preferences and biases for 587 public housing residents to compare the welfare effects of two prominent allocation policies. The UK mechanism batches arrivals and lets applicants sign up for acceptable options; the US mechanism makes sequential offers with rejection penalties. Under rationality, the US mechanism reduces congestion at the cost of weaker prioritization. The UK mechanism is more vulnerable to bias—yet some biases can improve average welfare, as high-priority applicants' mistakes free up better options for those below them.

*Naik: Haas School of Business, UC Berkeley (email: canishk@berkeley.edu). Thakral: Department of Economics, Brown University, Box B, Providence, RI 02912 (email: neil_thakral@brown.edu). We are grateful to Doug Bernheim, Kwok-Hao Lee, Matthew Levy, Daniel Reck, Alvin Roth, Johannes Spinnewijn, and seminar participants at Brown, LSE, Stanford and the North American Meeting of the Urban Economics Association for helpful comments and discussions. We also thank Santiago Hermo, Matthew Kutam, Steven Lee, Andres Mena, and Sidharth Satya for excellent research assistance. The study was approved by the Institutional Review Board at Brown University (Protocol #2111003161). This material is based upon work supported by the National Science Foundation under Grant No. 2149371.

1 Introduction

Rules governing how applicants exercise choice lie at the heart of policy design in public housing allocation. On the one hand, policymakers seek to provide better matches to households from vulnerable groups or with urgent needs. On the other hand, housing authorities aim to prevent applicants from holding out indefinitely for highly desirable units and creating unnecessary vacancies, especially in the face of persistent excess demand.

In the UK, over 131,000 households live in temporary accommodation as of March 2025, costing councils over £1 billion annually (Ministry of Housing, Communities & Local Government, 2025), even as nearly 270,000 homes sit long-term empty across England (Action on Empty Homes, 2024). These patterns reflect a central challenge in public housing allocation: how much freedom applicants should have when accepting or rejecting offers. Choice restrictions, such as capping rejections or limiting which homes may be accepted, provide a natural way to balance congestion against prioritization—that is, the extent to which higher-priority households can translate their rank into better outcomes. However, the effectiveness of these restrictions depends on how applicants navigate the allocation process. Errors in belief formation and misunderstandings of program rules can prompt applicants to reject, or fail to bid on, offers that would otherwise improve their welfare (Kullberg, 2002; Marsh, 2004; Brown and King, 2005).

This paper studies how different ways of restricting choice, as exemplified by housing allocation mechanisms used in the UK and the US, shape welfare when applicants exhibit empirically grounded behavioral biases. We combine administrative data on housing arrivals with a discrete choice experiment conducted with public housing residents to separately identify preferences over housing characteristics and key behavioral distortions, including misperceived offer probabilities, present focus, loss aversion, and decision noise. We embed these estimates in a dynamic model of housing allocation and simulate outcomes under alternative mechanisms to evaluate welfare. In the UK mechanism, units arriving within each week are batched and offered simultaneously, and applicants may bid on at most three units per cycle.[1] In the US mechanism, units are offered sequentially as they arrive, with applicants allowed to reject up to two offers before losing their waitlist position. These alternative restrictions generate different trade-offs between prioritization and congestion and interact with behavioral distortions through different margins.

To build intuition, we consider a stylized environment with two applicants and two units. Behavioral deviations interact with choice restrictions through different margins. The limited

[1] In the UK mechanism, "bidding" refers to expressing interest in a unit and does not involve monetary payments.

rejection allowance of the US mechanism tends to produce more balanced outcomes across the priority distribution by preventing high-priority households from holding out indefinitely. Biases that raise rejection rates can trigger large individual losses by pushing applicants into worse future allocations after exhausting their limited rejections, while biases that lower rejection rates speed up allocations but reduce match quality. The batching structure of the UK mechanism allows applicants to condition choices on a richer menu and target desirable units more selectively, which strengthens prioritization but slows matching further down the queue. Biases that reduce bidding or increase selectivity amplify congestion, whereas biases that increase bidding or shorten waiting reduce it. Small amounts of choice error can generate positive externalities for lower-priority applicants by reducing competition for oversubscribed units.

Evaluating these mechanisms quantitatively requires separating preferences from decision errors, which observational choice data alone cannot achieve. Thus, we combine administrative housing data with experimentally elicited preferences and behavioral measures. On the housing side, we use administrative data from a large London council, which provide detailed information on the timing and characteristics of housing arrivals as well as the priority levels of successful applicants. These data allow us to estimate an arrival process for housing and to calibrate applicants' beliefs about the probability of receiving an offer based on historical outcomes.

On the applicant side, we estimate preferences via a Bayesian Adaptive Choice Experiment (BACE) with 587 current or recent social housing residents in the UK (Drake et al., 2025a). Respondents make repeated pairwise choices between housing bundles, with each successive choice set selected to maximize information gain given the respondent's posterior. This approach allows us to identify preferences over housing characteristics in controlled choice sets, and any inconsistencies yield an individual-level measure of decision noise. We separately collect survey measures of biased beliefs about offer probabilities, present focus, and loss aversion (Stango and Zinman, 2023).

We embed these estimates in a dynamic model of housing allocation in which applicants condition their decisions on their priority points and on beliefs derived from historical letting outcomes. We solve for optimal policies via backward induction under both rational and biased behavior, and simulate each mechanism over a five-year horizon. Welfare is computed as the net present value of receiving housing relative to the outside option, and we compare mechanisms in terms of both average outcomes and a priority-weighted measure.

Under full rationality, the US mechanism delivers modestly higher average welfare effects than the UK mechanism (£2,300 vs. £2,150) by clearing the waitlist efficiently and preventing queue congestion. The UK mechanism generates higher welfare for the top priority decile

by allowing selective waiting, but at the cost of growing backlogs that harm lower-priority applicants.

While performing comparably under rationality, the two mechanisms respond very differently to bias. Misperceiving offer probabilities substantially harms UK welfare by directing bids toward unlikely outcomes and increasing congestion, whereas the US mechanism remains largely insulated because limited rejections bound the consequences of mistaken decisions. Loss aversion reduces welfare in both systems through different channels: in the UK, applicants become overly selective and bid less; in the US, it distorts both the acceptance threshold and the continuation value. Decision noise causes high-priority applicants' mistakes to free up better options for those below.

Unlike the other biases, present focus improves welfare for almost all applicants in the UK mechanism by resolving an underlying coordination failure. Offer beliefs are calibrated from historical data, but high-priority applicants wait selectively and rarely accept, so observed outcomes underrepresent high-priority winners. As a result, optimistic beliefs and waiting persists in equilibrium. Present-focused applicants break this cycle by discounting continuation values more heavily, accepting sooner and bidding more actively, which reduces congestion and improves outcomes across the priority distribution. This effect is muted in the US mechanism, where continuation values depend primarily on remaining rejections rather than time. The optimal mechanism therefore depends not only on whether applicants are biased, but on *which* biases prevail.

Our findings contribute most directly to the literature on public housing allocation and, more broadly, to empirical welfare analysis of dynamic matching mechanisms. In public housing, existing research emphasizes the trade-off between match quality and targeting in waiting list systems. Waldinger (2021) demonstrates that choice restrictions can improve social welfare under sufficiently strong preferences for redistribution.[2] Leshno (2022) highlights the value of allowing applicants to decline mismatched offers, and Murra-Anton and Thakral (2025) introduce a new mechanism in this vein. Theoretical work by Bloch and Cantala (2017) and Schummer (2021) characterizes optimal selectivity thresholds, while Thakral (2019) illustrates the difficulty of ensuring ex-post efficiency in dynamic matching. Beyond housing, related empirical work on dynamic mechanisms includes Verdier and Reeling (2020) on hunting licenses, Agarwal et al. (2021) on kidney allocation, and Johnston et al. (2023) on electricity interconnection queues.

In this literature, preferences are typically inferred from choice data, but the possible presence of bias creates challenges in characterizing underlying preferences. Our approach addresses this limitation through a discrete choice experiment that cleanly isolates preferences.

[2]Also see Arnosti and Shi (2020) for a related analysis of the trade-off between matching and targeting.

We contribute to the literature on public housing along two dimensions. First, we compare different *ways* of restricting choice, exemplified by the UK and US mechanisms. Second, we quantify how behavioral distortions interact with these restrictions to shape welfare. We find that the UK mechanism exhibits substantially greater sensitivity to behavioral biases, while the US mechanism's limited rejection structure tends to do a better job of insulating outcomes from misoptimization. As a result, the optimal design of choice restrictions depends not only on preference heterogeneity, but on the distribution of behavioral biases in the applicant pool.

This paper also contributes to recent work on behavioral biases in market design. Empirical evidence documents substantial biases in participants' reported preferences within high-stakes allocation settings (Rees-Jones, 2018; Hassidim et al., 2021).[3] In lab experiments, Meisner and Von Wangenheim (2023) and Dreyfuss et al. (2021, forthcoming) provide evidence of loss aversion in reported preferences. Kapor et al. (2020) document widely dispersed beliefs about offer probabilities in school choice, while Arteaga et al. (2022) and Fabre et al. (2023) show students underestimate their probability of being unassigned. Recent work on centralized education markets incorporates biases in assessing continuation values, with Wang et al. (2025) implying that true continuation values are more favorable than students assume, and Idoux (2023) implying an optimistic assessment of continuation values.

Our paper departs from prior work that infers a single structural bias from observed choices within a given mechanism. Instead, we separately identify preferences in controlled choice sets, measure a range of commonly studied behavioral biases using standardized survey tasks (Stango and Zinman, 2023), and map each bias into a transparent distortion of the applicant's dynamic decision problem. This strategy allows us to quantify how each bias interacts with the specific choice restrictions imposed by different mechanisms.

Finally, our empirical approach adds to a growing literature that uses theory-driven survey instruments (e.g., Ameriks et al., 2020; Andrew and Adams, 2025). In our setting, realized decisions confound preferences with beliefs and behavioral frictions that may be induced by the mechanism itself. The Bayesian adaptive discrete-choice design lets us discipline match values and waiting costs in controlled choice sets while measuring choice consistency and separately eliciting related behavioral parameters. This combination enables welfare counterfactuals without relying on strong restrictions that observational choice data typically require about information, rationality, or equilibrium beliefs.

[3]Also see Hassidim et al. (2017), Rees-Jones and Skowronek (2018), Shorrer and Sóvágó (2023), and Shorrer and Sóvágó (2024).

2 Choice restrictions in public housing

This section describes the policy setting, the US and UK mechanisms that our analysis focuses on, and presents a simple example to build intuition for how the effect of behavioral bias on welfare depends on the choice restrictions imposed by the mechanism.

2.1 Policy context

Public housing allocation in both the United States and the United Kingdom is administered by local agencies operating long waiting lists under persistent excess demand. In the US, applicants are typically offered units sequentially in priority order and may refuse only a limited number before losing their place. In the UK, the vast majority of councils (around 92 percent) operate a system known as *Choice-Based Lettings*, in which applicants "bid" on vacancies that become available each week. Bids are not associated with monetary payments, but are rationed for each applicant to a fixed maximum number each week (usually 3). As far as we know, this mechanism has never been empirically evaluated.

In Camden, our empirical setting, social housing is allocated through a points-based system that translates housing need into a priority score.[4] Points are awarded for circumstances including homelessness, overcrowding, health conditions exacerbated by housing, and harassment or domestic violence, with higher awards reflecting greater urgency. Ties are broken by waitlist entry date.

The central distinction between the two systems is what they restrict. The US mechanism limits the number of rejections; the UK mechanism limits the set of units an applicant makes herself eligible for. Batching is what creates a meaningful choice problem in the UK: if units were offered immediately as they arrived, the mechanism would collapse to a sequential system analogous to the US.

2.2 Defining public-housing allocation mechanisms

A public-housing authority faces the problem of assigning housing units to applicants on a waiting list. Units differ along dimensions such as quality, rent, and location. Applicants differ in both their priority, which reflects the urgency of their need for housing, and their preferences over housing. In this section, we set up a model of this allocation problem.

Environment. Time is discrete, $t = 0, 1, 2, \ldots$. In each period, a (possibly empty) set of housing units B^t becomes available. Each unit $b \in B^t$ is characterized by observable

[4] See Appendix Table 1 for a full description of the points schedule.

attributes (e.g., rent, location, property type, and ownership). Applicants $a \in A$ are ranked by an exogenous priority score $p(a) \in \mathbb{R}$, where higher values indicate greater need for housing. Applicants observe their own priority points and their waitlist entry date. They do not observe the full queue or the realized priorities of other applicants, so they know their absolute priority but not their exact rank relative to others.

If applicant a is assigned unit b in period t, they obtain lifetime utility $U_a(b, t)$. Utilities are normalized relative to each applicant's outside option: if they remain unassigned in period t, they receive flow utility 0 and remain on the waiting list. Thus, welfare is measured as the discounted value of assignment relative to continued waiting.

US-k mechanism (rejection cap). In the US mechanism, housing units are processed sequentially as they become available. When a unit becomes available, it is offered to applicants on the waiting list in descending priority order. An applicant who is offered a unit chooses whether to accept or reject it. If an applicant rejects k times, they lose their position on the waiting list and receive their outside option. Thus, the US mechanism restricts rejections.

UK-k mechanism (accept cap with batching). In the UK mechanism, vacancies that arise within a week are batched and advertised together. Each week t, the housing authority posts the batch of arrivals B^t, and each applicant may bid on k vacancies, indicating interest in specific units. For each unit $b \in B^t$, an offer is made to the highest-priority bidder. Thus, the UK mechanism restricts choice by limiting the set of units an applicant makes herself eligible for within a bidding cycle (or potential *acceptances*). We refer to this rule as the UK-k mechanism.

2.3 Illustrative example

This section sets up a stylized model to build intuition for how different ways of restricting applicants' choices shape welfare outcomes across the priority distribution, especially when applicants make mistakes. For the US mechanism, we consider the case in which applicants may reject at most one offer ($k = 1$). For the UK mechanism, we consider the case in which applicants may bid on at most one unit in each round ($k = 1$). The model clarifies how two mechanisms that share the same underlying priority ordering can still differ in how much priority translates into better outcomes, in the congestion they generate, and in their sensitivity to behavioral distortions. This provides a lens for interpreting the subsequent simulation results by showing which margins each mechanism changes and why the same bias can redistribute welfare differently across the priority distribution.

2.3.1 Setup

Two applicants, A and B, compete for two housing units, X and Y. Applicant A has higher priority. Both applicants prefer X to Y.

Time unfolds discretely and is structured into subperiods to reflect the batching nature of the UK mechanism. Specifically, the initial period (period 0) is broken into two subperiods. At period 0.1, unit Y arrives with certainty. At period 0.2, unit X arrives with probability p. If unit X does not arrive by the end of period 0 (i.e., after period 0.2), it arrives with certainty in period 1.

The reason for subdividing period 0 into subperiods 0.1 and 0.2 is to clearly capture the distinction between the two allocation systems. Under the UK mechanism, all units that arrive within the same larger period (period 0) are batched and allocated in a single round of simultaneous bidding. Under the US mechanism, by contrast, each unit is offered sequentially as soon as it becomes available. This arrival process is the simplest setup that captures the essential trade-off faced by applicants: deciding whether to accept a less desirable unit immediately or to risk waiting for a potentially better match. At the same time, it is rich enough to generate non-trivial bidding behavior under the UK mechanism when multiple units are available simultaneously.

Future utilities are discounted with a per-subperiod discount factor of δ. Receiving a unit in subperiod 0.1 yields the full undiscounted payoff. Receiving a unit in subperiod 0.2 discounts payoffs by δ, and receiving a unit in the subsequent subperiod (period 1) discounts payoffs by δ^2. For simplicity, we set $p = \frac{1}{2}$ and $\delta = \frac{1}{2}$. Both applicants value unit X at 1 and unit Y at $\frac{1}{3}$.

We first analyze outcomes under rational decision-making, then examine how they change when applicants deviate from the optimal decision with some fixed probability ε.

2.3.2 Rational benchmark

US mechanism. Under the US mechanism, A rejects the early offer of Y and waits for X, while B receives Y. Intuitively, because A has high priority and prefers X, it is worthwhile for A to hold out for the better unit rather than accept the worse one immediately. Unit Y then passes to B, who accepts it immediately. Thus, the US mechanism lets the high-priority applicant wait for the better unit, while the low-priority applicant takes the earlier inferior one.

More specifically, accepting Y in period 0.1 gives A utility $\frac{1}{3}$. Rejecting Y and waiting for X gives expected utility $p\delta + (1-p)\delta^2 = \frac{3}{8}$, since X arrives in period 0.2 with probability p and in period 1 otherwise. Because $\frac{3}{8} > \frac{1}{3}$, A optimally rejects Y. Applicant B then accepts

Y in period 0.1 and receives $\frac{1}{3}$. The resulting expected utilities are therefore $\frac{3}{8}$ for A and $\frac{1}{3}$ for B, yielding total welfare of $\frac{17}{24}$.

UK mechanism. Under the UK mechanism, the allocation depends on whether X arrives in period 0 or period 1. If both units are available in period 0, then A bids on X and B bids on Y, so A receives X and B receives Y at the end of the period. If only Y is available in period 0, then B bids on Y while A waits for X in period 1. Thus the allocation looks similar to the US mechanism, but batching delays the assignment of Y to B until the end of the period rather than immediately upon arrival.

More specifically, if X arrives in period 0.2, then A receives X at the end of period 0 for utility δ and B receives Y at the end of period 0 for utility $\frac{1}{3}\delta$. If instead X does not arrive until period 1, then B receives Y at the end of period 0 for utility $\frac{1}{3}\delta$, while A receives X in period 1 for utility δ^2. Expected utilities are therefore $p\delta + (1-p)\delta^2 = \frac{3}{8}$ for A and $\frac{1}{3}\delta = \frac{1}{6}$ for B, yielding total welfare of $\frac{13}{24}$. Relative to the US mechanism, total welfare is lower because batching makes B wait longer to receive Y.

2.3.3 Choice error

Each applicant now deviates from the optimal action with fixed probability ε, allowing us to compare how the same rate of choice error translates into welfare losses across the two mechanisms. This reduced-form perturbation captures choice noise around the optimal decision rule: an applicant sometimes selects the wrong option even when the optimal ranking of actions stays fixed. Later sections also consider directional distortions that systematically shift behavior, e.g., accepting or bidding too much or too little. We do not model the source of bias explicitly here; the goal is simply to compare how the same error rate ε translates into welfare losses across mechanisms. The same bias might affect the two mechanisms differently because the structure of each mechanism determines which mistakes are possible and how costly they are.

US mechanism. Section 2.3.3 plots expected utilities under the US mechanism as a function of ε (see Proposition 1 for a complete derivation). With choice error, applicant A's welfare is U-shaped under the US mechanism. At low error rates, A faces the risk of correctly rejecting Y but then also mistakenly rejecting X—leaving A with nothing. At high error rates, A falsely accepts Y often enough that this insures against the outcome of rejecting both units. Applicant B's welfare declines monotonically: B's optimal action is to accept Y, and any deviation risks rejection and an empty outcome.

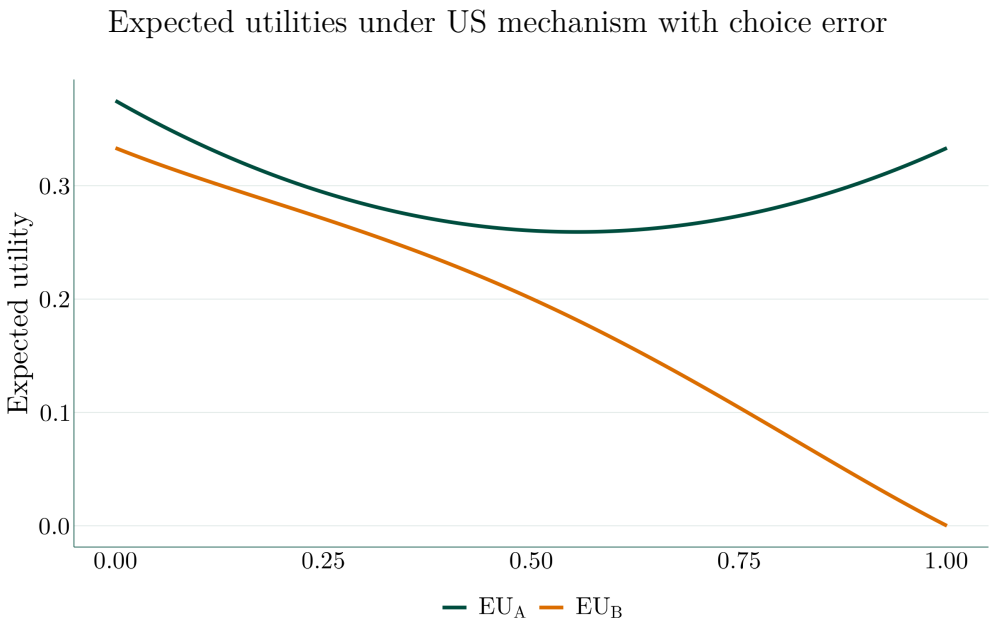

Note: The figure plots expected utilities under the US mechanism as a function of the choice error probability $\varepsilon \in [0, 1]$.

The first-order effects of choice error on expected payoffs under the US mechanism are as follows. For applicant A, mistakenly accepting Y in period 0.1 results in an expected payoff of $\frac{1}{3}$ instead of $\frac{3}{8}$ from waiting for X. Similarly, mistakenly rejecting X when it arrives results in an expected payoff of 0 instead of $\frac{3}{8}$. This leads to a first-order effect of $(\frac{1}{3} - \frac{3}{8}) - \frac{3}{8} = -\frac{5}{12}$ for applicant A, which can be interpreted as the welfare loss from bias for the higher-priority applicant under the US mechanism.

In addition, applicant A mistakenly accepting Y in period 0.1 imposes a positive externality on B, who receives an expected payoff of $\frac{3}{8}$ instead of $\frac{1}{3}$. However, if B mistakenly rejects Y in period 0.1, then B receives a payoff of 0 instead of $\frac{1}{3}$. This leads to a first-order effect of $(\frac{3}{8} - \frac{1}{3}) - \frac{1}{3} = -\frac{7}{24}$ for applicant B. In our simple example, the congestion externality shows up in the fact that A's decision about whether to wait for X determines whether B ever has a chance to receive that unit; bias that pushes A toward accepting Y instead of waiting relaxes this congestion for B.

UK mechanism. Section 2.3.3 plots expected utilities under the UK mechanism (see Proposition 2 for a complete derivation). Applicant A's welfare declines monotonically: A should wait for X, and any error that diverts A toward Y is costly with no offsetting benefit. Applicant B's welfare, however, *increases* with moderate ε. When A mistakenly bids on Y, B inherits a chance at X—a unit B would never access under rational play. This positive externality from A's mistakes more than offsets B's own errors.

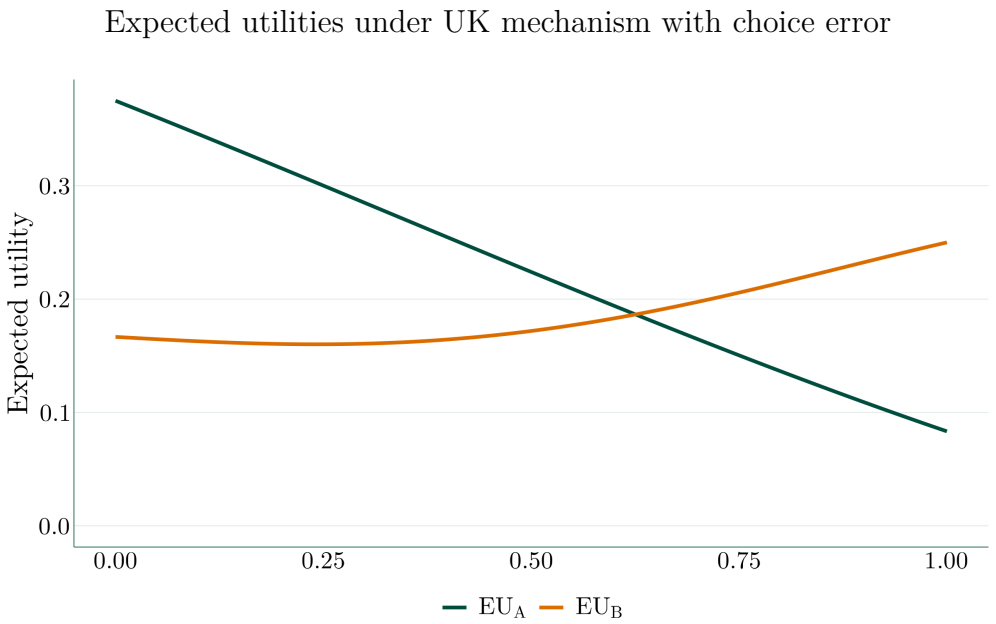

Note: The figure plots expected utilities under the UK mechanism as a function of the choice error probability $\varepsilon \in [0, 1]$.

Bias under the UK mechanism leads to the following first-order effects. For applicant A, if only $\{Y\}$ arrives in period 0, mistakenly bidding on Y results in an expected payoff of $\frac{\delta}{3} = \frac{1}{6}$ instead of $\delta^2 = \frac{1}{4}$ for A. Likewise, mistakenly not bidding on X in period 1 results in a payoff of 0 instead of $\delta^2 = \frac{1}{4}$. If $\{X, Y\}$ arrives in period 0, mistakenly not bidding on X in period 0 results in a payoff of $\delta^2 = \frac{1}{4}$ instead of $\delta = \frac{1}{2}$ due to having to wait until the next period for X. This leads to a first-order effect of $\frac{1}{2}(\frac{1}{6} - \frac{1}{4} + 0 - \frac{1}{4}) + \frac{1}{2}(\frac{1}{4} - \frac{1}{2}) = -\frac{7}{24}$ for applicant A.

For applicant B, if only $\{Y\}$ arrives in period 0, then A mistakenly bidding on Y imposes a positive externality on B, who receives an expected payoff of $\delta^2 = \frac{1}{4}$ instead of $\frac{\delta}{3} = \frac{1}{6}$. However, if applicant B mistakenly does not bid on Y in period 0, then B receives a payoff of $\frac{\delta^2}{3} = \frac{1}{12}$ instead of $\frac{\delta}{3} = \frac{1}{6}$. If $\{X, Y\}$ arrives in period 0, then B mistakenly bidding on X results in a payoff of $\frac{\delta^2}{3} = \frac{1}{12}$ instead of $\frac{\delta}{3} = \frac{1}{6}$. This leads to a first-order effect of $\frac{1}{2}(\frac{1}{4} - \frac{1}{6} + \frac{1}{12} - \frac{1}{6}) + \frac{1}{2}(\frac{1}{12} - \frac{1}{6}) = -\frac{1}{24}$ for applicant B. Relative to the US mechanism, the first-order loss for the lower-priority applicant is therefore substantially smaller in magnitude under the UK mechanism.

2.3.4 Key takeaways

Even in this stylized example, several patterns emerge that preview the richer simulation results that follow.

First, the US mechanism delivers higher total welfare under rationality. Batching in the

UK mechanism delays matching. The US mechanism also protects lower-priority applicants: by limiting how long A can hold out, it prevents A from crowding out B.

Second, the two mechanisms expose applicants to qualitatively different risks from error. In the US mechanism, the key danger is rejecting an offer that cannot be recalled. In the UK mechanism, the key danger is bidding on the wrong unit. Systematically bidding on too few units prolongs waiting; bidding on too many speeds matching and relaxes congestion. These structural differences mean the same error rate generates different welfare consequences across mechanisms.

Third, bias can improve welfare through externalities. When A mistakenly takes Y, this is a loss for A but frees X for B. Under the UK mechanism, this externality is strong enough that moderate choice error raises B's welfare above the rational benchmark.

The example does not capture all of the forces at play in our empirical results. For example, when simulating belief formation, applicants calibrate offer beliefs from historical data, creating a feedback loop between behavior and beliefs that generates coordination failures. Present focus, absent from the example, turns out to resolve this failure by inducing earlier acceptance—a channel that requires the richer environment of Section 3.

3 Simulating allocation mechanisms

We now turn to the paper's main exercise: comparing the welfare effects of the US and UK mechanisms. We combine administrative data on housing vacancies from a large London council with experimentally elicited preferences and biases from 587 public housing residents to simulate both mechanisms over a five-year horizon. The main components of the environment are the arrival process of applicants, the arrival process of houses, and the allocation mechanisms themselves (Section 3.1).

3.1 Setup

An allocation mechanism φ is a procedure which chooses which housing-unit b (or no housing-unit $b = 0$) to allocate each applicant a in each time period depending on the arrival process of housing-units, the arrival process of applicants, applicants' preferences $\{U_a(\cdot)\}_{a \in A}$, beliefs about being offered building-units and exogenous priority orderings $\{p(a)\}_{a \in A}$. It is visualized in Section 3.1.

We compare different allocation mechanisms φ. All mechanism comparisons hold fixed a common environment: the same applicant sample (described in Section 3.2), the Camden housing stock and arrival process (Section 3.3) and the baseline information environment

Illustrating Public Housing Allocation

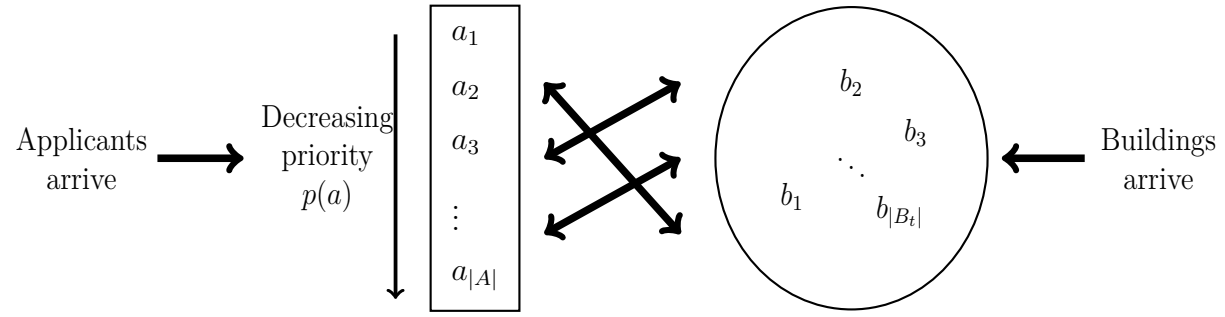

Notes: Applicants $a \in A$ are ordered in a waitlist by decreasing priority $p(a)$. Each period, a flow of houses arrive B^t. An allocation mechanism φ determines the allocation in each period of houses to applicants depending on applicants' priorities and preferences.

constructed from Camden's historical letting outcomes (see below). In other words, the simulations answer the question of what would happen if the respondents of our choice experiment, carrying beliefs formed from Camden-style historical data, were processed over a five-year horizon under the UK or US mechanism using the same underlying arrival process. We solve each applicant's Bellman equation (with and without bias) via backward induction, then simulate the allocation forward 100 times, reporting means and Monte-Carlo standard errors.

One key component of the simulations of allocation mechanisms is the specification of how applicants form beliefs about the probability they will be offered a particular housing unit.

Beliefs about being offered a unit. Since applicants do not observe the full state of the queue, they condition decisions on their own priority points and on beliefs derived from historical letting outcomes. They summarize competition for each unit b by a subjective belief

$$\theta_a(b) \coloneqq \Pr(a \text{ is offered } b)$$

which, in the UK mechanism, is interpreted as the probability of receiving an offer conditional on bidding on b. Applicants are assumed to condition their choices on these beliefs rather than on the full state of the waiting list.[5]

Let p_b^* be the random variable representing the priority score of the (hypothetical) applicant who receives unit b. Because allocation is priority-based, applicant a is offered b if and only

[5]The websites of the largest local authorities operating this mechanism in the UK almost all provided only the following information to applicants: which properties were available and data on historic lets.

if $p(a) \geq p_b^*$.[6] Therefore,

$$\theta_a(b) = \Pr(a \text{ is offered } b) = F_{p_b^*}(p(a)), \tag{1}$$

where $F_{p_b^*}$ denotes the distribution of recipient priority scores for unit type b. In our simulations, we estimate these distributions from historical letting outcomes. We also consider belief distortions by replacing θ with a probability-weighted transformation $w(\theta)$.

3.2 Applicants

We elicit preferences and behavioral biases from 587 public housing residents in England via an adaptive discrete choice experiment. This section describes the preference estimates, the bias measures and their parameterization, and the decision problems applicants solve under each mechanism.

3.2.1 Applicant preferences

We elicit preferences over housing characteristics from a sample of public housing residents in England, recruited through the platform Prolific. Table 1 shows some summary statistics on the demographics of the population ($N = 587$). 60 percent of our sample are women, 40 percent have at most a high-school degree and 40 percent are low income. We conduct a discrete choice experiment (DCE) in which applicants choose among 25 pairs of hypothetical housing units that differ along several key dimensions: rent, location, property type, and property ownership. Specifically, we implement a Bayesian Adaptive Choice Experiment (BACE), following Drake et al. (2025a,b), to obtain more precise estimates at the individual level by tailoring each choice set based on the respondent's previous answers. Starting from a uniform prior over preference parameters, the algorithm updates each respondent's posterior after every choice and selects the next pair of housing profiles to maximize expected information gain.

The DCE intentionally strips the choice task to controlled pairwise comparisons. The goal is not to emulate the emotional and cognitive pressures of a live housing offer — it is to isolate stable valuations over housing characteristics in a familiar domain. To the extent that temptation, time pressure, or salience distort real-world choices, those forces enter the framework through the separately measured behavioral parameters, not through the preference estimates.

[6] This follows directly from the fact that all mechanisms we consider are priority-based: if a has higher priority than the recipient of b, then a would have been offered b; and if a were offered b, then the recipient must have weakly lower priority than a.

We assume a utility function of the form

$$u_{i,j,t} = \delta_i^{t_j}(X_j\beta_i - r_j) + \nu_i\varepsilon_{i,j,t},$$

where δ_i is applicant i's monthly discount factor, X_j collects observable characteristics of housing unit j, β_i is a vector of individual-specific tastes, r_j denotes rental costs, and t_j denotes the number of months until unit j's availability. The building characteristics consist of indicators for being in the applicant's preferred location, property type (low-rise flat, high-rise flat, or house), and property ownership (council-owned, housing association-owned or private-rented). The idiosyncratic component $\varepsilon_{i,j,t}$ follows an i.i.d. type-I extreme-value distribution. The scale parameter ν_i governs the dispersion of this idiosyncratic component and therefore pins down respondent i's choice consistency, with larger ν_i implying noisier choices. In general, $\nu_i\varepsilon_{i,j,t}$ absorbs residual utility from any form of survey response noise. In the simulations, we map this to the reduced-form "decision noise" bias. This mapping treats within-task stochasticity as choice error rather than as welfare-relevant unobserved quality differences across realized units, which a DCE holds fixed by design.

As in standard discrete-choice models, we estimate preferences relative to omitted baseline categories for each attribute (non-preferred location, high-rise flat for property type, and private rental for property ownership). For the simulations, we additionally normalize each applicant's utility so that their outside option $oo(i)$, defined as the residence type they report living in during their most recent spell on a waitlist, yields zero utility.[7] We implement this normalization by redefining utility relative to the outside option as

$$\tilde{u}_{i,j,t} = u_{i,j,t} - u_{i,oo(i),t}$$

Because rent enters utility with coefficient -1, each element of β_i admits a willingness-to-pay interpretation as the rent increase (in the units of r_j) that leaves applicant i indifferent to gaining the corresponding housing attribute, holding timing fixed. Table 2 displays summary statistics for the estimated preferences. On average, respondents exhibit the strongest willingness to pay for residing in a house rather than a high-rise flat, and for residing in their preferred location. Respondents also value public-sector ownership (council or housing-association) relative to private rental, and low-rise flats relative to high-rise flats. The estimated monthly discount factors exhibit substantial dispersion in waiting costs

[7]Each respondent in the BACE data reports their outside-option rent $r_{oo(i)}$. However, the resulting distribution differs from the rent distribution in the Camden arrivals data, likely because the respondents do not necessarily live in Camden (see Figure 5). We therefore rescale reported values so that the two distributions share the same mean, and we impute rents that fall below the threshold of £20 per week. For more details, see Appendix C.1.

across applicants, and the estimated choice-consistency parameters also show considerable heterogeneity in decision noise.

Key takeaways: Two features of the estimates in Table 2 matter for the simulation. First, willingness-to-pay magnitudes are large relative to rent: the average WTP for preferred location is £37.5/week, roughly 25 percent of mean rent. Applicants care substantially about which house they receive. Second, the standard deviations of the WTP estimates are of similar magnitude to the means, indicating large preference heterogeneity across applicants. Together, these imply that how a mechanism matches applicants to units—not just how quickly it clears the queue—has first-order welfare consequences.

3.2.2 Behavioral biases

We complement the preference estimates with direct measures of behavioral biases from the same survey. This section describes the biases we elicit, how we parameterize them for simulation, and how each modifies the rational Bellman equations for each allocation mechanism.

We directly elicit a range of behavioral biases using hypothetical choices following Stango and Zinman (2023). Present bias is measured by choices between an immediate unhealthy snack and a delayed healthy alternative. Loss aversion is captured by willingness to accept small-stakes lotteries. Decision noise is measured by choice inconsistency in the BACE itself — the estimated scale parameter ν_i from the preference model. We additionally elicit biased beliefs about offer probabilities: we show respondents historical letting data from a hypothetical council and ask them to estimate their probability of receiving an offer, allowing us to estimate a probability weighting function $w(\theta)$ for each demographic cell (Appendix C.3.5). We also measure exponential growth bias, gambler's fallacy, non-belief in the law of large numbers, limited attention, and narrow bracketing. We focus on the first four biases because they map directly onto the allocation mechanisms: present bias and loss aversion shape the trade-off between accepting now and waiting for a better match, biased beliefs distort which units applicants bid on, and decision noise captures errors in evaluating multi-attribute housing options.

In our data, bias is prevalent: 82 percent of respondents exhibit loss aversion, 18 percent are present-focused, and the estimated probability weighting functions indicate systematically distorted offer beliefs. These rates are broadly comparable to those in the American Life Panel (Stango and Zinman, 2023), suggesting our sample is not unusually biased. The high prevalence of bias, combined with the large preference stakes documented above, motivates the simulation exercise that follows: if applicants care a lot about which house they get and

frequently err when choosing, the welfare consequences of mechanism design may depend critically on the interaction between choice restrictions and choice errors. Full details on elicitation and prevalence comparisons are in Appendix C.2.1.

Parameterization. The survey identifies *whether* each respondent exhibits a bias. To simulate allocation mechanisms, we also require the *extent* of each bias — for example, a β parameter rather than a binary indicator of present focus. We map binary survey responses to continuous parameters using a truncated-expectations approach. Assume the bias parameter follows a known distribution G in the population. Within a demographic cell, if a fraction q exhibit the bias, we assign those individuals $\hat{\beta} = \mathbb{E}[\beta \mid \beta \leq G^{-1}(q)]$ and the remainder $\hat{\beta} = \mathbb{E}[\beta \mid \beta \geq G^{-1}(q)]$. We select cell definitions by LASSO: for each bias, we regress the binary indicator on demographics and retain the covariates selected most frequently across outcomes. Gender, age group, and income group are each retained for the majority of biases, so we estimate parameters within cells of gender $\times$ age group $\times$ income group.

For present focus, we assume $G = U[0.7, 0.9]$ (Frederick et al., 2002; Bradford et al., 2019); for loss aversion, $G = U[1.82, 2.10]$ (Brown et al., 2024). The disaggregated model fits significantly better than a single parameter for loss aversion ($\chi^2 = 13.7$, $p = 0.01$) but not for present focus. Decision noise ν_i is estimated individually from any choice inconsistencies in the main BACE. For biased beliefs, we fit a probability weighting function $w(\theta)$ for each demographic cell using the elicited offer probabilities described above; the disaggregated model fits significantly better than a single function ($F = 3.12$, $p < 0.001$). Full results are in Table 5 and Figure 1; Appendix C.3 details the calibration procedure for each bias, including the belief elicitation design and cell-level estimates.

The Stango-Zinman instruments are calibrated on a large representative sample, which supports external validity. However, mapping binary survey indicators to continuous bias parameters requires distributional assumptions (G). Directional conclusions — which biases help or hurt which mechanism — depend on the sign of the bias effect and are likely robust even if the exact magnitudes are sensitive to model specification.

3.2.3 Decision problems

We now formalize how applicants make decisions under each mechanism. The Bellman equations take as given preferences $U_a(b, t)$ estimated in Section 3.2.1 and biases estimated in Section 3.2.2.

US mechanism. Each period, at most one unit b arrives and is offered sequentially down the waitlist. A rational applicant with r rejections used solves:

$$V^t(b,r) = \underbrace{\theta_a(b)}_{\text{offered } b} \max\left\{ \underbrace{U_a(b,t)}_{\text{accept}},\ \underbrace{u_a(0,t) + \delta\, \mathbb{E}_\pi\Big[V^{t+1}(b', r+1)\,\Big|\, b\Big]}_{\text{reject}} \right\} + \underbrace{\Big(1-\theta_a(b)\Big)}_{\text{not offered}} \cdot \Big[u_a(0,t) + \delta\, \mathbb{E}_\pi\Big[V^{t+1}(b', r)\,\Big|\, b\Big]\Big] \tag{2}$$

with terminal condition $V^t(b,k) = u_a(0,t)$ once k rejections are exhausted. The applicant accepts b if and only if the immediate payoff $U_a(b,t)$ exceeds the continuation value from waiting.

UK mechanism. Each period, a batch B^t of units arrives and each applicant may bid on at most k units, limiting the number of potential acceptances per cycle. With $k = 3$, rank bids so that $U_a(b_1,t) \geq U_a(b_2,t) \geq U_a(b_3,t)$ without loss of generality. A rational applicant's value function satisfies:

$$\begin{aligned} V_a^t(B^t) = \max_{\substack{b_1 \in B^t,\\ b_2 \in B^t \setminus \{b_1\},\\ b_3 \in B^t \setminus \{b_1, b_2\}}} \Bigg\{ & \underbrace{\theta_a(b_1)\, U_a(b_1,t)}_{\text{1st bid wins}} + \underbrace{(1-\theta_a(b_1))\, \theta_a(b_2)\, U_a(b_2,t)}_{\text{2nd bid wins}} \\ & + \underbrace{(1-\theta_a(b_1))(1-\theta_a(b_2))\, \theta_a(b_3)\, U_a(b_3,t)}_{\text{3rd bid wins}} \\ & + \underbrace{(1-\theta_a(b_1))(1-\theta_a(b_2))(1-\theta_a(b_3))}_{\text{no bid wins}} \Big[u_a(0,t) + \delta\, \mathbb{E}_\pi\Big[V_a^{t+1}(\tilde{B}^{t+1})\Big]\Big] \Bigg\} \end{aligned} \tag{3}$$

where $b_i = 0$ represents passing on the i-th bid, in which case all subsequent bids are also zero. If no bid succeeds, the applicant receives the continuation value from waiting.

The two mechanisms differ in two key ways. First, the US mechanism penalizes waiting through the rejection limit, while the UK mechanism allows unrestricted passing. Second, the UK mechanism batches arrivals, giving applicants a choice set each period rather than a single take-it-or-leave-it offer.

Biased Bellman equations. Each bias modifies the rational decision problem in a specific way. We describe how each enters the Bellman equations; in each case, we show the US mechanism and note that the UK mechanism is modified analogously.

1. **Decision noise:** recall that utility is modeled as:

$$u_{i,j,t} = \delta_i^{t_j}(X_j\beta_i - r_j) + \underbrace{\nu_i \varepsilon_{i,j}}_{\text{Decision noise}}$$

where utility without noise is given by the first component only.

2. **Biased-beliefs (probability weighting):** we estimate a probability weighting function $w(\theta)$ capturing biased-beliefs about the probability of offer θ. Therefore, we assume applicants solve Bellman equations using these weighted beliefs. For example, in the US mechanism, Equation (2) becomes:

$$\begin{aligned} V^t(b,r) = \underbrace{w\Big(\theta_a(b)\Big)}_{\text{offered } b} \max\Bigg\{ \underbrace{U_a(b,t)}_{\text{accept}}, \underbrace{u_a(0,t) + \delta \mathbb{E}_\pi\Big[V^{t+1}(b', r+1) \,\Big|\, b\Big]}_{\text{reject}} \Bigg\} \\ + \underbrace{\Big(1 - w\Big(\theta_a(b)\Big)\Big)}_{\text{not offered}} \cdot \Big[u_a(0,t) + \delta \mathbb{E}_\pi\Big[V^{t+1}(b', r) \,\Big|\, b\Big]\Big] \end{aligned} \tag{4}$$

where r denotes the number of rejections already used, and after k rejections the applicant exits to the outside option:

$$V^t(b,k) = u_a(0,t). \tag{5}$$

The UK mechanism's Bellman equation incorporates $w(\theta)$ similarly.

3. **Present focus (See e.g. Harris and Laibson, 2001):** We assume that applicants are fully naive about their present focus. Therefore, in each mechanism, W is the value function which discounts with present bias, and V is the exponentially discounting value function. For example, in the US mechanism:

$$\begin{aligned} W^t(b,r) = \underbrace{\theta_a(b)}_{\text{offered } b} \max\Bigg\{ \underbrace{U_a(b,t)}_{\text{accept}}, \underbrace{u_a(0,t) + \beta\delta \mathbb{E}_\pi\Big[V^{t+1}(b', r+1) \,\Big|\, b\Big]}_{\text{reject}} \Bigg\} \\ + \underbrace{(1 - \theta_a(b))}_{\text{not offered}} \cdot \Big[u_a(0,t) + \beta\delta \mathbb{E}_\pi\Big[V^{t+1}(b', r) \,\Big|\, b\Big]\Big] \end{aligned} \tag{6}$$

where

$$V^t(b,r) = \underbrace{\theta_a(b)}_{\text{offered } b} \max\left\{\underbrace{U_a(b,t)}_{\text{accept}}, \underbrace{u_a(0,t) + \delta\mathbb{E}_\pi\left[V^{t+1}(b', r+1)\,\middle|\, b\right]}_{\text{reject}}\right\}$$
$$+ \underbrace{(1-\theta_a(b))}_{\text{not offered}} \cdot \left[u_a(0,t) + \delta\mathbb{E}_\pi\left[V^{t+1}(b', r)\,\middle|\, b\right]\right] \tag{7}$$

The UK mechanism's Bellman equation incorporates present focus similarly. In each case, β discounts everything from the next period onward, regardless of how long a period is in calendar time.

4. **Loss aversion:** we follow Andersen et al. (2022) and model reference dependence over prices (here, rent). Therefore: the utility of house j relative to i's outside option $oo(i)$ is:

$$u_{i,j,t} - u_{i,oo(i),t} = \delta_i^{t_j}\left(X_j - X_{oo(i)}\right)\beta_i - \begin{cases} \delta_i^{t_j}\left(r_j - r_{oo(i)}\right) & \text{if } r_{oo(i)} \geq r_j \\ \delta_i^{t_j} \underbrace{\lambda}_{\text{Loss Aversion}} \left(r_j - r_{oo(i)}\right) & \text{if } r_{oo(i)} < r_j \end{cases}$$

3.3 Building-units

In this section we describe our data on building arrivals, observed in administrative data from Camden Council in London, England. To contextualize our analysis, we supplement these data with publicly available council-level data from the UK government[8] as well as the 2019 UK Housing Survey Data.[9]

We use administrative arrivals data from Camden, a large London council. All weekly advertisements of available social housing in Camden were publicly available for the time period Jan 2016–Feb 2021. Figure 2a shows what an individual listing in an advertisement from Camden looks like. According to the Local Authority Housing Data, Camden has a slightly longer waitlist and a much larger housing stock than the median local authority (Figure 3). On average, around 20 units are available each week (Figure 4). Table 6 displays some summary statistics about the arrivals each week. Camden employs a "Choice-Based Lettings" allocation mechanism with applicants allowed to bid on at most 3 units (i.e. UK-3).

The Camden arrivals data contain more information about houses than we have preferences estimated for in the BACE data. For example, we see whether an apartment has lift access in the Camden data although do not estimate each person's utility of having a lift in the

[8] https://www.gov.uk/government/collections/local-authority-housing-data
[9] https://beta.ukdataservice.ac.uk/datacatalogue/studies/study?id=8920

BACE data due to computational constraints. Therefore, our first step is to encode each arrival in the Camden data as a house-type on which we have preference data. Through this process, we end up with 432 house-types which vary based on the characteristics, as in the BACE data (preferred location, property type, property ownership and rent).

In addition to information about all arrivals during the time period, the data also contain information about the priority points of the winning applicants for previous ads (see Figure 2b). We use this information to calibrate the applicant beliefs about probability of offer, following Equation (1). Specifically, we observe the priority points for the winners of 4,193 past advertisements (mean = 443, sd = 306). We use this data to fit a truncated normal (min = 0, max = 2670) distribution for each of the 432 house-types, yielding $\hat{F}_{p_b^*} \forall b$. Each applicant then evaluates this function at their specific priority points to calculate $\theta_a(b)$. Consistent with the steady-state interpretation of Equation (1), we treat these $\hat{F}_{p_b^*}$ as time-invariant and do not update them with the simulated evolution of the queue.

Finally, for each successful let where the priority points of the winning applicant are displayed, there is also information on how *many* people bid on that house. We assume a house is in a generally-preferred location if, on average, it receives over 200 bids in the arrivals data (corresponding to the median number of bids).[10]

3.4 Allocation mechanisms

We simulate public housing allocation to a hypothetical applicant pool constructed based on the 587 people we observe in the BACE data.

3.4.1 Applicants

We do not observe any information about the arrival process of applicants. Therefore, we assume that applicants arrive at the same rate as houses, and copies are drawn uniformly from the original waitlist. The waiting list is ordered by priority points and, in case of ties, by applicants' arrival times. This conservative assumption focuses our analysis on how mechanisms differ in their treatment of choice restrictions rather than on supply-demand imbalances.

The BACE data contains individuals who reside in local authorities in the UK which do not necessarily operate a points-based system.[11] Therefore, to implement the allocation as it would be in Camden, we need to assign a priority $p(a)$ for each a in the data. We present our

[10] This assumes that all applicants agree which are the good locations, although may have heterogeneous preferences for how much they trade off location with other house characteristics. For more information on the buildings data see Section 3.3.

[11] Other ways of prioritizing individuals include bands, a coarser version of points.

method for proxying for $p(a)$ in Appendix C.2. Only 26 percent of participants in the BACE report enough information for us to proxy their priority points. Otherwise, we draw their quantile $q(a)$ randomly according to a uniform distribution on $[0, 1]$. This also contributes to the Monte-Carlo standard errors across simulation runs.

3.4.2 Building-units

We simulate the arrival process of houses as a Poisson process with arrival rate $\lambda = 18.7$, which is the mean arrival rate per week in the Camden data. Conditional on a house arriving, its type is drawn from the empirical distribution of types. This modelling assumption assumes stationarity of the arrival process of houses. We denote this stationary arrival process by π, and expectations $\mathbb{E}_\pi[\cdot]$ in the Bellman equations are taken with respect to this process. Figure 4 shows the number of arrivals per week in the arrivals data with associated mean.

In Camden, applicants are ineligible for properties that do not match their stated requirements for bedrooms (Camden Council, 2018). Therefore, in both US and UK mechanisms we set $\theta_a(b) \equiv 0$ if b does not match a's required bedrooms, and in the simulation, the US mechanism skips applicants whose bedrooms do not match when making sequential offers.

3.4.3 Other assumptions

UK mechanism We abstract from viewings by assuming that applicants do not reject units they bid on when offered in the UK mechanism. We make two further simplifying assumptions for computational feasibility. First, due to batching, the state variable in the UK mechanism (B in Equation (3)) can take any one of 2^{432} values. Instead of enumerating all possibilities, we assume that applicants perform a Monte-Carlo simulation of 100 different scenarios for the set of available houses next period, according to the empirical distribution observed in the arrivals data.

Second, the optimal bidding policy function in the UK mechanism is computationally taxing to compute each period. Suppose that 18 houses are available in a given period. Ranking each potential set of bids would require ranking $\binom{18}{3} = 816$ different possibilities for each applicant. Instead, we assume that applicants employ the following heuristic when choosing their bids. Applicants' bids are the top-3 out of the following set: $\{\theta_a(b) \cdot U_a(b)\}_{b \in B} \cup \left\{u_a(0,t) + \delta \mathbb{E}_\pi\left[V_a^{t+1}(\tilde{B}^{t+1})\right]\right\}$. This heuristic essentially picks bids "independently" according to the index $\theta_a(b) \cdot U_a(b)$, where it is also possible not to use one of your bids, and not bidding is associated with the passing-value $u_a(0,t) + \delta \mathbb{E}_\pi\left[V_a^{t+1}(\tilde{B}^{t+1})\right]$. In theory, this could result in a different set of bids to the true optimal policy function if applicants ideally prefer to have a risky and safe option in their portfolio if choosing bids as a set, but when choosing

bids one at a time unintentionally choose two risky options. The heuristic matches the true optimum in our simulations.

US Mechanism The US mechanism is designed to be implemented practically in continuous time (one arrival per period). We implement this by choosing our time-period to be short enough that the likelihood that two houses arrive per period is less than 0.01. This amounts to the mechanism operating at the frequency of approximately every 1.5 hours. If a house is not allocated to an applicant within the processing period, then it goes to waste (which represents it being allocated to an applicant pool that we do not observe, or staying void). Aside from the difference in batching, we simulate both allocation mechanisms on the same arrival process of applicants and houses.

3.4.4 Comparison metrics

We compare mechanisms on two aggregations of ex-post simulated welfare, as well as some report some other outcomes to understand mechanisms. Welfare is defined as follows.

Definition 1. Individual Welfare: An allocation mechanism φ defines an allocation $\mu^t(a)$ for each applicant a and time t. Therefore, individual a's lifetime utility under mechanism φ is given by:

$$U_a(\varphi) = \sum_{t=0}^{\infty} u_a(\mu^t(a), t) \tag{8}$$

Using Definition 1, we can compare the US and UK mechanisms based on ex-post social welfare functions of the form W_g defined as in Definition 2.

Definition 2. Social Welfare: Every function $g : A \to \mathbb{R}$ defines the welfare weights of a social welfare function Ψ_g as follows:

$$\Psi_g(\varphi) = \sum_{a \in A} g(a) \cdot U_a(\varphi) \tag{9}$$

We are particularly interested in two specific cases:

1. **Average Welfare**: $g(a) = \frac{1}{|A|}$ for all $a \in A$.
2. **Priority-weighted Welfare**: $g(a) \propto p(a)$ for all $a \in A$.

We use administrative priority points as a weighting device that tilts welfare comparisons toward applicants assessed to have greater housing need. This allows us to conduct the exercise of asking how mechanism rankings change when the planner puts greater welfare weight on higher-need households.

We report welfare effects in terms of NPVs: averages across applicants of the lump-sum payment that makes each applicant indifferent between receiving this payment at the moment they enter the waitlist and staying where they are vs participating in the mechanism and potentially getting allocated a new house.[12]

We evaluate allocation mechanisms primarily based on their effects on applicant welfare, which is the object most directly affected by choice restrictions. A comprehensive policy evaluation would also require incorporating fiscal costs borne by local governments, which we do not attempt to estimate in this paper; however, we report the plausibly important determinants of these costs. In particular, we track the number of void properties (empty housing units that do not generate rental income) and the size of the waiting list backlog. Both are commonly cited drivers of fiscal pressure in public housing systems, since prolonged vacancies reduce rental income and large waiting lists are often associated with greater reliance on temporary accommodations, which incur substantial costs for local housing councils.[13]

4 Simulated welfare effects

4.1 Welfare effects under rationality

Table 7 displays the main results of the simulations under full rationality. We estimate average welfare effects of $W_{\text{US-2}} \approx$ £2,305 and $W_{\text{UK-3}} \approx$ £2,150 and priority-weighted welfare effects of $W_{\text{US-2}} \approx$ £2,427 and $W_{\text{UK-3}} \approx$ £2,595. Although both mechanisms respect the same priority ordering, they translate that ordering into outcomes differently. The US-2 mechanism performs better when welfare is averaged equally across applicants, whereas the UK-3 mechanism performs better if higher-need households receive greater welfare weight. Across simulations, we can reject that $W_{\text{US-2}} = W_{\text{UK-3}}$ using Monte-Carlo standard errors. In terms of magnitudes, Moving from UK-3 to US-2 generates a NPV of £150 on average, an increase of 7 percent.

Interpretation of magnitudes: The average welfare effects are equivalent to yearly transfers of approximately £882 (UK-3) and £945 (US-2), or roughly 11–12 percent of yearly rent and 15–16 percent of yearly income (Table 1).[14] These magnitudes are comparable to the surplus gains from neighborhood quality in Cook et al. (2023) (£1,200/year). This

[12] To convert NPV estimates to yearly transfers, we multiply by 41 percent, derived from an average yearly discount factor of 0.804 and an assumed tenancy of three years: $\frac{1-0.804}{1-0.804^3} \approx 0.41$.

[13] The mechanisms also differ in processing frequency. Although operating allocation mechanisms more frequently may entail higher administrative costs, this channel is likely to be limited in magnitude with modern digital systems for allocation.

[14] Using the conversion rule in Section 3.4.4.

fact that they are smaller likely reflects the UK housing context: unlike many US settings, where public housing applicants may face outright homelessness, most applicants in our sample already occupy temporary accommodation or private rentals. The welfare gain from assignment therefore captures the improvement in housing quality rather than the full value of shelter. The 7 percent welfare difference between mechanisms is comparable to the range of counterfactual comparisons between mechanisms in Waldinger (2021), whose level estimates are larger because welfare is measured relative to no housing assignment.

Drivers of the rationality effect: Why is the US-2 mechanism more effective than UK-3 on average under full rationality? A key mechanism can be seen in Figure 6b: the US mechanism clears the waitlist while the UK mechanism does not. In general, more congestion is detrimental to welfare, as people have to wait longer before being allocated.

Figure 6c shows the welfare effects of each allocation mechanism, split by deciles of the priority points distribution. While the UK mechanism does worse on average, it performs better for the highest priority applicants and thus is superior to the US mechanism when using a priority-weighted social welfare function. Since the UK mechanism does not penalize waiting, and gives substantial choice to each applicant. This is good for those at the top of the waitlist, who can bid on their most preferred house, or wait if nothing good is available. This comes at the cost of increasing backlog. The main benefit of the US mechanism, is that it is *fast*—higher priority applicants get only moderately better allocations because if they receive low quality offers, they are limited in their ability to wait for something better. This helps clear the waitlist.

The US and UK mechanisms differ along two key dimensions that can affect both congestion and the extent to which higher-priority applicants can translate rank into better outcomes: (i) the US mechanism penalizes waiting through rejection limits, and (ii) the UK mechanism batches arrivals and lets applicants choose among the units in each batch. To isolate the contribution of these two features under rationality, we also simulate US-∞, which is identical to the US mechanism except that applicants face no penalty for rejecting offers. Further note that US-∞ is the same as UK-k **without** batching and therefore serves as a useful middle point to understand the different forces.[15]

Appendix Figure 2a shows that batching contributes approximately 25 percent to congestion, while penalizing waiting contributes the rest. Appendix Figure 2b shows that the UK mechanism's match-quality advantages for higher-priority households are primarily due to allowing for waiting, rather than batching.

[15]If there is no batching, then only one house arrives per period and it does not matter how many bids applicants have in the UK mechanism.

4.2 Robustness to behavioral bias

How does each mechanism's welfare depend on behavioral bias? We consider four forms: belief bias, decision noise, present focus, and loss aversion. Figures 7a and 7b summarize the results. The UK mechanism exhibits substantially more variance across bias types, indicating greater sensitivity. Under rationality, switching from UK-3 to US-2 raises average welfare by 7.2 percent. This gap widens or narrows—and in one case reverses—depending on the bias.

Belief bias. Belief bias is the most damaging bias for the UK mechanism (Figure 7a), widening the UK–US gap to 36.3 percent. Belief bias harms the UK mechanism not by reducing the maximum number bids, but by misallocating them to units the applicant is unlikely to receive. Congestion rises to approximately 4,500 applicants (Figure 8a), with losses concentrated at the bottom of the priority distribution (Figure 8b). Figure 1 shows that the estimated probability weighting functions generate optimism on average. Most beliefs about offer probabilities are small, and all misperceptions of small probabilities are biased upward. This causes arrivals to go to "waste" (not allocated to anybody, see Figure 8c). All in all, this means that applicants do not bid less—average number of bids stays around 1.65 (Figure 8d)—but bid on the *wrong* properties (resulting in more waste).

The US mechanism is less affected. In the biased-belief Bellman equation (Equation (4)), $w(\theta_a(b))$ enters only through the offer probability, leaving the accept/reject comparison undistorted apart from through the raised continuation value through optimistic beliefs. Overall, belief-bias produces few mistakes (Figures 9a and 9b).

Loss aversion. Loss aversion narrows the UK–US gap to 3.9 percent. UK average welfare reduces to approximately £2,145 (Figure 7a) but is statistically indistinguishable to welfare effects under rationality. UK congestion increases to approximately 4,100 (Figure 8a), but through a different channel than belief bias. Reference dependence over rent raises the perceived value of the status quo, reducing average number of bids to 1.52 (Figure 8d). The contrast with belief bias—which generates worse congestion despite similar frequency of bidding—shows that misallocation of bids drives waitlist accumulation more than reduced bidding.

In the US mechanism, loss aversion distorts both margins: the λ penalty raises the acceptance threshold and inflates the continuation value. The majority of the welfare costs of loss aversion are from false rejections—rejecting offers is extremely costly for the US mechanism due to the reject cap (Figures 9a and 9b). This problem is especially pronounced with loss aversion because it makes applicants prefer the status quo.

Present focus. Present focus is the only bias that reverses the mechanism ranking: UK-3 average welfare rises to over £3,000 and outperforms US-2 by 31.5 percent under present focus (Figure 7a). The welfare of the highest-priority decile is very slightly smaller under present focus, however lower priority deciles do considerably better. Present-focused applicants accept sooner and bid more actively (1.75 per period, Figure 8d), clearing the waitlist to $\approx 2,500$ applicants. We interpret this improvement as follows. Offer beliefs are calibrated from historical data. High-priority applicants wait selectively and rarely accept, so observed outcomes underrepresent high-priority winners. As a result, optimistic beliefs and waiting persists in equilibrium. Present-focused applicants break this cycle by discounting continuation values more heavily, accepting sooner and bidding more actively, which reduces congestion and improves outcomes across the priority distribution. Therefore, by accepting sooner, present focussed applicants accidentally solve a coordination problem and the resulting lower congestion benefits almost everyone.

In the US mechanism, present focus has small overall effects (loss in average welfare of £20). Present-focus unilaterally makes applicants accept more frequently - there are no false rejections (Figures 9a and 9b). The pattern described above for the UK mechanism does not play out in the US because continuation values are mainly determined by the rejection cap as opposed to beliefs about future offers.

Decision noise. Finally, decision noise closes the UK–US gap to 5.8 percent. Decision noise perturbs flow utility symmetrically via the $\nu_i \varepsilon_{i,j}$ term, shocking both acceptance and rejection decisions. Under UK-3, decision noise statistically insignificantly reduces average welfare but improves priority-weighted welfare (Figures 7a and 7b): high-priority applicants occasionally bid on non-preferred units, imposing internalities on themselves but generating positive externalities by freeing competition for desirable units (Figure 8b). The gains accrue primarily to the middle of the priority distribution, while lower-priority applicants also lose, so average welfare falls; priority weighting discounts these losses and amplifies the middle's gains.

In the US mechanism, decision noise generates both false acceptances and false rejections (Figures 9a and 9b). Its welfare effect is intermediate between belief bias (which inflates only the continuation value) and loss aversion (which reduces both acceptance utility and the continuation value).

Synthesis. The UK mechanism is more sensitive to bias because high-priority applicants' decisions have larger downstream effects. Unrestricted passing and batched choice mean that bidding distortions cascade into waitlist accumulation, harming lower-priority applicants. But

this cuts both ways: internalities imposed by biased high-priority applicants can generate positive externalities—and in the case of present focus, solve a coordination problem that improves welfare for everyone. The US mechanism bounds this spillover through the rejection cap. Biases generate false rejections that impose individual losses, but rejected units flow to lower-priority applicants.

5 Conclusion

This paper examines how different forms of choice restrictions affect welfare in dynamic allocation mechanisms. The US mechanism limits the number of rejections applicants may make, while the UK mechanism restricts which units applicants may bid on. We show that these restrictions create different trade-offs between prioritization—allowing high-priority applicants to wait for better matches—and congestion—the accumulation of waitlists that harm lower-priority applicants. Under rationality, the US mechanism reduces congestion and delivers higher average welfare, while the UK mechanism provides stronger match-quality advantages for higher-priority households.

The UK mechanism is more sensitive to bias because high-priority applicants' decisions have larger downstream effects. Unrestricted passing and batched choice mean that distortions in bidding cascade into waitlist accumulation, harming lower-priority applicants. But this cuts both ways: internalities imposed by biased high-priority applicants can generate positive externalities. Present focus, which reduces strategic waiting, solves a coordination problem and improves welfare for all priority groups. Decision noise generates positive externalities for lower-priority applicants but does not improve average welfare. The US mechanism bounds these spillovers through the rejection cap. Biases generate false rejections that impose individual losses, but rejected units flow to lower-priority applicants, limiting aggregate damage.

The US mechanism performs better when the council places most weight on average welfare, in part because it clears the queue more quickly and because, when applicants mistakenly reject offers they would otherwise accept (e.g., due to belief bias or loss aversion), those rejections pass units down the queue to lower-priority households. The UK mechanism tends to perform better when the council places greater weight on higher-need households because it lets those households convert their priority into better matches by waiting selectively and targeting preferred units. Effective mechanism choice therefore requires knowing the distribution of biases in the applicant pool.

Our approach extends to other dynamic allocation problems including organ transplantation, adoption, and school choice. Several questions remain open: optimal calibration of the

choice restriction parameters for a given bias profile, whether information interventions that disclose offer probabilities can recover welfare losses from biased beliefs, how endogenous belief updating interacts with changed behavior, and the welfare consequences of multiple simultaneous biases. More broadly, identifying settings where behavioral biases generate beneficial externalities through coordination effects is a promising direction for mechanism design.

References

Action on Empty Homes, "Empty Homes Data," 2024. Accessed December 2024. 1

Agarwal, Nikhil, Itai Ashlagi, Michael Rees, Paulo Somaini, and Daniel Waldinger, "Equilibrium Allocations under Alternative Waitlist Designs: Evidence from Deceased Donor Kidneys Kidneys," *Econometrica*, 2021. 3

Ameriks, John, Joseph Briggs, Andrew Caplin, Matthew D Shapiro, and Christopher Tonetti, "Long-term-care utility and late-in-life saving," *Journal of Political Economy*, 2020, *128* (6), 2375–2451. 4

Andersen, Steffen, Cristian Badarinza, Lu Liu, Julie Marx, and Tarun Ramadorai, "Reference dependence in the housing market," *American Economic Review*, 2022, *112* (10), 3398–3440. 19

Andrew, Alison and Abi Adams, "Revealed beliefs and the marriage market return to education," *The Quarterly Journal of Economics*, 2025, *140* (3), 2107–2162. 4

Arnosti, Nick and Peng Shi, "Design of Lotteries and Wait-Lists for Affordable Housing Allocation," *Management Science*, 2020. 3

Arteaga, Felipe, Adam J Kapor, Christopher A Neilson, and Seth D Zimmerman, "Smart matching platforms and heterogeneous beliefs in centralized school choice," *The Quarterly Journal of Economics*, 2022, *137* (3), 1791–1848. 4

Bloch, Francis and David Cantala, "Dynamic assignment of objects to queuing agents," *American Economic Journal: Microeconomics*, 2017, *9* (1), 88–122. 3

Bradford, W David, Paul Dolan, and Matteo M Galizzi, "Looking ahead: Subjective time perception and individual discounting," *Journal of Risk and Uncertainty*, 2019, *58* (1), 43–69. 16, 44

Brown, Alexander L, Taisuke Imai, Ferdinand M Vieider, and Colin F Camerer, "Meta-analysis of empirical estimates of loss aversion," *Journal of Economic Literature*, 2024, *62* (2), 485–516. 16, 44

Brown, Tim and Peter King, "The power to choose: effective choice and housing policy," *European Journal of Housing Policy*, 2005, *5* (1), 59–97. 1

Camden Council, "Housing Allocation Scheme 2018: Full Scheme," August 2018. London: Camden Council. 21

Cook, Cody, Pearl Z Li, and Ariel J Binder, *Where to build affordable housing?: Evaluating the tradeoffs of location*, US Census Bureau, Center for Economic Studies

Rochester, NY, 2023. 23

Drake, Marshall, Fernando Payró, Neil Thakral, and Linh T Tô, "Bayesian Adaptive Choice Experiments," *Mimeo*, 2025. 2, 13

_ , Neil Thakral, and Linh T Tô, "Wage Differentials and the Price of Workplace Flexibility," *Mimeo*, 2025. 13

Dreyfuss, Bnaya, Ofer Glicksohn, Ori Heffetz, and Assaf Romm, "Deferred acceptance with news utility," *Management Science*, forthcoming. 4

_ , Ori Heffetz, and Matthew Rabin, "Expectations-Based Loss Aversion May Help Explain Seemingly Dominated Choices in Strategy-Proof," *American Economic Journal: Microeconomics (forthcoming)*, 2021. 4

Fabre, Anaïs, Tomás Larroucau, Manuel Martınez, Christopher Neilson, and Ignacio Rios, "Application Mistakes and Information Frictions in College Admissions," Technical Report, Working paper 2023. 4

Frederick, Shane, George Loewenstein, and Ted O'donoghue, "Time discounting and time preference: A critical review," *Journal of economic literature*, 2002, *40* (2), 351–401. 16, 44

Harris, Christopher and David Laibson, "Dynamic choices of hyperbolic consumers," *Econometrica*, 2001, *69* (4), 935–957. 18

Hassidim, Avinatan, Assaf Romm, and Ran I Shorrer, "The limits of incentives in economic matching procedures," *Management Science*, 2021, *67* (2), 951–963. 4

_ , Déborah Marciano, Assaf Romm, and Ran I Shorrer, "The mechanism is truthful, why aren't you?," *American Economic Review*, 2017, *107* (5), 220–224. 4

Idoux, Clemence, "Integrating new york city schools: The role of admission criteria and family preferences," *Mimeo*, 2023. 4

Johnston, Sarah, Yifei Liu, and Chenyu Yang, "An empirical analysis of the interconnection queue," Technical Report, National Bureau of Economic Research 2023. 3

Kapor, Adam J, Christopher A Neilson, and Seth D Zimmerman, "Heterogeneous beliefs and school choice mechanisms," *American Economic Review*, 2020, *110* (5), 1274–1315. 4

Kullberg, Jeanet, "Consumers' responses to choice-based letting mechanisms," *Housing studies*, 2002, *17* (4), 549–579. 1

Leshno, Jacob D, "Dynamic matching in overloaded waiting lists," *American Economic Review*, 2022, *112* (12), 3876–3910. 3

Marsh, Alex, "The inexorable rise of the rational consumer? The Blair government and the reshaping of social housing," *European Journal of Housing Policy*, 2004, *4* (2), 185–207. 1

Meisner, Vincent and Jonas Von Wangenheim, "Loss aversion in strategy-proof school-choice mechanisms," *Journal of Economic Theory*, 2023, *207*, 105588. 4

Ministry of Housing, Communities & Local Government, "Statutory Homelessness in England: January to March 2025," Technical Report, UK Government July 2025. Accredited official statistics. 1

Murra-Anton, Zeky and Neil Thakral, "The Public Housing Allocation Problem," *Mimeo*, 2025. 3

Rees-Jones, Alex, "Suboptimal behavior in strategy-proof mechanisms: Evidence from the residency match," *Games and Economic Behavior*, 2018, *108*, 317–330. 4

_ **and Samuel Skowronek**, "An experimental investigation of preference misrepresentation in the residency match," *Proceedings of the National Academy of Sciences*, 2018, *115* (45), 11471–11476. 4

Schummer, James, "Influencing waiting lists," *Journal of Economic Theory*, 2021, *195*, 105263. 3

Shorrer, Ran I and Sándor Sóvágó, "Dominated choices in a strategically simple college admissions environment," *Journal of Political Economy Microeconomics*, 2023, *1* (4), 781–807. 4

_ **and** _ , "Dominated choices under deferred acceptance mechanism: The effect of admission selectivity," *Games and Economic Behavior*, 2024, *144*, 167–182. 4

Stango, Victor and Jonathan Zinman, "We are all behavioural, more, or less: A taxonomy of consumer decision-making," *The Review of Economic Studies*, 2023, *90* (3), 1470–1498. 2, 4, 15, 42

Thakral, Neil, "Matching with stochastic arrival," *AEA Papers and Proceedings*, 2019, *109*, 209–12. 3

Verdier, Valentin and Carson Reeling, "Welfare effects of dynamic matching: An empirical analysis," *The Review of Economic Studies*, 2020. 3

Waldinger, Daniel, "Targeting in-kind transfers through market design: A revealed preference analysis of public housing allocation," *American Economic Review*, 2021, *111* (8), 2660–96. 3, 24

Wang, Ao, Shaoda Wang, and Xiaoyang Ye, "Option Value Neglect: Evidence from Centralized College Admission," Technical Report, National Bureau of Economic Research 2025. 4

Figure 1: Probability weighting Functions by demographic cell

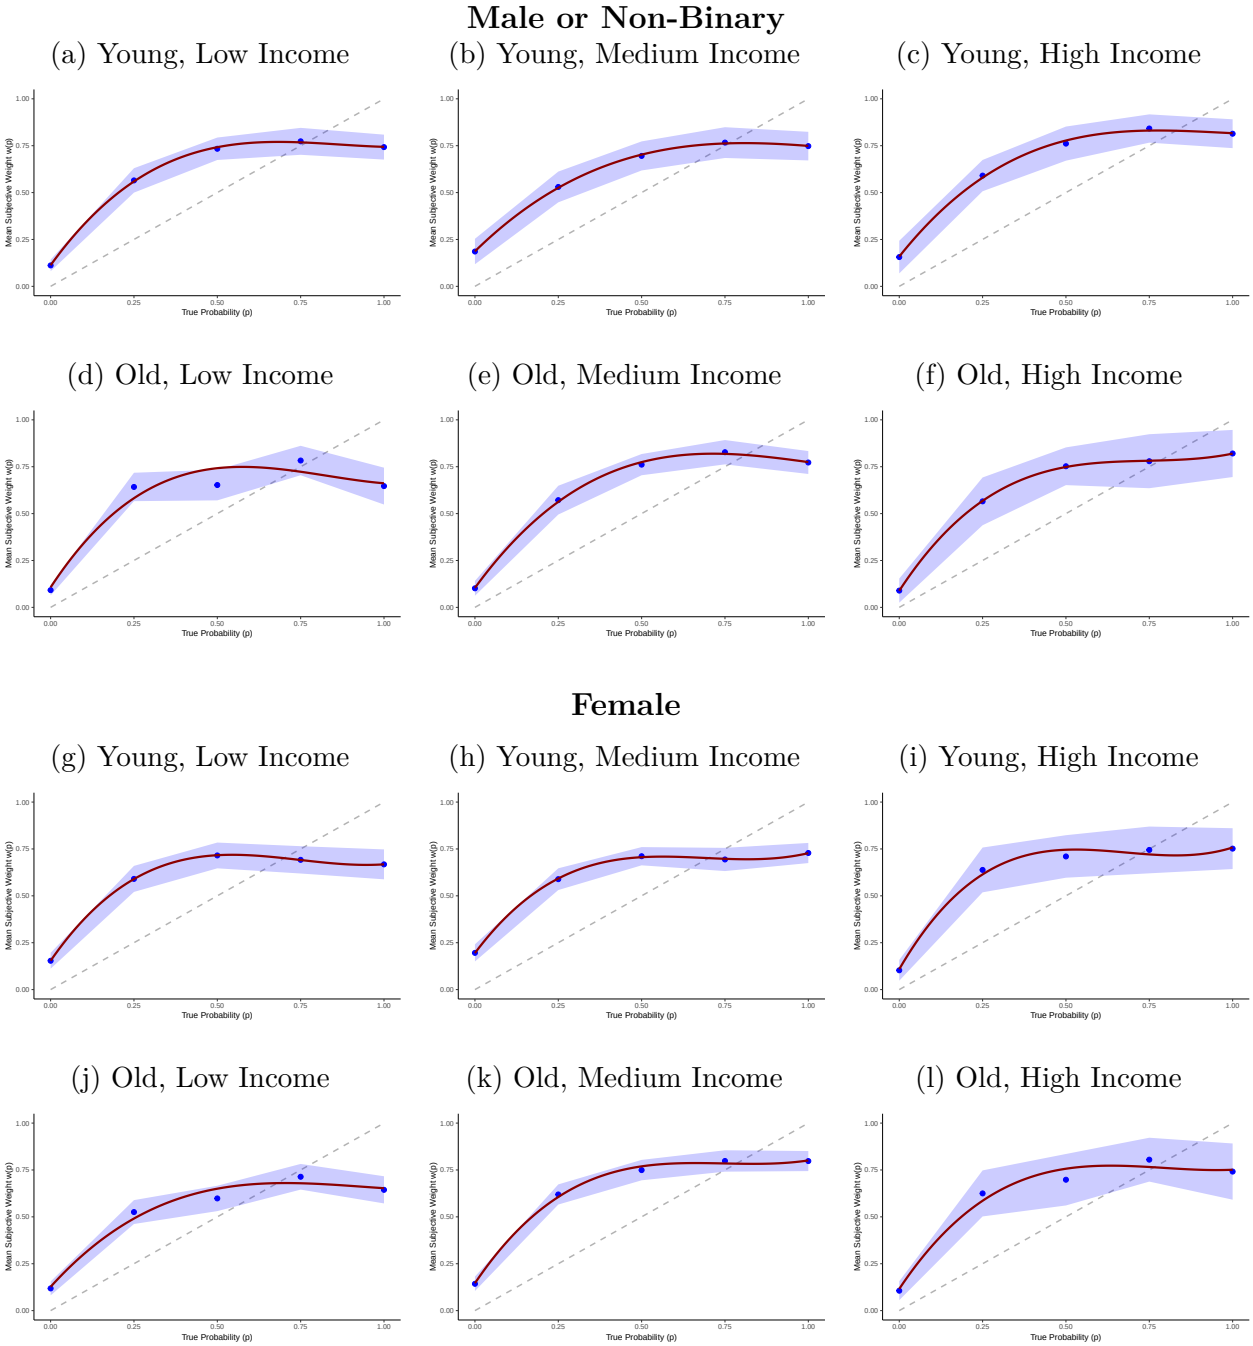

Note: Each panel shows the estimated probability weighting function $w(p) = b_0 + b_1 p + b_2 p^2 + b_3 p^3$ for the corresponding demographic cell. For full set of parameter estimates, see Table 5.

Figure 2: Camden Choice-Based Lettings advertisements

(a) Example of advertisement in Camden arrivals data.

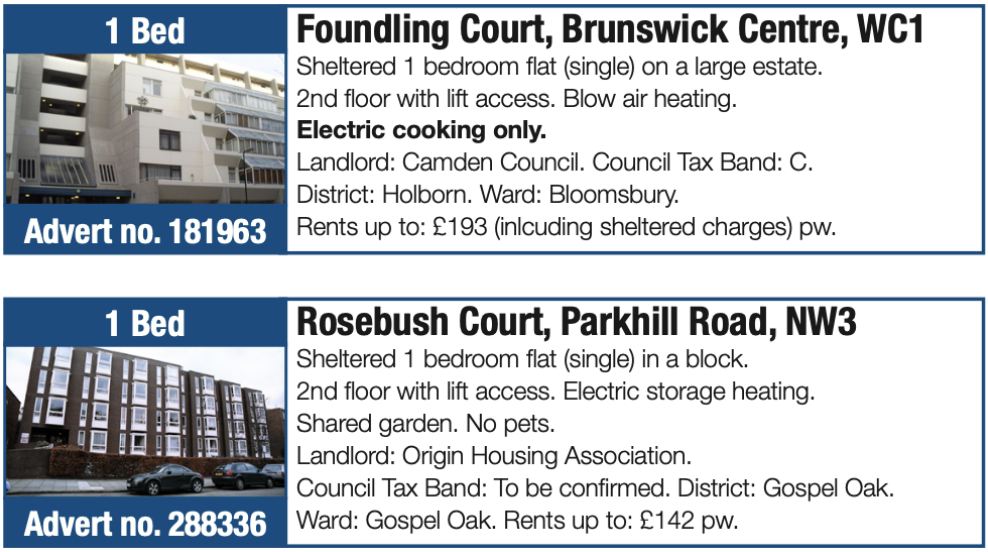

(b) Example of information on previous week's adverts

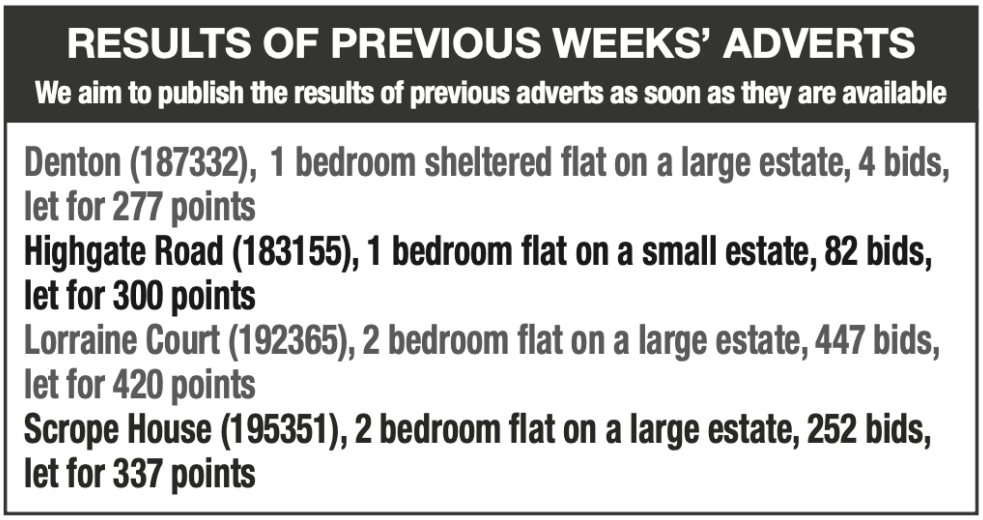

Note: Panel (a) shows a weekly advertisement listing available social housing units in Camden, including property characteristics such as rent, property type, location, and bedroom count. Panel (b) shows the results of the previous week's advertisements, including the priority points of the winning applicant and the number of sign-ups received. These data are publicly available for the period January 2016–February 2021.

Figure 3: Camden vs England on waitlist length and housing stock

(a) Waitlist Length

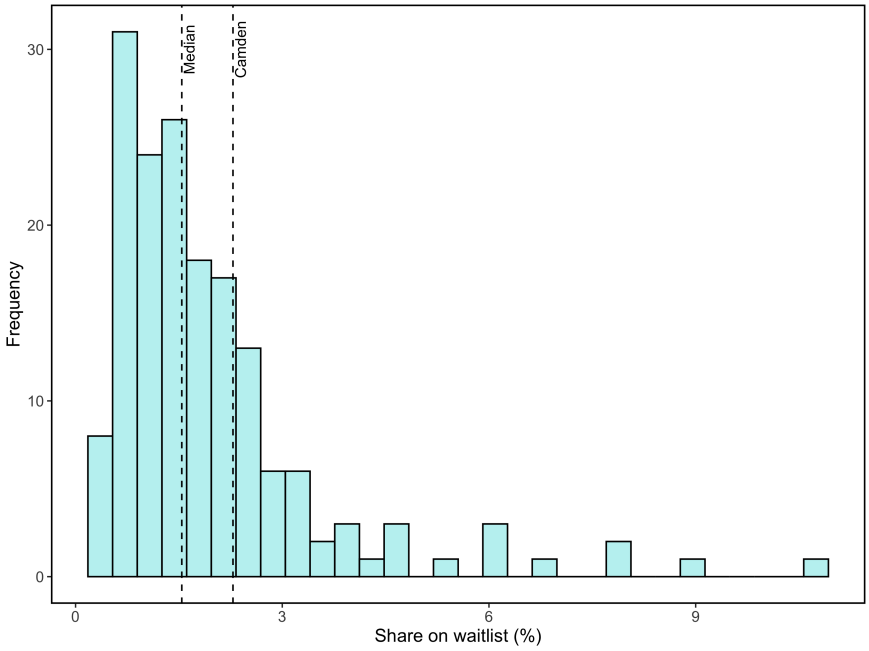

(b) Size of Housing Stock

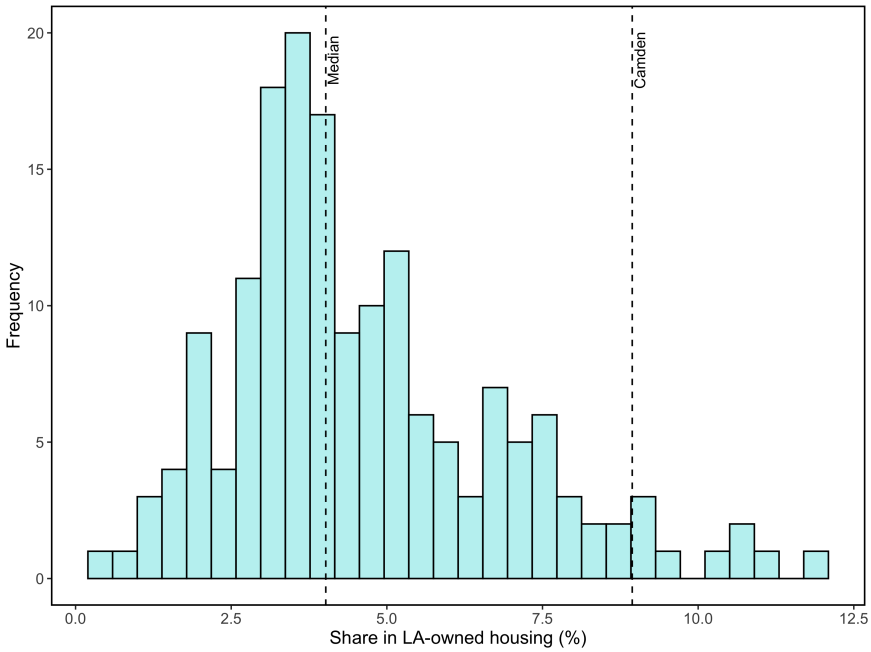

Note: Figures show histograms of (a) share of local authority population on waitlist and (b) share of local authority population living in local-authority owned housing (size of housing stock) across local authorities in England, with Camden denoted. Data come from local authority housing statistics and the 2019 UK Housing Survey Data.

Figure 4: Arrivals in Camden data

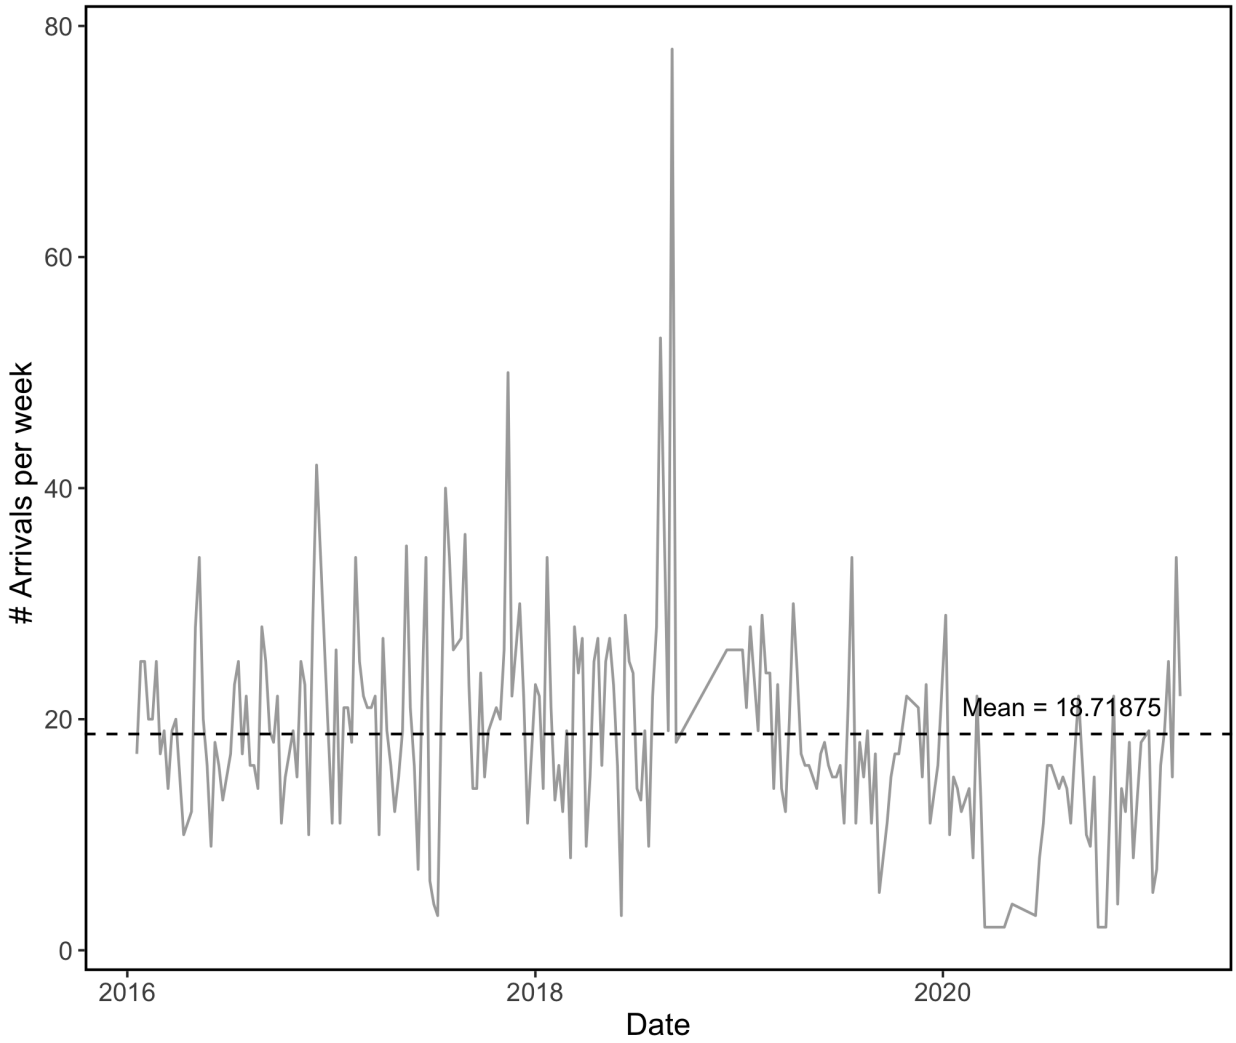

Note: Number of advertisements each week in Camden council, plotted over time where data is available (2016–2020). Mean over the period also shown.

Figure 5: Comparison of rent distribution

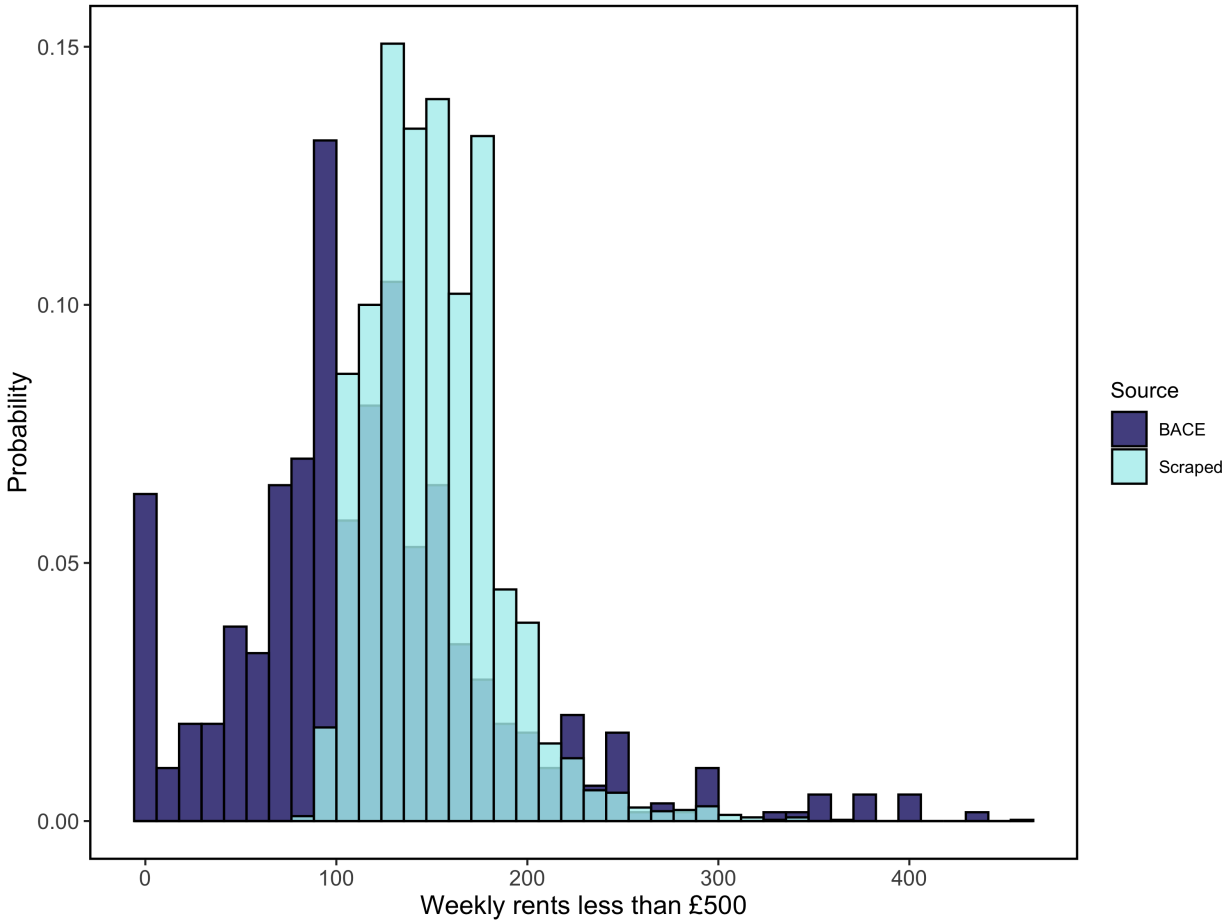

Note: Histogram of rent from Camden advertisements data, and from BACE data.

Figure 6: UK vs US: congestion and prioritization under rationality

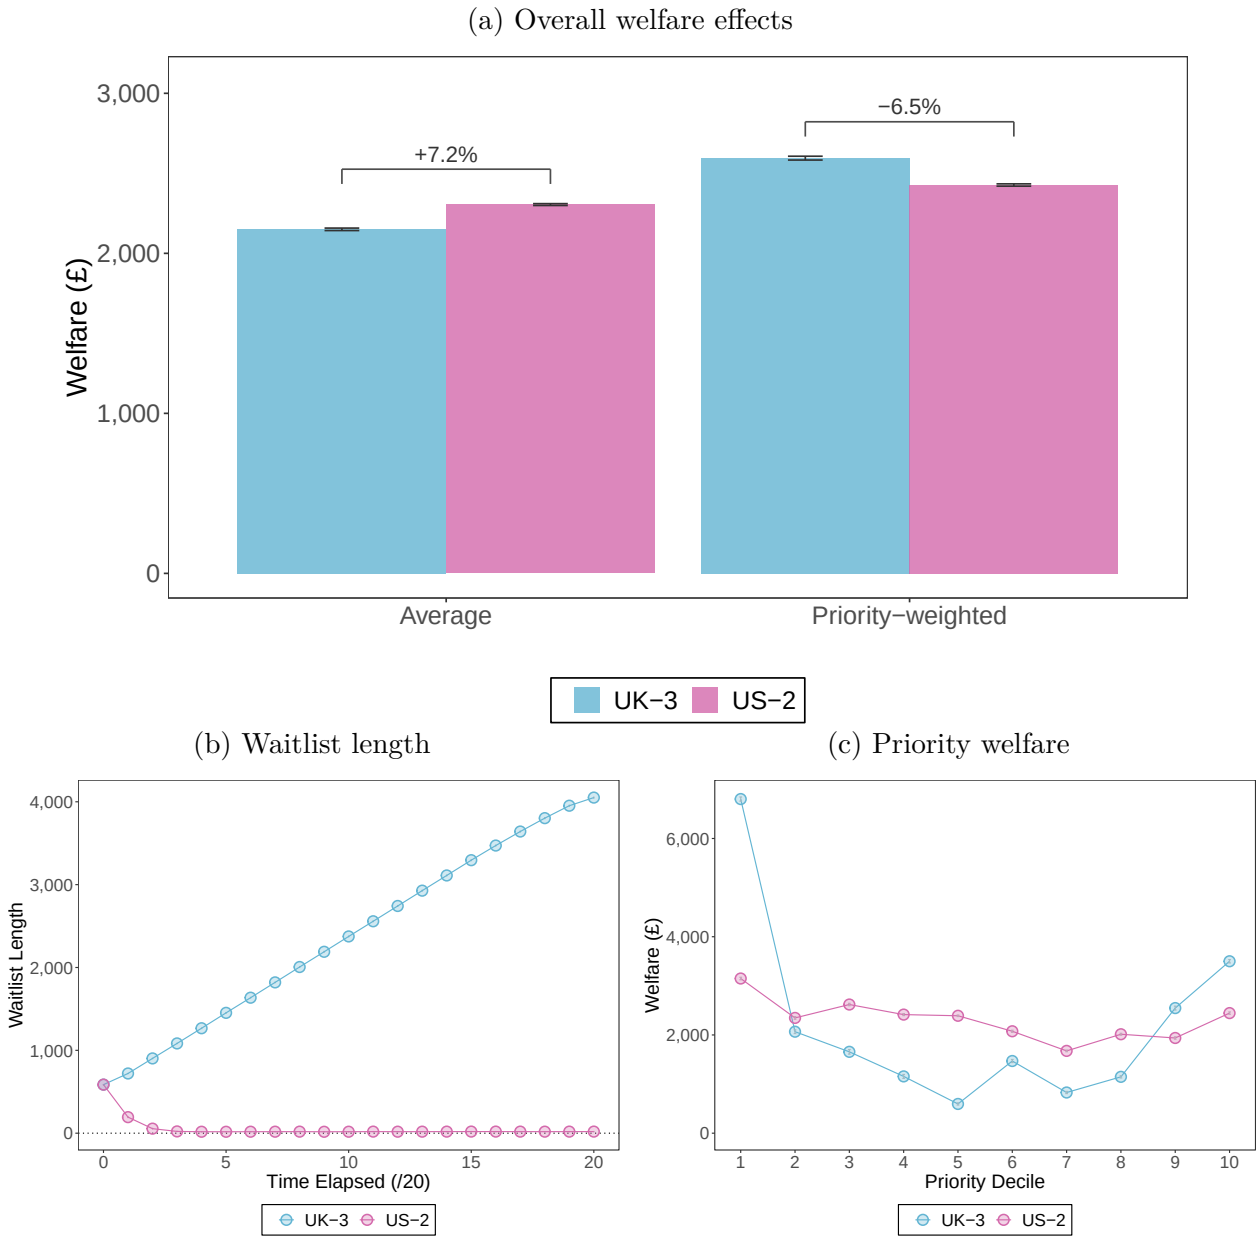

Note: Plot shows results of simulations of $UK-3$ and $US-2$ mechanisms. Figure 6a shows the overall welfare estimates from Table 7. In Figures 6b and 6c, each circle represents an across-simulation average and across-simulation Monte-Carlo standard errors are also shown. Figure 6b shows the length of the waitlist over time (split into 20 stages) and Figure 6c shows welfare split by deciles of the priority distribution.

Figure 7: UK vs US: overall effects of bias

(a) Average welfare

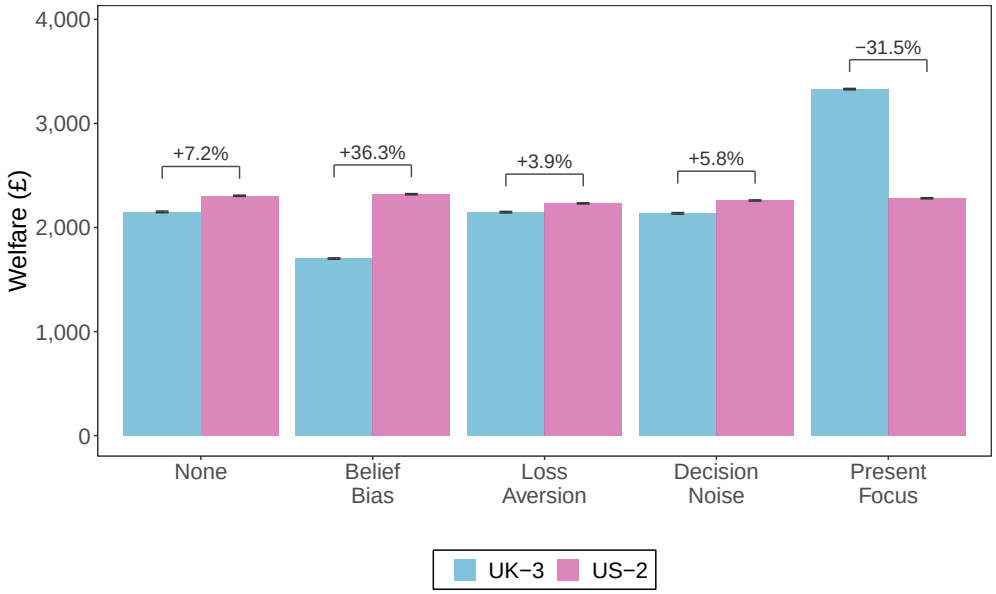

(b) Priority-weighted welfare

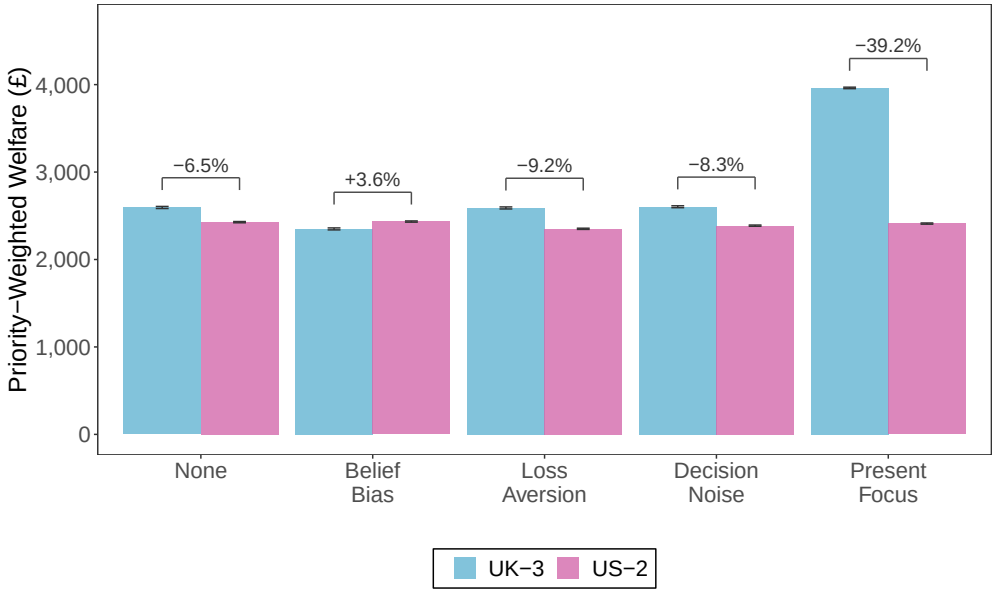

Note: Plot shows welfare effects under different behavioral biases for $UK-3$ and $US-2$ mechanisms. Each bar represents an across-simulation average and error bars show Monte-Carlo standard errors. Figure 7a shows average welfare where all applicants are weighted equally. Figure 7b shows priority-weighted welfare where weights are proportional to priority points.

Figure 8: UK-3 mechanism performance under different behavioral biases

(a) Waitlist length over time

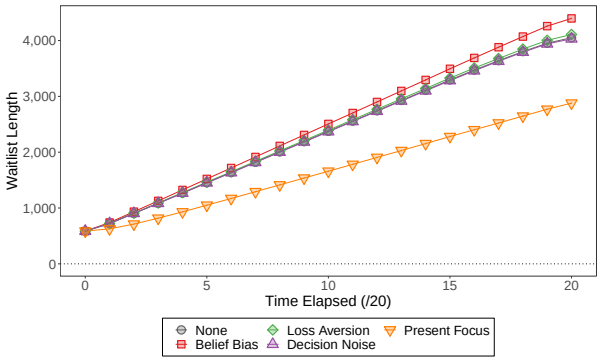

(b) Welfare by priority decile

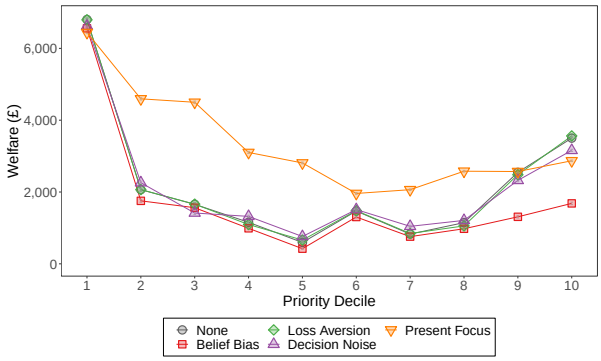

(c) Wasted arrivals

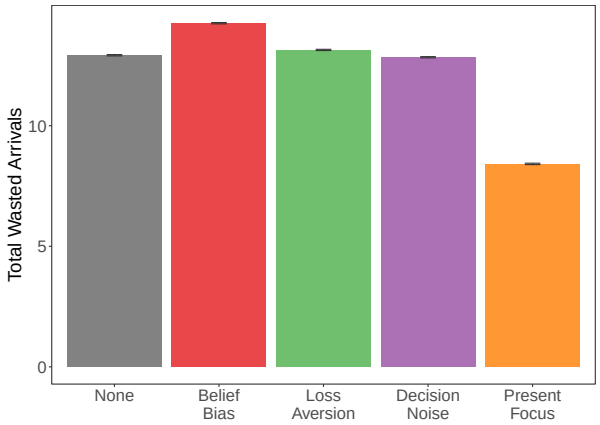

(d) Average number of bids

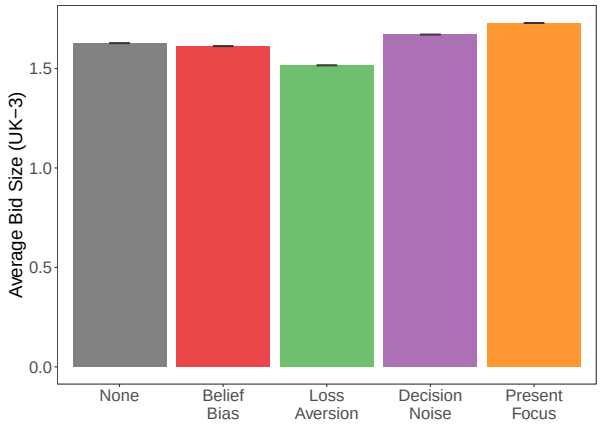

Note: Plot shows performance of the $UK-3$ mechanism under different behavioral biases. Figure 8a shows waitlist accumulation over time (split into 20 stages). Figure 8b shows welfare across priority deciles, where D1 is highest priority. Figure 8c shows the average number of arrivals per week which go to waste (not allocated to any applicant). Figure 8d shows the average number of bids placed per period under each bias type. Error bars and shaded regions represent Monte-Carlo standard errors across simulations.

Figure 9: Welfare loss from false acceptances and false rejections in US-2 mechanism

(a) Number of False Acceptances / Rejections (per person)

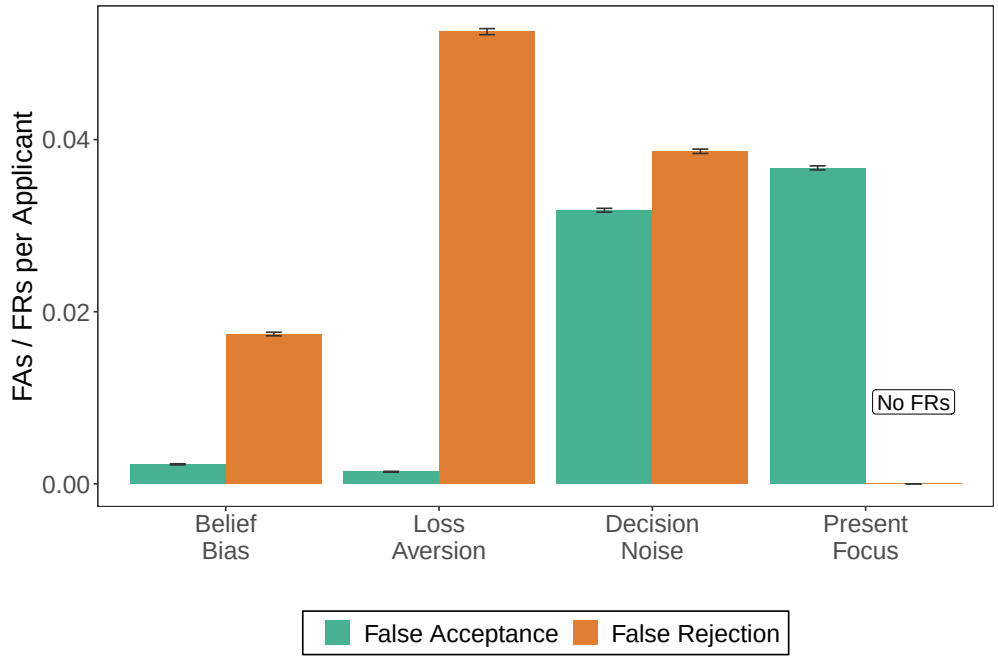

(b) Average welfare loss per decision

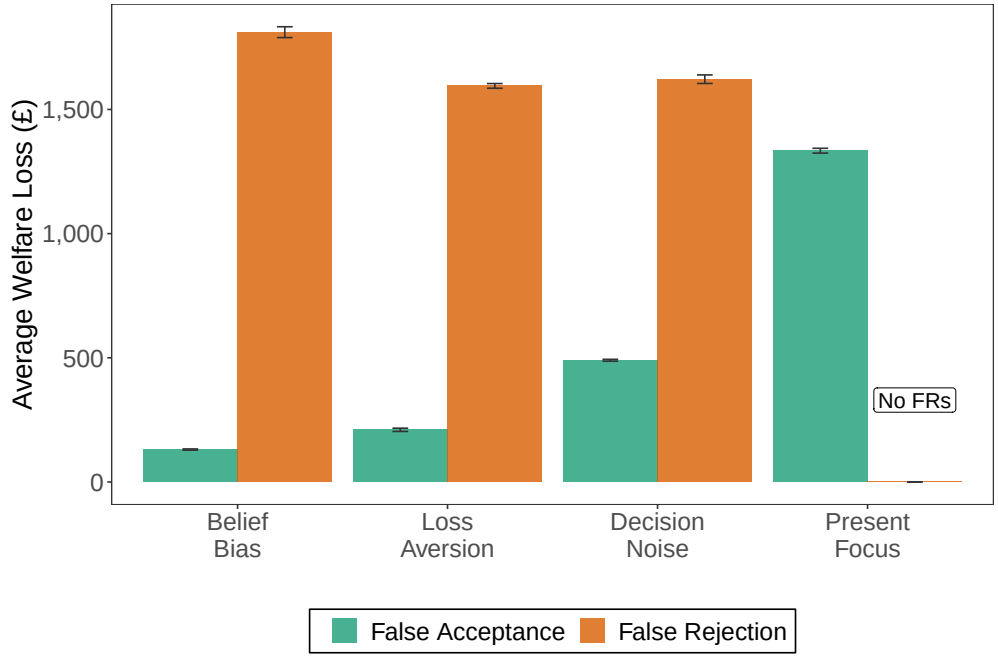

Note: Plot decomposes welfare losses in the $US-2$ mechanism under different behavioral biases. False acceptances occur when applicants accept offers they should optimally reject. False rejections occur when applicants reject offers they should optimally accept. Figure 9a shows the number of false acceptances / rejections per simulated applicant. Figure 9b shows average welfare loss per wrong decision. Error bars represent Monte-Carlo standard errors across simulations.

Table 1: Demographic summary statistics

Characteristic	Percentage
Gender	
Man	38.8
Woman	59.5
Non-Binary / Other	1.7
Age	
Mean (SD)	39.2 (13.3)
Range	18–82
Education	
High school or less	40.2
University degree	59.8
Marital status	
Single	57.9
Employment	
Employed	39.9
Monthly income	
Low (<£750)	40.5
Medium (£750–£2,500)	47.2
High (>£2,500)	12.3
Weekly rent	
Mean (SD)	151.7 (74.5)
Sample size	587

Note: This table presents demographic characteristics of the BACE respondents. Income categories are based on monthly household income. Education is classified as high school or less (including further education) versus university degree (bachelor's or master's degree). Employment refers to full-time employment status.

Table 2: Summary statistics of estimated preferences

	Mean	St. Dev.	P25	Median	P75
Panel A: Willingness to pay (units of weekly rent)					
Preferred location	37.47	37.27	12.86	43.06	67.15
Property type (vs. high-rise flat)					
House	48.61	37.72	28.63	60.33	77.78
Low-rise flat	12.53	40.71	-17.10	15.03	43.60
Property ownership (vs. private)					
Council	22.60	42.40	-7.86	26.74	57.74
Housing association	17.77	42.97	-15.47	21.11	53.42
Panel B: Other choice parameters					
Monthly discount factor δ_i	0.982	0.019	0.978	0.989	0.994
Choice consistency ν_i	4.86	1.25	4.02	4.92	5.69

Note: The top panel reports willingness-to-pay (WTP) implied by the random-utility model with rent coefficient normalized to -1. These values represent the weekly rent increase that leaves the respondent indifferent to gaining the stated attribute relative to the omitted category (non-preferred location, high-rise flat, private rental). The bottom panel reports the estimated monthly discount factor and the individual-level choice-consistency scale. St. Dev. denotes the standard deviation, P25 denotes the lower quartile, and P75 denotes the upper quartile.

Table 3: Prevalence of behavioral biases

Measure	UK Housing	Stango-Zinman (%)
Present bias		
Present bias (consumption)	18.2	15.1
Loss aversion		
Reject at least one lottery	81.6	63.5
Narrow bracketing		
Narrow bracketing (Task 1)	31.9	29.1
Non-belief in the law of large numbers		
$\in [0, 20)$	22.8	9.5
$\in [20, 40)$	43.6	41.0
$\in [40, 60)$	18.1	28.7
$\in [60, 80)$	6.0	7.8
$\in [80, 100)$	5.6	8.0
$= 100$	3.9	5.0
Gambler's and hot hand fallacies		
Gambler's fallacy	21.5	26.3
Hot-hand	17.6	13.6
Exponential growth bias (EGB)		
< 240	17.4	22.3
240	27.9	24.0
242	46.3	43.7
> 242	8.35	10.0
Limited attention		
Day-to-day	26.4	26.0
Medium-run	28.3	22.6
Long-run	38.7	34.9

Note: This table compares the prevalence of behavioral biases in our UK housing sample of 587 respondents with estimates from Stango and Zinman (2023) using the American Life Panel. Present bias is measured through intertemporal consumption choices. Loss aversion is identified by rejecting at least one of two small-stakes lotteries. Narrow bracketing is detected through dominance-violating paired choices. Non-belief in the law of large numbers (NBLLN) is measured based on beliefs about the percent chance that the number of heads out of 1000 fair coin tosses is a distance of (strictly) less than 20 away from 500. Gambler's fallacy and hot-hand beliefs are elicited through probability assessments after streaks. Exponential growth bias is measured through a compound interest calculation (correct answer: 242). Limited attention is self-reported across three time horizons of the response "Yes, and I/we often regret not paying greater attention."

Table 4: Relationship between housing preferences and behavioral biases

	Loss Av.		EGB		Hot Hand		Inatt.	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Location	3.1481	2.4283	-8.7824	-10.0938	-4.9537	-7.0662	1.9746	1.7317
	(4.2275)	(4.5978)	(4.2794)	(4.5725)	(4.1091)	(4.4938)	(5.2688)	(5.6091)
Disc. factor	-0.0013	-0.0010	-0.0009	0.0006	0.0044	0.0033	0.0025	0.0024
	(0.0020)	(0.0022)	(0.0020)	(0.0023)	(0.0018)	(0.0019)	(0.0024)	(0.0027)
Property type (relative to high rise)								
House	11.0875	10.7675	1.1162	-1.8987	-4.3485	-3.9579	1.7965	4.6626
	(4.7602)	(4.6814)	(4.3730)	(4.6298)	(4.5985)	(4.7710)	(5.2019)	(4.9077)
Low-rise	11.3972	9.8867	-7.2894	-9.8387	-3.2815	-1.9376	-7.9331	-7.7463
	(4.5812)	(4.9955)	(4.7024)	(5.0282)	(4.7221)	(5.1376)	(5.5549)	(5.8156)
Property ownership (relative to private)								
Assoc.	0.8845	2.1858	1.8879	0.6242	2.5544	1.7472	-19.8261	-15.8314
	(4.5218)	(4.5151)	(5.1291)	(5.3960)	(5.0395)	(5.1273)	(5.6163)	(6.1080)
Council	0.3564	0.5831	-3.2301	-4.4878	0.2493	0.1824	-13.1238	-10.2641
	(4.4941)	(5.0528)	(5.1447)	(5.6461)	(4.8773)	(5.3953)	(6.1196)	(6.4408)
Controls		X		X		X		X

Note: This table reports coefficients from regressions of housing preference parameters on behavioral bias measures, using the sample of 587 BACE respondents. The dependent variables are as follows: an indicator for residing in their preferred location, monthly discount factor, preference for house and low-rise flat (relative to high-rise flat), and preference for housing association and council ownership (relative to private ownership). Independent variables are indicators for loss aversion (rejecting at least one lottery), exponential growth bias (incorrect compound interest calculation), hot-hand fallacy (believing in streaks), and limited attention (regretting inattention to finances). Even-numbered columns include controls for age (quadratic), household size and composition, gender, marital status, ethnicity, nativity, monthly income, education, and employment status. Standard errors in parentheses.

Table 5: Behavioral parameter estimates by demographic cell

	Probability Weighting				Present Bias		Loss Aversion	
Cell	b_0	b_1	b_2	b_3	Share	β	Share	λ
Man/NB, Young, Low	0.113	2.454	-2.963	1.139	0.140	0.807	0.807	1.960
	(0.031)	(0.316)	(0.800)	(0.525)	(0.011)			
Man/NB, Young, Medium	0.187	1.737	-1.649	0.474	0.205	0.807	0.795	1.960
	(0.039)	(0.398)	(1.011)	(0.664)	(0.014)			
Man/NB, Young, High	0.159	2.232	-2.404	0.830	0.111	0.807	0.500	1.960
	(0.042)	(0.426)	(1.082)	(0.711)	(0.019)			
Man/NB, Old, Low	0.107	2.714	-3.615	1.455	0.167	0.807	0.833	1.960
	(0.039)	(0.402)	(1.021)	(0.671)	(0.013)			
Man/NB, Old, Medium	0.104	2.438	-2.622	0.854	0.213	0.807	0.770	1.960
	(0.031)	(0.315)	(0.802)	(0.528)	(0.012)			
Man/NB, Old, High	0.089	2.718	-3.605	1.619	0.125	0.807	0.688	1.960
	(0.058)	(0.600)	(1.525)	(1.001)	(0.020)			
Woman, Young, Low	0.154	2.565	-3.709	1.657	0.143	0.807	0.873	1.960
	(0.034)	(0.350)	(0.890)	(0.585)	(0.011)			
Woman, Young, Medium	0.195	2.372	-3.559	1.720	0.253	0.807	0.874	1.960
	(0.027)	(0.279)	(0.709)	(0.466)	(0.010)			
Woman, Young, High	0.109	3.051	-4.710	2.308	0.059	0.807	0.941	1.960
	(0.054)	(0.559)	(1.421)	(0.933)	(0.016)			
Woman, Old, Low	0.127	1.965	-2.240	0.801	0.197	0.807	0.776	1.960
	(0.032)	(0.330)	(0.839)	(0.551)	(0.010)			
Woman, Old, Medium	0.146	2.629	-3.557	1.583	0.143	0.807	0.896	1.960
	(0.026)	(0.268)	(0.681)	(0.448)	(0.010)			
Woman, Old, High	0.116	2.664	-3.496	1.468	0.286	0.807	0.667	1.960
	(0.061)	(0.632)	(1.612)	(1.060)	(0.021)			

Note: This table reports behavioral parameter estimates for each demographic cell. Probability weighting function takes the form $w(p) = b_0 + b_1 p + b_2 p^2 + b_3 p^3$, where p is the true probability of receiving an offer. Standard errors in parentheses for probability weighting parameters. Share (Present Bias): Proportion of cell exhibiting present bias based on consumption choices. β: Mean present focus parameter from quasi-hyperbolic discounting model, calibrated assuming $\beta \sim \mathcal{U}[0.7, 0.9]$ following Frederick et al. (2002) and Bradford et al. (2019). Share (Loss Aversion): Proportion of cell exhibiting loss aversion based on lottery rejection. λ: Loss aversion parameter, calibrated assuming $\lambda \sim \mathcal{U}[1.82, 2.10]$ following Brown et al. (2024). Cell format: Gender, Age Group (Young/Old split at median), Income Level (Low/Medium/High). Man/NB includes men and non-binary respondents. An F-test rejects homogeneity across cells for probability weighting ($F = 3.12$, $p < 0.001$) and loss aversion ($\chi^2 = 13.7$, $p = 0.01$), but not for present bias.

Table 6: Summary statistics: Housing arrivals

	Mean	Min	Max
Panel A: Arrival and bidding characteristics			
Arrivals per week	18.72	2	78
	(9.04)		
Number of bids per let property	189.81	1	810
	(158.37)		
Priority points (winning bidder)	443.29	0	2670
	(306.01)		
Avg. number of bids in year 2016	221.79	–	–
	(166.52)		
Avg. number of bids in year 2017	169.97	–	–
	(156.19)		
Avg. number of bids in year 2018	168.69	–	–
	(148.35)		
Avg. number of bids in year 2019	178.10	–	–
	(152.33)		
Avg. number of bids in year 2020	194.30	–	–
	(134.32)		
Panel B: Property characteristics			
Weekly rent (£)	152.05	82	1317
	(39.91)		
Number of bedrooms	1.55	1	6
	(0.86)		
Flats	99.0%		
Houses	1.0%		
Sheltered housing	24.5%		
Lift access	46.1%		
Council-owned	72.5%		
Panel C: Rent by number of bedrooms			
1 bedroom	145.85	82	631
	(36.33)		
2 bedrooms	157.24	106	1317
	(50.11)		
3 bedrooms	168.30	109	342
	(28.81)		
4 bedrooms	179.91	139	236
	(17.84)		
5 bedrooms	186.33	171	222
	(13.67)		
6 bedrooms	188.33	176	221
	(17.18)		

Note: This table presents summary statistics for housing unit arrivals from Camden Council administrative data, covering weekly advertisements from January 2016 to February 2021. Standard deviations in parentheses where applicable. Panel A shows arrival rates and competitive intensity. Panel B describes physical property characteristics. Panel C breaks down rents by property size. Priority points measure need-based ranking, with higher values indicating greater need. The UK-3 mechanism operates weekly, batching all arrivals within a week for simultaneous bidding.

Table 7: Welfare Estimates by Mechanism and Bias Type

	US-2		UK-3	
Bias	Average	Priority-wtd	Average	Priority-wtd
None	2304.8	2427.0	2150.1	2594.9
	(5.8)	(7.2)	(7.7)	(11.9)
Belief Bias	2319.9	2434.0	1701.6	2349.8
	(5.8)	(7.2)	(6.7)	(11.4)
Loss Aversion	2231.5	2351.0	2146.8	2589.5
	(5.7)	(7.0)	(7.6)	(11.7)
Decision Noise	2259.3	2387.1	2135.7	2603.4
	(5.8)	(7.0)	(7.7)	(10.9)
Present Focus	2280.7	2410.9	3329.5	3962.2
	(5.5)	(6.8)	(6.8)	(8.4)

Note: This table reports welfare estimates (in £) for the UK-3 and US-2 mechanisms under different behavioral biases, measured as the net present value of participating in the mechanism versus remaining in the outside option (private housing). "None" represents the rational baseline. "Belief Bias" implements probability weighting functions estimated from offer probability elicitation tasks. "Loss Aversion" models reference-dependent preferences over rents with parameters estimated from lottery choices. "Decision Noise" adds i.i.d. choice noise to flow utilities, calibrated from choice inconsistency in the BACE. "Present Focus" implements naive quasi-hyperbolic discounting with β parameters estimated from intertemporal consumption choices. Average" weights all applicants equally. "Priority-weighted" uses welfare weights proportional to priority points: $g(a) = p(a)/\sum_{a'} p(a')$. Monte-Carlo standard errors in parentheses, computed across 100 simulation runs. Each simulation covers a 5-year period with applicants drawn from the BACE sample ($N = 587$) and housing arrivals following the empirical distribution from Camden administrative data. The discount factor is calibrated from the BACE discrete choice experiment. Each bias is implemented separately (not jointly) to isolate its welfare effect.

Table 8: Probability weighting function parameters by demographic cell

Cell	b_0	b_1	b_2	b_3	R^2	N
Man or NB, Young, Low	0.113	2.454	-2.963	1.139	0.533	57
	(0.031)	(0.316)	(0.800)	(0.525)		
Man or NB, Young, Medium	0.187	1.737	-1.649	0.474	0.411	44
	(0.039)	(0.398)	(1.011)	(0.664)		
Man or NB, Young, High	0.159	2.232	-2.404	0.830	0.677	18
	(0.042)	(0.426)	(1.082)	(0.711)		
Man or NB, Old, Low	0.107	2.714	-3.615	1.455	0.458	42
	(0.039)	(0.402)	(1.021)	(0.671)		
Man or NB, Old, Medium	0.104	2.438	-2.622	0.854	0.551	61
	(0.031)	(0.315)	(0.802)	(0.528)		
Man or NB, Old, High	0.089	2.718	-3.605	1.619	0.587	16
	(0.058)	(0.600)	(1.525)	(1.001)		
Woman, Young, Low	0.154	2.565	-3.709	1.657	0.372	63
	(0.034)	(0.350)	(0.890)	(0.585)		
Woman, Young, Medium	0.195	2.372	-3.559	1.720	0.359	95
	(0.027)	(0.279)	(0.709)	(0.466)		
Woman, Young, High	0.109	3.051	-4.710	2.308	0.558	17
	(0.054)	(0.559)	(1.421)	(0.933)		
Woman, Old, Low	0.127	1.965	-2.240	0.801	0.352	76
	(0.032)	(0.330)	(0.839)	(0.551)		
Woman, Old, Medium	0.146	2.629	-3.557	1.583	0.535	77
	(0.026)	(0.268)	(0.681)	(0.448)		
Woman, Old, High	0.116	2.664	-3.496	1.468	0.452	21
	(0.061)	(0.632)	(1.612)	(1.060)		

Note: This table reports estimated parameters of the probability weighting function $w(p) = b_0 + b_1 p + b_2 p^2 + b_3 p^3$ for each demographic cell, where p is the true probability of receiving a housing offer. Parameters are estimated by eliciting beliefs about offer probabilities for hypothetical properties and comparing them to objective probabilities calculated from Camden administrative data on past allocations. Respondents were shown historical allocation data (priority points of winning bidders) and asked to assess their likelihood of receiving offers for properties of varying desirability. Standard errors in parentheses. R^2 measures within-cell fit of the probability weighting function. Total $N = 587$.