The Public Housing Allocation Problem

Zeky Murra-Anton
ISO New England

Neil Thakral
Brown University

May 2024*

Abstract

This paper introduces a dynamic model of public housing allocation in which units that arrive stochastically must be matched upon arrival with applicants on a waiting list. We establish a lower bound on the number of possible violations of ex-post Pareto efficiency or envy-freeness that grows linearly with the applicant pool. We construct a novel strategy-proof direct mechanism that yields interim Pareto efficient and envy-free outcomes by optimizing applicants' preferences over buildings and expected waiting times through building-specific queues. Using data on preferences for public housing, we estimate that adopting the proposed mechanism would improve welfare by $2,300–$4,900 per applicant.

*Murra-Anton: Market Development, ISO New England, 1 Sullivan Road, Holyoke, MA 01040 (email: zmurraanton@iso-ne.com). Thakral: Department of Economics, Brown University, Box B, Providence, RI 02912 (email: neil_thakral@brown.edu). We thank Bernadette Hicks and Clive Johnston for excellent research assistance. We are grateful to Kfir Eliaz, Scott Duke Kominers, Maciej Kotowski, Canishk Naik, Giacomo Rubbini, Ran Shorrer, Linh Tô, and Daniel Waldinger for helpful comments and discussions. This paper supersedes and replaces earlier versions that circulated under the same title by Thakral (2016). This paper has not been funded or endorsed by ISO New England. This material is based upon work supported by the National Science Foundation under Grant No. 2149371.

1 Introduction

Dynamic stochastic allocation problems arise in many settings, such as disaster-affected communities waiting for emergency supplies, prospective adoptive parents seeking children, parents waiting for openings in childcare facilities, patients awaiting organ transplants, patients waiting for appointments with medical specialists, freelance professionals seeking job opportunities, ride-share drivers seeking passengers, and refugee families awaiting humanitarian sponsors (Farajzadeh et al., 2023). To illustrate and analyze the general features of such problems, we consider the allocation of public housing. On one side of the market, apartment units in different buildings, which hold priorities over applicants, arrive stochastically over time. On the opposite side, applicants express preferences for apartment buildings and wait for units to become available.

Public housing authorities (PHAs) aim to satisfy several practical objectives when allocating units to applicants. First, due to extreme levels of excess demand for public housing units, policymakers and the general public alike tend to view systems that allow vacancies as "unacceptable."[1] Second, PHAs strive to ensure housing allocations align closely with applicants' needs.[2] Third, PHAs aim to provide affordable housing to those with the greatest need according to a system of priorities. A meaningful priority system ensures that higher-priority applicants do not desire lower-priority applicants' assignments, and it safeguards against the possibility of lower-priority applicants manipulating the system to their advantage. Balancing these objectives in practice presents a formidable challenge for policymakers.[3]

[1] The quote is from the Massachusetts Secretary of Housing and Livable Communities, who further goes on to say, "we need to do everything we can to make sure that every single one of our precious public housing units is filled and the amount of time between tenants is as short as is humanly possible" (Wallack and Willmsen, 2023). Housing authorities also incur a pure financial loss due to vacancies; the District of Columbia Housing Authority calculates that the "loss per month for a vacant unit is $337 in rent and $674 in federal subsidy" (Thompson, 2022).

[2] A report by the US Department of Housing and Urban Development notes "the serious transportation and employment barriers that would be created by forcing people to move from one end of the county to the other. For this reason, applicants can indicate the part of the county that they prefer. Even PHAs with restrictive policies make exceptions for applicants who would suffer a significant hardship if they were unable to choose particular locations" (Devine, Rubin and Gray, 1999). Similarly, a policy report from the UK states, "the overall aim for choice-based lettings schemes is to ensure that: the most effective and efficient use is made of accessible housing stock; and disabled people are allocated accommodation which meets their needs, while giving them the widest possible choice and greater say over where they live" (Department for Communities and Local Government, 2011).

[3] The executive director of the Chelmsford Housing Authority denounced the system for selecting

We construct a dynamic matching model, capturing practical features of the public housing allocation problem, and use it to study the design of fair and efficient *non-wasteful allocation mechanisms*—that is, mechanisms that do not hold units vacant when a suitable match exists. From an ex-post perspective, we first establish that every non-wasteful mechanism creates the potential for widespread efficiency or fairness violations. To overcome this limitation while preserving the essence of the underlying institutional objectives, we develop *interim* notions of efficiency and fairness, which take the perspective of the realized assignment timing as it unfolds within the mechanism. Then we propose a novel strategy-proof mechanism, which accounts for applicants' priorities by building on the idea of the "you request my house—I get your turn" algorithm of Abdulkadiroğlu and Sönmez (1999), and show that it satisfies the desired axioms from our interim perspective. Finally, leveraging data on preferences for public housing, we conduct simulations to empirically assess the welfare consequences of adopting the proposed mechanism.

Uncertainty about future arrivals generates tension between a mechanism's efficiency or fairness on the one hand and non-wastefulness on the other. A non-wasteful mechanism must assign a unit at or before its arrival time without perfect foresight, creating the potential for a Pareto-improving reallocation or justified envy involving present and future assignees. This tension leads to our first result, which establishes the impossibility of designing a non-wasteful allocation mechanism that is ex-post Pareto efficient or envy-free (Theorem 1). While this parallels classic impossibility results in the market design literature (e.g., Roth, 1982; Balinski and Sönmez, 1999), our result offers a unique perspective that emerges from the dynamic stochastic environment.

Previous work emphasizes how imposing restrictions on the model's primitives restores possibility.[4] In contrast, our ex-post impossibility result arises fundamentally due to uncertainty. Even under extreme restrictions on preferences and priorities, we show that fairness issues persist. Our result also complements the recent dynamic

potential tenants in Massachusetts as "the most horrible, horrible, inefficient program" (Wallack and Willmsen, 2023). The Productivity Commission of the Australian Government states that "Australia's social housing system is broken. Eligible tenants have little choice over the home they live in and can face a lengthy wait to access housing. . . " (Productivity Commission, 2017).

[4]For example, the top dominance condition on preferences (Alcalde and Barberà, 1994) overcomes (a stronger version of) the impossibility result in Roth (1982), and the acyclicity condition on priorities (Ergin, 2002) overcomes the impossibility result in Balinski and Sönmez (1999). The mechanism design literature offers many such examples: for instance, the single-peakedness condition on preferences (Moulin, 1980) overcomes the Gibbard-Satterthwaite theorem.

two-sided matching literature that explores the trade-off between withholding matches to thicken the market and increasing waiting times (e.g., Baccara, Lee and Yariv, 2020; Ashlagi, Nikzad and Strack, 2023). In our setting, ex-post efficiency necessitates generating waste by discarding—but not delaying—units for which a suitable match exists.

The inability to achieve efficiency and fairness from an ex-post perspective naturally raises the question of whether a mechanism can confine violations to only a small share of applicants. Once again, an impossibility arises: No mechanism can guarantee limiting violations to fewer than (the integer part of) half of the number of applicants (Theorem 2).

Formulating suitable axioms that account for fairness and efficiency in a dynamic stochastic context poses significant challenges. The axiomatic paradigm that exemplifies canonical market design mechanisms (Sönmez, 2023) has thus found relatively limited traction in settings with uncertainty about future arrivals. To make progress on this problem, we investigate whether alternative notions of efficiency and fairness suffer from the same uncertainty effects that generate the impossibility result. Defining these properties requires specifying a reference time for comparing a given applicant's assignment with alternatives, such as another applicant's or another mechanism's assignment. Without fixing the reference time in these comparisons, issues similar to those uncovered in our impossibility result arise because a non-wasteful mechanism cannot fully account for the randomness in the arrival process.

We introduce interim notions of efficiency and fairness to overcome the ex-post impossibility result in our setting. These criteria use the applicant's assignment time as the reference time for comparisons. Formally, our interim notions compare whether an applicant prefers, from the time the mechanism selects them, waiting for their assignment to waiting for a different one. Our main result characterizes an efficient and fair non-wasteful mechanism from the described interim perspective (Theorem 3 and Theorem 4). Moreover, we show that alternative perspectives, such as ex-ante and allowing applicants to consider whether they prefer their allocations over swapping places with other applicants, allow the non-wastefulness of a mechanism to interact with the uncertainty about future arrivals in a way that leads to impossibility (Theorem 5 and Theorem 6).

We propose a non-wasteful strategy-proof mechanism, the choice-based waiting list (CBWL) mechanism, which satisfies our proposed interim fairness and efficiency properties, as long as the set of building priorities are acyclic (Theorem 3). Furthermore, we establish that acyclicity of the priority structure also constitutes a necessary condition and thus fully characterizes the existence of a mechanism satisfying these properties (Theorem 4). An easy-to-implement algorithm achieves the CBWL mechanism matching outcome. First, applicants report their preferences and join an unordered general waiting list. As units arrive, the mechanism populates building-specific queues by optimizing applicants' preferences—accounting for their expected waiting time for units in each building—while respecting applicants' priorities at different buildings. In particular, when an unmatched unit becomes available, the algorithm aims to assign its highest-priority applicant to their preferred queue while ensuring that any applicant with a higher priority for that queue takes precedence.

Uncertainty is equally crucial for facilitating interim possibility as it is for generating ex-post impossibility. When the space of distributions over arrivals contains only deterministic arrivals, all applicants' interim comparisons are reduced to ex-post comparisons, a situation known in the literature to create impossibility (Roth, 1982). This highlights a dimension along which having a richer state space can enable desirable properties, contrasting with the conventional insight that restricting a model's primitives enhances implementability.

The central importance of uncertainty also distinguishes our work from existing papers that study dynamic matching markets from an axiomatic perspective. Most of these papers focus on dynamic aspects of stability (Damiano and Lam, 2005; Pereyra, 2013; Du and Livne, 2014; Kadam and Kotowski, 2018*a*,*b*; Kotowski, 2020; Bando and Kawasaki, 2021; Doval, 2022; Liu, 2023), but those that analyze efficiency also do not attribute any significant role to uncertainty (Kennes, Monte and Tumennasan, 2014; Kurino, 2014, 2020). Our approach also stands out from the broader dynamic allocation literature, which predominantly considers utilitarian objectives.[5]

Notwithstanding our emphasis on interim axioms, the tractable nature of our model allows for a detailed empirical analysis of ex-post welfare in the context of public housing allocation. Public housing authorities in the US tend to use waiting

[5]See, for instance, Derman, Lieberman and Ross (1972); Albright (1974); Su and Zenios (2004, 2005, 2006); Gershkov and Moldovanu (2009); M. Utku Ünver (2010); Bloch and Houy (2012); Gershkov, Moldovanu and Strack (2015); Bloch and Cantala (2017); Akbarpour, Li and Gharan (2020); Baccara, Lee and Yariv (2020); Schummer (2021); Leshno (2022); Loertscher, Muir and Taylor (2022); Che and Tercieux (2023); Nikzad and Strack (2023); Ashlagi, Monachou and Nikzad (forthcoming).

list policies that impose significant penalties on applicants who refuse offers, leaving households with little choice over their place of residence (Lui and Suen, 2011; Thakral, 2017) and creating suboptimal allocations (Theorem 7). To demonstrate the scale of the welfare gains achievable through changing the allocation system, we simulate arrival processes and matchings under counterfactual mechanisms. We do so by combining information on public housing developments operated by some of the largest public housing authorities in the United States with stated-preference data from Naik and Thakral (2023) on valuations of housing characteristics. Computing the monetary transfers that households would require under the existing mechanism to achieve the utility level of their allocation under CBWL, we find that changing the allocation mechanism to CBWL improves welfare by an average of \$2,300 to \$4,900 per household.

The ability to perform this empirical analysis not only complements our theoretical framework but also constitutes a distinct contribution in and of itself. Few empirical papers focus on design-related aspects of US housing assistance programs. This includes work that simulates the effects of counterfactual housing-voucher policies (Galiani, Murphy and Pantano, 2015) and public housing choice or priority systems (Waldinger, 2021; Naik and Thakral, 2023). The previous work on public housing explores alternative policy designs that exist in other contexts, with Naik and Thakral (2023) comparing welfare under offer-based and bidding-based allocation mechanisms, and Waldinger (2021) finding that housing authorities can improve the targeting of public housing benefits by appropriately exploiting priorities rather than by restricting choice. Our paper expands this research area by analyzing a novel policy design developed from theoretical principles.

More broadly, our paper contributes to a large body of research that characterizes the properties of and improves the design of matching mechanisms in a variety of applied domains such as labor markets (Roth and Peranson, 1999), school choice (Balinski and Sönmez, 1999; Abdulkadiroğlu and Sönmez, 2003), organ exchange (Roth, Sönmez and Ünver, 2004; Ergin, Sönmez and Ünver, 2020), course allocation (Sönmez and Ünver, 2010*a*; Budish and Cantillon, 2012), airport landing slots (Schummer and Vohra, 2013), vehicle licenses (Li, 2018), ridesharing (Liu, Wan and Yang, 2019), hunting licenses (Verdier and Reeling, 2022), teacher assignment (Combe, Tercieux and Terrier, 2022; Combe et al., 2022), and refugee resettlement (Delacrétaz, Kominers and Teytelboym, 2023), among others. Our work enriches this literature by introducing a theory for dynamic stochastic allocation problems and highlighting its application in the market for public housing.

The paper is organized as follows. Section 2 describes the dynamic allocation problem and introduces various properties. Section 3 establishes and assesses the impossibility of ex-post efficiency and fairness properties of strategy-proof direct mechanisms. Section 4 determines the necessary conditions for the existence of efficient and fair ex-ante strategy-proof direct mechanisms, introduces the CBWL algorithm, and shows that it has the desired properties. Section 5 applies the framework to public housing allocation and uses preference estimates from a discrete choice experiment to evaluate welfare. Section 6 concludes.

2 A Model of Public Housing Allocation

Time is continuous, denoted by $t \in \mathcal{T} := [0, \infty]$. There is a countably infinite **set of potential applicants** $\mathcal{I}$, with typical element $i \in \mathcal{I}$. There is also a countably infinite **set of potential buildings** $\mathcal{B}$, with typical element $\beta \in \mathcal{B}$. We denote by $I \subset \mathcal{I}$ and $B \subset \mathcal{B}$ arbitrary finite subsets of applicants and buildings. Finally, $\mathbb{N}$ is the set of natural numbers, with $\mathbb{N}_0 = \mathbb{N} \cup \{0\}$.

2.1 Buildings' Priorities

Each building has preferences or **priorities** over applicants, denoted by $\triangleright_\beta$, where $i \triangleright_\beta i'$ means that *building* β *prioritizes applicant* i *over applicant* i'. Applicant i is **acceptable** to β if $i \triangleright_\beta \beta$. The space of complete and transitive strict building priority profiles $(\triangleright_\beta)_{\beta \in \mathcal{B}}$ is denoted by $\triangle$. We note that priorities do not necessarily determine the order in which applicants receive their units; the allocation mechanism's rules do.[6] Although initially we do not restrict $\triangle$, we introduce the following notion of acyclicity based on Ergin (2002):

Definition 1 (Acyclicicity)**.** A priority profile $\triangleright$ is **acyclic** if there are no applicants $i_1, i_2, i_3 \in \mathcal{I}$ and buildings $\beta_1, \beta_2 \in \mathcal{B}$ such that $i_1 \triangleright_{\beta_1} i_2 \triangleright_{\beta_1} i_3 \triangleright_{\beta_1} \beta_1$ and $i_3 \triangleright_{\beta_2} i_1 \triangleright_{\beta_2} \beta_2$. We say that priority space $\triangle$ is **acyclic** when every $\triangleright \in \triangle$ is acyclic.

[6] Even if a building prioritizes low-income applicants, a high-income applicant may be assigned before in a first-come, first-serve mechanism, depending on their sign-up dates.

2.2 Unit Arrivals

Each building has multiple apartments or *units*, which are identical. Units become available or *arrive* sequentially over time. For simplicity, we do not restrict the number of units within each building.[7] We index units within a building according to their arrival order by $r \in \mathbb{N}$. We write $\langle \beta, r \rangle$ to denote unit r in building β.

Fixing a subset of buildings B, $a_t \in B \cup \{\emptyset\}$ is time t's **arrival**, representing the building the unit belongs to; $a_t = \emptyset$ represents the case where no unit arrived at t.[8] The set of all **stochastic arrival processes** over units in buildings B featuring at most one arrival per period is denoted by $\Pi(B)$, with typical element $\pi \in \Pi(B)$. The probability distribution of arrival a_t is denoted by $\pi(a_t)$. Notably, we allow non-Markovian arrival processes. Finally, $A(\pi)$ is the set of all **arrival sequences** $(a_t)_{t \in \mathcal{T}}$ such that $\pi(a_t) > 0$ for all $t \in \mathcal{T}$.[9]

The following arrival process is useful to prove some of our results:

Definition 2 (MAP)**.** A stochastic arrival process $\pi \in \Pi(B)$ is a **multinomial arrival process (MAP)** when the following conditions hold:

i. $\pi(a_t \in B) = 1$ when $t \in \mathbb{N}_0$ and $\pi(a_t = \emptyset) = 1$ otherwise.
ii. Whenever $t, t' \in \mathbb{N}_0$, a_t and $a_{t'}$ are independent and identically distributed.
iii. For each $\beta \in \mathcal{B}$, there exists $0 < p_\beta \leq 1$ such that $\pi(a_t = \beta) = p_\beta$ for all $t \in \mathbb{N}_0$.

2.3 Applicants' Preferences

To address the trade-off between the value applicants place on being matched with specific buildings and their associated waiting times, we posit that each applicant has **preferences** $\succ_i$ over building-delivery time pairs, where $(\beta, t) \succ_i (\beta', t')$ means that *applicant i strictly prefers receiving a unit in building β at time t over receiving one in building β' at time t'*. Building β is **acceptable** to i when $(\beta, t) \succ_i i$ for all $t \in \mathcal{T}$. We assume that the building's acceptability to applicants is atemporal in the sense that $(\beta, t) \succ_i i$ for some $t \in \mathcal{T}$ if and only if $(\beta, t) \succ_i i$ for all $t \in \mathcal{T}$. Then, we write $\beta \succ_i i$ without ambiguity. In some instances, for clarity of exposition, we compare

[7] We can interpret the problem as all units being occupied initially and becoming available over time, per the arrival process.

[8] For simplicity, a_t omits the unit number, but the history tracks this information.

[9] Our setup allows for a finite number of arrivals through the distribution π. For instance, the case where an arrival only happens at $t = 0$ is accounted for with $\pi(a_0 = \emptyset) = 0$ and $\pi(a_t = \emptyset) = 1$ for all $t > 0$.

units, in which case we write $(\langle\beta, r\rangle, t) \succ_i (\langle\beta', r'\rangle, t')$. Given that units are identical within each building, $(\langle\beta, r\rangle, t) \succ_i (\langle\beta', r'\rangle, t')$ is equivalent to $(\beta, t) \succ_i (\beta', t')$. We impose the following stationarity (dynamic consistency) and future discounting (costly waiting) conditions on preferences:

Definition 3 (Dynamic Consistency)**.** The preference $\succ_i$ satisfies **dynamic consistency** when $(\beta, t) \succ_i (\beta', t')$ if and only if $(\beta, t+k) \succ_i (\beta', t'+k)$ for all $k > 0$, $\beta \in \mathcal{B}$, and $t, t' \in \mathcal{T}$.

Definition 4 (Costly Waiting)**.** The preference $\succ_i$ satisfies **costly waiting** when $(\beta, t) \succ_i (\beta, t')$ for any $\beta \in \mathcal{B}$ and any $t, t' \in \mathcal{T}$ such that $t' > t$.

Finally, Θ is the space of complete, transitive, and risk neutral strict preference profiles $(\succ_i)_{i \in \mathcal{I}}$ satisfying dynamic consistency and costly waiting.[10]

2.4 The Public Housing Allocation Problem and Allocation Mechanisms

Definition 5 (PHAP)**.** A **public housing allocation problem (PHAP)** is a list $(I, B, \succ, \triangleright, \pi, a)$ such that $I \subset \mathcal{I}$, $B \subset \mathcal{B}$, $\succ \in \Theta$, $\triangleright \in \triangle$, $\pi \in \Pi(B)$, and $a \in A(\pi)$. We refer to the space of all such lists as the **PHAP space**, denoted $\mathcal{P}$.

A **matching** in problem $P = (I, B, \succ, \triangleright, \pi, a)$ is an injective function $\varphi_P \colon I \to (B \times \mathbb{N} \times \mathcal{T}) \cup I$ such that $\varphi_P(i) \in (B \times \mathbb{N} \times \mathcal{T}) \cup \{i\}$ for all $i \in I$. We refer to $\varphi(P) = (\varphi_P(i))_{i \in I}$ as a **matching outcome** and to $\varphi_P(i)$ as **applicant *i*'s matching outcome**.

When $\varphi_P(i) = i$, the applicant is unmatched. Otherwise, applicant i's matching outcome $\varphi_P(i) = (\mu_{\varphi(P)}(i), t_{\varphi(P)}(i))$ consists of an **assignment** $\mu_{\varphi(P)}(i) \in \mathcal{B} \times \mathbb{N}$ and a **delivery time** $t_{\varphi(P)}(i)$ for the assigned unit. In addition, we write $\mu_{\varphi(P)}(i) = \langle\beta_{\varphi(P)}(i), r_{\varphi(P)}(i)\rangle$ to denote that applicant i is assigned to unit $r_{\varphi(P)}(i)$ in building $\beta_{\varphi(P)}(i)$. Our matching definition pairs applicants with units, not arrivals, to allow for rules that assign future units in the present. Then, $T_{\varphi(P)}(i)$ denotes applicant i's **assignment time** and $\tau_{\varphi(P)}(\langle\beta, r\rangle, t')$ denotes the **expected waiting time** from

[10] Applicants' preferences can be represented by utility functions of the form $u(\beta, t) = u_\beta - ct$ for $u_\beta \in \mathbb{R} \cup \{-\infty\}$ and $c > 0$, with $u_\beta = -\infty$ representing the case where the applicant prefers to remain unmatched instead of a unit in building β. A tiebreaker is needed to keep the preferences strict over time-building pairs; our results, however, do not hinge on a specific tiebreaker.

t' until the delivery of unit $\langle\beta, r\rangle$.[11] Finally, when $a_t \neq \emptyset$, we denote the applicant receiving arrival a_t by $i_{\varphi(P)}(a_t) \in I \cup \{a_t\}$, with $i_{\varphi(P)}(a_t) = a_t$ representing the case where the arrival goes unassigned.

Definition 6 (Allocation Mechanism)**.** A **public housing allocation mechanism** φ, or simply "allocation mechanism," is a systematic procedure that assigns to each PHAP $P = (I, B, \succ, \triangleright, \pi, a) \in \mathcal{P}$ an outcome $\varphi(P)$ with the following properties:

- **Informational Constraint:** For every $i \in I$, $\mu_{\varphi(P)}(i)$ depends on, at most, history up to time $T_{\varphi(P)}(i)$ and future deterministic arrivals.[12]
- **Individual Rationality:** For all $i \in I$ such that $\varphi_P(i) \neq i$, i is acceptable for $\beta_{\varphi(P)}(i)$, and vice versa.

The informational constraint allows us to discipline our definition by stopping the mechanism from "peeking into the future," given that a PHAP contains arrivals that may be uncertain at a given time.[13]

Finally, given a matching outcome $\varphi(P)$, we denote by $\overline{\varphi}_P(i,t)$ applicant i's **interim expected matching outcome from time *t***, defined by replacing in a matching outcome the unit's delivery time with the sum of time t and the expected waiting time from t until delivery. Formally, $\overline{\varphi}_P(i,t) = i$ when the applicant is left unmatched and $\overline{\varphi}_P(i,t) = (\mu_{\varphi(P)}(i), t + \tau_{\varphi(P)}(\mu_{\varphi(P)}(i), t))$ otherwise. We discuss the interpretation of this definition in Section 2.5.4, in the context of our axioms

2.5 Axioms

A recent annual work plan from the Auditor General in Toronto (Romeo-Beehler, 2019) provides the following top recommendations for improving the waiting list to effectively manage access to subsidized housing:

> "1. Improve the integrity of the data and use of the waiting list to know exactly who is actively waiting and eligible for RGI assistance — so that units can be filled fairly and as quickly as possible 2. Improve matching of RGI applicants with vacant and available social housing units"

[11] The expected waiting time $\tau_{\varphi(P)}(\langle\beta, r\rangle, t')$ accounts for all the history up to time t', including $a_{t'}$.

[12] History up to time t includes B, I, $\succ$, $\triangleright$, π, as well as the history of play.

[13] To avoid unnecessary notation, we did not define an allocation mechanism as a sequence of per-period mechanisms (Bergemann and Välimäki, 2019; Doval, 2022). Such an approach structurally restricts the mechanism to use only presently available information.

We proceed to introduce a series of axiomatic properties that capture such objectives as articulated in policy documents of public housing authorities.

2.5.1 Strategic Robustness

A strategically robust mechanism is one where applicants always report their preferences truthfully:

Definition 7 (Strategy Proofness)**.** An allocation mechanism φ is **strategy-proof** if no applicant can ever benefit by unilaterally misrepresenting their preferences.[14]

This aligns with the desire expressed by policymakers in Toronto that the "City must ensure people understand the impacts of their housing choices" (Romeo-Beehler, 2019).

2.5.2 Efficiency

We study two aspects of efficiency: efficiency in the typical Pareto sense (where no situation exists where every applicant is weakly better off, with at least one applicant strictly better off) and efficiency in an administrative sense.

We introduce two notions of Pareto efficiency, which capture the desire for an improved matching between applicants and vacancies that become available (the second recommendation highlighted earlier):

Definition 8 (Ex-Post Pareto Efficiency)**.** Matching outcome $\varphi(P)$ is **ex-post Pareto dominated** by matching outcome $\phi(P)$ if $\phi_P(i) \succeq_i \varphi_P(i)$ for all $i \in I$ and there is some $i' \in I$ such that $\phi_P(i') \succ_{i'} \varphi_P(i')$. Matching outcome $\varphi(P)$ is **ex-post Pareto efficient** if it is not ex-post Pareto dominated by any other matching outcome. A public housing allocation mechanism φ is **ex-post Pareto efficient** if $\varphi(P)$ is an ex-post Pareto efficient matching outcome for all $P \in \mathcal{P}$.

Definition 9 (Interim Pareto Efficiency)**.** Matching outcome $\varphi(P)$ is **interim Pareto dominated** by matching outcome $\phi(P)$ if $\overline{\phi}_P(i, T_{\varphi(P)}(i)) \succeq_i \overline{\varphi}_P(i, T_{\varphi(P)}(i)))$ for all $i \in I$ and there is some $i' \in I$ such that $\overline{\phi}_P(i', T_{\varphi(P)}(i'))) \succ_{i'} \overline{\varphi}_P(i', T_{\varphi(P)}(i')))$. Matching outcome $\varphi(P)$ is **interim Pareto efficient** if it is not interim Pareto dominated by any other matching outcome. A public housing allocation mechanism φ is **interim**

[14] Equivalently, truth-telling is a weakly-dominant strategy for every applicant in the preference revelation game induced by φ (Balinski and Sönmez, 1999).

Pareto efficient if $\varphi(P)$ is an interim Pareto efficient matching outcome for all $P \in \mathcal{P}$.

The second aspect of efficiency we investigate is in an administrative sense:

Definition 10 (Non-Wastefulness)**.** A matching outcome $\varphi(P)$ is **non-wasteful** when the following is true for all $a_t \neq \emptyset$: If there exists $i \in I \setminus \left\{i' \in I : T_{\varphi(P)}(i') < t\right\}$ such that $a_t \succ_i i$ and $i \rhd_{a_t} a_t$, then (i) $i_{\varphi(P)}(a_t) \neq a_t$ and (ii) $t_{\varphi(P)}(i_{\varphi(P)}(a_t)) = t$. A public housing allocation mechanism φ is **non-wasteful** if $\varphi(P)$ is non-wasteful for all $P \in \mathcal{P}$.

Specifically, a non-wasteful mechanism is one that only produces matching outcomes where, if a suitable unmatched applicant exists, (i) units are not discarded, and (ii) units are not withheld after their arrival date.[15] The non-wastefulness property captures realistic institutional objectives in public housing allocation.[16] As the Auditor General's report in Toronto states, "The demand for financial assistance for housing in the City far exceeds the supply of RGI social housing units… Any vacancies or delays in filling available RGI housing… means people, including individuals and families designated as a priority, are going longer without stable housing."[17]

2.5.3 Fairness

We study fairness in the following sense: if an applicant i desires the applicant's i' assignment over their own, then i' must have a higher priority.

Definition 11 (Ex-Post Elimination of Justified Envy)**.** Applicant i **ex-post justifiably envies** applicant i' in matching outcome $\varphi(P)$ if $\varphi_P(i') \succ_i \varphi_P(i)$ and $i \rhd_{\beta_{\varphi(P)}(i')} i'$. A matching outcome $\varphi(P)$ is **ex-post free of justified envy** if no applicant ex-post

[15]Our notion of non-wastefulness is different from that of Balinski and Sönmez (1999). It relates more closely to models that assume units must be allocated immediately upon arrival (David and Yechiali, 1990; Gershkov and Moldovanu, 2010, 2012; Andersson, Ehlers and Martinello, 2018; Arnosti and Shi, 2020), including settings in which units are perishable such as organs (Shi and Yin, 2022). The computer science literature (e.g., Karp, Vazirani and Vazirani 1990) uses the term "online matching" to describe such problems.

[16]Efficiency considerations play an important role in motivating this property since leaving units vacant clearly decreases utilitarian welfare (Leshno, 2022).

[17]Consistent with the importance of this objective, the state of Massachusetts "levied $4.1 million in fines against 212 agencies from 2019 to 2022" because "nearly 2,300 of 41,500 state subsidized units were vacant, even as more than 184,000 people languished on the waitlist for public housing" (WBUR Newsroom, 2024).

envies another one. A public housing allocation mechanism φ **eliminates ex-post justified envy** or is **ex-post free of justified envy** if it is ex-post free of justified envy for all $P \in \mathcal{P}$.

Definition 12 (Interim Elimination of Justified Envy)**.** Applicant i **ad interim justifiably envies** applicant i' in matching outcome $\varphi(P)$ if $\overline{\varphi}_P(i', T_{\varphi(P)}(i)) \succ_i \overline{\varphi}_P(i, T_{\varphi(P)}(i))$ and $i \triangleright_{\beta_{\varphi(P)}(i')} i'$. A matching outcome $\varphi(P)$ is **interim free of justified envy** if no applicant ad interim envies another one. A public housing allocation mechanism φ **eliminates interim justified envy** or is **interim free of justified envy** if it is interim free of justified envy for all $P \in \mathcal{P}$.

This aligns with the first recommendation highlighted earlier, which emphasizes the importance of having accurate information about who is actively waiting and eligible for housing to facilitate making informed decisions about who should receive housing based on their needs and priorities.

2.5.4 Discussion of the Axioms

Our ex-post notions are based on comparing whether applicants prefer others' assignments and delivery times over their own. We call this view "ex-post" because a complete comparison requires realizing every applicant's assigned unit and delivery time. In contrast, our interim notions take the perspective of an applicant's assignment time. Then, by considering the expected waiting from that point until delivery allows the applicant to evaluate their preference over future uncertain events. More specifically, one way of thinking about our interim notions is that applicants consider, once selected for assignment, whether they would prefer to wait for a different unit. We acknowledge that different interim or ex-ante notions can be defined; in Section 4.2, we compare our interim notions to alternatives.

Our definition of ex-post Pareto efficiency is similar in spirit to the typical definition in the static literature (c.f. Abdulkadiroğlu and Sönmez, 1999; Ergin, 2002; Balinski and Sönmez, 1999): no matching outcome exists where at least one applicant can be strictly improved without harming the rest. However, our framework features a significant difference with the static literature stemming from the dynamic nature of the problem. Unlike in static problems where applicants' outcomes depend solely on the unit they receive, the unit's delivery time in our framework also affects applicants. Our definition of ex-post Pareto efficiency captures the potential for Pareto improvements resulting from this dynamic consideration.

The time dimension in our model, in fact, creates an interesting connection between Pareto efficiency and non-wastefulness. To see it, consider any unit, say a_t, that is assigned to an applicant, say i, under a given mechanism. Beyond assignment, the mechanism has two options for a_t: to deliver it upon arrival, at t, or to deliver it after arrival, at $t' > t$. However, if the mechanism delivers a_t at $t' > t$, it cannot be ex-post Pareto efficient due to the costly-waiting property of applicants' preferences. The reason is that a matching outcome involving a delayed delivery is ex-post Pareto dominated by the outcome in which every other applicant's assignment and delivery time is unchanged but i receives a_t at t, leading to a strict improvement for i while the rest of the applicants are unaffected. The following Lemma arises immediately:

Lemma 1. *Suppose $\varphi(P)$ is an ex-post Pareto efficient outcome. For all $a_t \neq \emptyset$ such that there exists $i \in I \setminus \left\{i' \in I : T_{\varphi(P)}(i') < t\right\}$ with $a_t \succ_i i$ and $i \rhd_{a_t} a_t$: (i) $i_{\varphi(P)}(a_t) \neq a_t$ implies (ii) $t_{\varphi(P)}(i_{\varphi(P)}(a_t)) = t$.*

A final observation is that by focusing on individually rational mechanisms, our notion of ex-post elimination of justified envy can be interpreted as a notion of stability, as it prevents blocking pairs from arising. A similar relationship appears in the context of a many-to-one student placement problem (Balinski and Sönmez, 1999; Romm, Roth and Shorrer, 2021).

3 Ex-Post Impossibility

Our first result investigates the ex-post possibility of an allocation mechanism with desirable efficiency and fairness properties:

Theorem 1 (Ex-Post Impossibility)**.** *From an ex-post perspective, a non-wasteful public housing allocation mechanism is neither Pareto efficient nor free of justified envy.*

We defer to the appendix the case of Pareto efficiency, which is analogous to the justified envy proof below.

Proof. Consider any given non-wasteful mechanism φ, along with the following setup:

- **Applicants.** $I^* = \{i_1, i_2\}$.

- **Buildings.** $B^* = \{\beta_1, \beta_2\}$.
- **Preferences.** $\succ^*$ is a given preference profile such that

$$i_1 : (\beta_1, 0) \succ^*_{i_1} (\beta_1, 1) \succ^*_{i_1} (\beta_2, 0) \succ^*_{i_1} (\beta_2, 1) \succ^*_{i_1} i_1$$
$$i_2 : (\beta_2, 0) \succ^*_{i_2} (\beta_2, 1) \succ^*_{i_2} (\beta_1, 0) \succ^*_{i_2} (\beta_1, 1) \succ^*_{i_2} i_2$$

- **Priorities.** $\rhd^*$ is a given priority profile such that $i_2 \rhd^*_\beta i_1 \rhd^*_\beta \beta$ for all $\beta \in B^*$.
- **Arrival Distribution.** π^* is a given MAP distribution over B^*.
- **Realized Arrivals.** $A^* = \{a \in A(\pi^*) : a_0 = \beta_1\}$.

Consider the PHAP subset

$$\mathcal{P}^* = \{P \in \mathcal{P} : I = I^*, B = B^*, \succ = \succ^*, \rhd = \rhd^*, \pi = \pi^*, a \in A^*\}.$$

Given that both applicants are acceptable at both buildings and vice versa, non-wastefulness of φ implies that in any $P \in \mathcal{P}^*$, arrival a_0 must be assigned and delivered at $t = 0$. Moreover, all the PHAPs in $\mathcal{P}^*$ have the same applicants, buildings, preferences, priorities, distribution over arrivals, and first arrival; they only differ in $(a_t)_{t \in \mathbb{N}}$. Consequently, the informational constraint of φ implies that $i_{\varphi(P)}(a_0) \in \{i_1, i_2\}$ must be the same for all $P \in \mathcal{P}^*$. This leads to only two possibilities: $i_{\varphi(P)}(a_0) = i_1$ for all $P \in \mathcal{P}^*$ or $i_{\varphi(P)}(a_0) = i_2$ for all $P \in \mathcal{P}^*$.

Case 1: $i_{\varphi(P)}(a_0) = i_1$ for all $P \in \mathcal{P}^*$. As i_2 is acceptable at both buildings and vice versa, φ's non-wastefulness implies that $i_{\varphi(P)}(a_1) = i_2$ for all $P \in \mathcal{P}^*$. This is particularly true for any $P_1^* \in \mathcal{P}^*$ featuring $a_1 = \beta_1$. It follows that i_2 ex-post justifiably envies i_1 in outcome $\varphi(P_1^*)$ because $\varphi_{P_1^*}(i_1) = (\langle \beta_1, 1 \rangle, 0) \succ^*_{i_2} (\langle \beta_1, 2 \rangle, 1) = \varphi_{P_1^*}(i_2)$ and $i_2 \rhd^*_{\beta_1} i_1$.

Case 2: $i_{\varphi(P)}(a_0) = i_2$ for all $P \in \mathcal{P}^*$. As i_1 is acceptable at both buildings and vice versa, φ's non-wastefulness implies that $i_{\varphi(P)}(a_1) = i_1$ for all $P \in \mathcal{P}^*$. This is particularly true for any $P_2^* \in \mathcal{P}^*$ featuring $a_1 = \beta_2$. It follows that i_2 ex-post justifiably envies i_1 in outcome $\varphi(P_2^*)$ because $\varphi_{P_2^*}(i_1) = (\langle \beta_2, 1 \rangle, 1) \succ^*_{i_2} (\langle \beta_1, 1 \rangle, 0) = \varphi_{P_2^*}(i_2)$ and $i_2 \rhd^*_{\beta_2} i_1$. □

Note that Theorem 1 is slightly stronger than the typical impossibility results in the market design literature. For instance, Roth (1982), Alcalde and Barberà (1994), and Balinski and Sönmez (1999) point out the tension that arises in strategy-proof mechanisms between Pareto efficiency and stability, which cannot hold simultaneously.

Meanwhile, we show that ex-post Pareto efficiency is impossible even in the absence of justified envy and vice versa. In the literature, the impossibility typically arises from the model's primitives. Our proof, however, highlights a novel economic tension between non-wastefulness and ex-post fairness and efficiency, resulting from the dynamic nature of the problem.

Specifically, due to the informational constraint, a non-wasteful allocation mechanism must assign and deliver a unit upon its arrival, at the latest, while uncertainty remains about future stochastic arrivals. As a result, the mechanism cannot always prevent the present cases of justified envy or Pareto-enhancing trading cycles arising from future arrivals.

In fact, non-wastefulness plays a prominent role in the result. For instance, assume that the mechanism in the proof of theorem 1 can withhold the first unit. Then, in case 2, if φ waits until $a_0 = \beta_1$ and $a_1 = \beta_2$ are realized and then, at $t = 1$, it assigns a_0 to i_1 and a_1 to i_2, applicant i_2's envy toward applicant i_1 vanishes. However, we should point out that violating non-wastefulness to avoid ex-post fairness or efficiency issues may come at the expense of new violations to these properties. In this example, delaying the delivery of a_0 opens up the possibility of a Pareto improvement: compared to the delayed delivery, applicant i_1 strictly prefers the matching outcome that, all else equal, delivers a_0 when it becomes available at $t = 0$, while applicant i_2 remains unaffected by the change.

A popular approach in the literature to address the impossibility of mechanisms with desirable properties is to restrict the model's primitives (for instance, Alcalde and Barberà, 1994; Ergin, 2002; Sönmez and Ünver, 2010*b*). This approach may not be realistic or fruitful in our framework from a practical perspective. On the one hand, a policymaker will likely handle only buildings' priorities. On the other hand, the impossibility in our framework stems from two sources: first, the stochastic nature of PHAPs and, second, the richness of the preference space, as is typically the case in the literature. Therefore, overcoming the impossibility in our framework likely requires controlling the stochastic nature of the unit arrival and the preference space to some degree.[18] For instance, in a related static deterministic assignment problem with priorities and quotas, Ergin (2002) shows that acyclicity is sufficient for a strategy-proof Pareto efficient and consistent mechanism. Meanwhile, the priority

[18]For instance, in the extreme and implausible case where the PHAP space trivially contains only homogeneous preference profiles, no ex-post Pareto-improving trading cycle ever arises.

profile in our proof is acyclic, and the impossibility is still attained.

A second insight from Theorem 1 pertains to the efficiency trade-offs that policy-makers face in practice. On the one hand, Theorem 1 implies that only mechanisms violating non-wastefulness can be ex-post Pareto efficient. On the other hand, Lemma 1 establishes that a mechanism that withholds units after arrival cannot be ex-post Pareto efficient. Therefore, a mechanism can only be ex-post Pareto efficient if it is wasteful in the sense of discarding units in at least one PHAP. The following result formalizes this insight:

Corollary 1. *If φ is ex-post Pareto efficient there exists $P \in \mathcal{P}$ such that for some $a_t \neq \emptyset$ and $i \in I \setminus \left\{i' \in I : T_{\varphi(P)}(i') < t\right\}$, we have $a_t \succ_i i$, $i \triangleright_{a_t} a_t$, and $i_{\varphi(P)}(a_t) = a_t$.*

A different and novel approach is to acknowledge the impossibility and design instead assignment mechanisms that lead to outcomes that minimize the violation of desirable properties (e.g. Pereyra, 2013; Kwon and Shorrer, 2020; Abdulkadiroğlu et al., 2020; Abdulkadiroğlu and Grigoryan, 2021; Doğan and Ehlers, 2021, 2022; Abdulkadiroğlu and Grigoryan, 2023; Combe, 2023). In this spirit, we investigate the extent of the impossibility in our framework, to determine whether the violations in the proof of Theorem 1 are a genericity or more of a systematic problem.

The function $F(\varphi(P))$ counts the instances of ex-post justified envy that arise in matching outcome $\varphi(P)$:

$$F(\varphi(P)) = \sum_{i \in I} |\{i' \in I : i' \text{ ex-post justifiably envies } i \text{ in outcome } \varphi(P)\}|.$$

On the other hand, the function $E(\varphi(P))$ counts, among Pareto-dominant matching outcomes, the largest number of applicants that can be strictly improved concerning matching outcome $\varphi(P)$. Letting $\mathcal{D}(\varphi(P))$ be the set of non-wasteful matching outcomes that ex-post Pareto dominate $\varphi(P)$,

$$E(\varphi(P)) = \max_{\phi(P) \in \mathcal{D}(\varphi(P)) \cup \{\varphi(P)\}} |\{i \in I : \phi_P(i) \succ_i \varphi_P(i)\}|.$$

By defining the optimization over $\mathcal{D}(\varphi(P)) \cup \{\varphi(P)\}$ we avoid an indeterminacy when $\mathcal{D}(\varphi(P)) = \emptyset$, in which case our definition naturally leads to $E(\varphi(P)) = 0$. Based on $F(\varphi(P))$ and $E(\varphi(P))$, we measure the significance of the impossibility of a non-wasteful mechanism as a function of the number of applicants involved in the following way:

Definition 13. The **extent of the ex-post impossibility** of a public housing allocation mechanism φ when it features n applicants is

$$V(\varphi, n) = \max_{\{P \in \mathcal{P} : |I| = n\}} \Big(E(\varphi(P)) + F(\varphi(P)) \Big).$$

Theorem 2 (Extent of the Ex-Post Impossibility)**.** *For any non-wasteful public housing allocation mechanism φ and number of applicants n, we have that $V(\varphi, n) \geq \lfloor \frac{n}{2} \rfloor$.*[19]

The proof of Theorem 2 iterates the idea behind Theorem 1 over replicas of a suitable applicant set. As such, it is involved and thus deferred to the appendix. Theorem 2 establishes that any non-wasteful mechanism has at least one possible outcome leading to as many violations to ex-post Pareto efficiency and elimination of justified envy as roughly half the applicants involved. Hence, regardless of the assignment mechanism, the impossibility is significant and has critical practical consequences, given that public housing authorities typically deal with assignment problems featuring thousands or millions of applicants.

4 Interim Possibility

4.1 The Choice-Based Waiting List Algorithm

Based on the extent of the ex-post impossibility, we focus on designing a mechanism with desirable properties from an interim perspective. This section outlines the conditions under which an interim fair and efficient mechanism exists. To that end, we define an algorithm, the **choice-based waiting list (CBWL) algorithm**, that consistently implements a matching outcome with the properties above. Given $P = (I, B, \succ, \rhd, \pi, A(\pi))$:

Step 0: Initialization

i. Collect all applicants' preferences over building-time pairs.

ii. Create an empty queue Q_β for each building $\beta \in B$ arriving with strictly positive probability under π and an unordered waiting list L with all the applicants in I.

Step 1: Arrival of units and individual rationality of potential matches

[19] $\lfloor x \rfloor$ is the floor function, representing the greatest integer number smaller than x.

i. Terminate the algorithm if $L = Q_\beta = \emptyset$ for all $\beta \in B$, no units are left to arrive with strictly positive probability, or no individually rational matching with respect to applicants' reported preferences exists in L for the units that still have a strictly positive probability of arrival.

ii. Otherwise, move forward to the next period with a unit arrival. Denote the arriving unit by $\langle \beta_A, r_A \rangle$.

Step 2: Process queue for building β_A

i. If $Q_{\beta_A} \neq \emptyset$, assign $\langle \beta_A, r_A \rangle$ to the applicant at the top of the queue, remove them from Q_{β_A}, and return to Step 1.

ii. If $Q_{\beta_A} = \emptyset$ and no individually-rational matching applicant for β_A exists in L, discard the unit and restart Step 1. Otherwise, proceed to Step 3.

Step 3: Locate β_A's highest-priority applicant i^* and their preferred building β^*

i. Locate the applicant $i^* \in L$ with the highest priority at building β_A. If i^* is the only applicant in L, add i^* at the end of Q_{β_A}, remove them from L, and return to Step 2.

ii. Otherwise, determine i^*'s presently highest-ranked queue among buildings where they are acceptable and denote the corresponding building by β^*. If $\beta^* = \beta_A$, add i^* at the end of Q_{β_A}, remove them from L, and return to Step 2. If $\beta^* \neq \beta_A$ instead, proceed to Step 4.

Step 4: Locate β^*'s highest-priority applicant i'

i. Locate the applicant $i' \in L$ with the highest priority at building β^*. If $i' \neq i^*$, place i' at the end of their presently highest-ranked queue among buildings where they are acceptable according to their reported preferences, remove i' from L, and return to Step 2. Otherwise, if $i' = i^*$, add i^* to the end of Q_{β^*}, remove them from L, and return to Step 3.

We denote by $\varphi^{\text{CBWL}}(P)$ the matching outcome implemented by the CBWL algorithm in PHAP P. Finally, we say that φ is a **choice-based waiting list (CBWL) mechanism** whenever $\varphi(P) = \varphi^{\text{CBWL}}(P)$ for all $P \in \mathcal{P}$. To show how the CBWL algorithm works, we present an example.

Consider the following PHAP $P \in \mathcal{P}$: $I = \{i_1, i_2, i_3, i_4\}$ and $B = \{B_1, B_2\}$; there are only three arrivals that happen with certainty: $a_0 = \beta_1$, $a_1 = \beta_2$, $a_2 = \beta_1$, and $a_3 = \beta_1$; buildings' priorities are $i_1 \rhd_{\beta_1} i_2 \rhd_{\beta_1} i_3 \rhd_{\beta_1} \beta_1$ and $i_2 \rhd_{\beta_2} i_1 \rhd_{\beta_2} i_3 \rhd_{\beta_2} i_4 \rhd_{\beta_2} \beta_2$; finally, both buildings are acceptable for all applicants, who have the following preferences:

$$i_1 : (\beta_2, 0) \succ (\beta_2, 1) \succ (\beta_1, 0) \succ (\beta_2, 2) \succ (\beta_1, 1) \succ (\beta_1, 2)$$
$$i_2 : (\beta_1, 0) \succ (\beta_1, 1) \succ (\beta_2, 0) \succ (\beta_1, 2) \succ (\beta_2, 1) \succ (\beta_2, 2)$$
$$i_3 : (\beta_2, 0) \succ (\beta_1, 0) \succ (\beta_2, 1) \succ (\beta_1, 1) \succ (\beta_2, 2) \succ (\beta_1, 2)$$
$$i_4 : (\beta_2, 0) \succ (\beta_2, 1) \succ (\beta_2, 2) \succ (\beta_1, 0) \succ (\beta_1, 1) \succ (\beta_1, 2)$$

A simplifying observation is that because π is a deterministic arrival process, the expected waiting time until a unit's arrival from time T is simply the remaining time from that point until the arrival. Mathematically, for $T > 0$ and a unit $\langle \beta, r \rangle$ arriving at time t, we have $\tau_{\varphi(P)}(\langle \beta, r \rangle, T) = t - T$.

We start the algorithm with Step 0. First, we collect the applicants' preferences; for now, we assume that applicants report their true preferences. Then, we initialize the buildings' queues $Q_{\beta_1} = Q_{\beta_2} = \emptyset$ and the general waiting list $L = \{i_1, i_2, i_3, i_4\}$. We proceed to Step 1. As there are four units left to arrive and four applicants in L are potential individually rational matches for them, the algorithm continues to $t = 0$ where $\langle \beta_A, r_A \rangle = \langle \beta_1, 1 \rangle$ arrives.

No applicant has yet been placed in the queue for β_A ($Q_{\beta_A} = \emptyset$), and an individually rational applicant for β_A still exists in L, so we proceed directly to Step 3. In Step 3, we locate the applicant with the highest priority at β_A among L, which is $i^* = i_1$. Along with i^*, there are three other applicants in L, so we proceed to Part (ii) of Step 3, where we determine i^*'s presently highest-ranked queue. On the one hand, if i^* is placed in Q_{β_A}, the applicant receives a unit in β_1 at $t = 0$, as the algorithm returns to Step 2, where the currently available unit is awarded to i^* now waiting in Q_{β_A}. On the other hand, if i^* is placed in Q_{β_B}, given the rules of the algorithm and the arrival process, the applicant receives a unit in β_2 at $t = 1$. According to i^*'s preferences, $(\langle \beta_2, 1 \rangle, 1) \succ_{i^*} (\langle \beta_A, r_A \rangle, 0)$. Therefore, $\beta^* = \beta_2 \neq \beta_A$, and we must proceed to Step 4.

In Step 4, we locate the applicant in L with the highest priority at β^*, which is $i' = i_2 \neq i^*$. Next, we must determine a queue that maximizes the preferences of i', who faces the same options outlined for i^*, as no applicant has been placed on either queue. In this case, however, $(\langle \beta_A, r_A \rangle, 0) \succ_{i'} (\langle \beta_2, 1 \rangle, 1)$. Therefore, we update $Q_{\beta_A} = i'$, $L = \{i_1, i_3, i_4\}$, and return to Step 2, where $\langle \beta_A, r_A \rangle$ is awarded to i', who is the top applicant in Q_{β_A}. Finally, we remove i' from Q_{β_A} and return to Step 1.

Per Step 1, as $L = \{i_1, i_3, i_4\}$ and there are three units left to arrive, we move on to $t = 1$, where $a_1 = \beta_2$ arrives. By repeating the algorithm, as described in the case of a_0, we assign a_1 and a_2 as the following table summarizes; the algorithm finalizes after two more iterations, when $L = \{i_4\}$, as only $a_3 = \beta_1$ is left to arrive with strictly positive probability but i_4 is not acceptable for it.[20]

i	$\mu_{\varphi^{\text{CBWL}}(P)}$	$T_{\varphi^{\text{CBWL}}(P)}$	$\tau_{\varphi^{\text{CBWL}}(P)}$	$t_{\varphi^{\text{CBWL}}(P)}$
i_1	$\langle \beta_2, 1 \rangle$	1	0	1
i_2	$\langle \beta_1, 1 \rangle$	0	0	0
i_3	$\langle \beta_1, 2 \rangle$	2	0	2
i_4	i_4	3	N/A	N/A

This example reveals helpful insights about the subsequently induced matching outcomes that generalize to any PHAP. First, the CBWL mechanism is strategy-proof. This follows because, in any PHAP, the CBWL algorithm chooses applicants mainly based on their exogenous priorities, and then, it attempts to assign applicants to their reported preference-maximizing queue. Therefore, any preference misreport cannot affect the order in which the mechanism chooses applicants and cannot strictly improve their outcome.[21] Second, any matching outcome produced by the CBWL algorithm is non-wasteful. This is true because the algorithm discards only the units for which no individually rational unmatched applicant is left; the rest of the units are assigned at or before their arrival and delivered upon arrival.

Third, the CBWL mechanism is interim free of justified envy. This is true because the only applicant who would prefer another applicant's expected matching outcome is i_3, who prefers i_2's matching outcome. However, i_2 holds a higher priority than i_3 at β_1, eliminating any justification for interim envy. Finally, the CBWL mechanism is interim Pareto efficient. In this example, we can easily see it among the class of non-wasteful mechanisms because outcomes of non-wasteful mechanisms are simple reassignments of the units without affecting their delivery times.[22] The analysis is simplified further because the deterministic arrival process reduces our interim notion

[20] As acceptability is atemporal for applicants, the time at which i_4 learns that they will remain unmatched is immaterial for the results. Here, we assume they are notified once all units have been awarded.

[21] At first glance, one may wonder if i^* can affect the CBWL algorithm's choice of i' by misreporting their preferences. In the proof of theorem 3 in the appendix, we show that due to acyclicity, doing so can lead, at most, to i^* being placed on a sub-optimal queue.

[22] The CBWL is interim Pareto efficient beyond the class of non-wasteful mechanisms. The proof, however, is beyond the scope of this example and, thus, deferred to the Appendix.

to ex-post.[23] In the example, the only applicant who would strictly improve from a reallocation is i_3. However, any potential reallocation involving i_2 is detrimental, as the CBWL mechanism awards them their most desirable option. Moreover, trade between i_3 and i_1 is not beneficial for either of them.

Generally, the CBWL mechanism's interim Pareto efficiency and elimination of justified envy hinge on the CBWL algorithm maximizing applicants' preferences to determine their queues. To see it, suppose that the CBWL algorithm assigns applicant i unit $\langle \beta, r \rangle$ at time T and applicant i' unit $\langle \beta', r' \rangle$ at time $T' > T$. Then, i strictly prefers their expected outcome over the unit of i' assigned at time T. This is true because when i is assigned to a queue, $\langle \beta', r' \rangle$ is still available. If $\beta = \beta'$, it means that $r < r'$, so the waiting time from T for r' is strictly longer and i strictly prefers $\langle \beta, r \rangle$ due to preferences' costly-waiting property. Meanwhile, if $\beta \neq \beta'$, the CBWL algorithm's queue choice for i reveals that the applicant strictly prefers $\langle \beta, r \rangle$ over $\langle \beta', r_i \rangle$ for some $r_i \leq r'$. Consequently, no applicant ever prefers the unit over their own of an applicant who received their assignment at a later date.

The above observation is crucial for the elimination of justified envy because Steps 3 and 4 of the CBWL algorithm ensure that applicants awarded a queue at an earlier date have a higher priority for their building than the remaining applicants, eliminating any possible case of interim justified envy. Likewise, the same observation is crucial for Pareto efficiency because it implies that in any PHAP, the first applicant, time-wise, whose expected matching outcome under the CBWL algorithm is different from their expected matching outcome under another mechanism strictly prefers the former matching outcome. This is true because the alternative mechanism assigns to them a unit that the CBWL algorithm assigns to a different applicant at a later time.

Theorem 3 formalizes the generality of the properties in the example:

Theorem 3 (Interim Possibility). *If $\triangle$ is an acyclic priority space, the CBWL mechanism is non-wasteful, strategy-proof, interim free of justified envy, and interim Pareto efficient.*

A final yet important point formalized by Theorem 4 is that an acyclic priority space is necessary for the existence of a non-wasteful, strategy-proof, and interim Pareto efficient and free of justified envy mechanism. The proof, deferred to the Appendix,

[23] We emphasize that the interim and ex-post comparisons generally differ when the arrival process is stochastic.

is constructive and shows that no matching outcome can have the aforementioned properties if a non-acyclic priority ordering exists. In the above example, it can be verified that if $i_1 \rhd_{\beta_1} i_2 \rhd_{\beta_1} i_3 \rhd_{\beta_1} \beta_1$ and $i_3 \rhd_{\beta_2} i_1 \rhd_{\beta_2} i_2 \rhd_{\beta_2} i_4 \rhd_{\beta_2} \beta_2$, the CBWL algorithm's outcome becomes the following:

i	$\mu_{\varphi^{\text{CBWL}}(P)}$	$T_{\varphi^{\text{CBWL}}(P)}$	$\tau_{\varphi^{\text{CBWL}}(P)}$	$t_{\varphi^{\text{CBWL}}(P)}$
i_1	$\langle\beta_2, 1\rangle$	1	0	1
i_2	$\langle\beta_1, 2\rangle$	2	0	2
i_3	$\langle\beta_1, 1\rangle$	0	0	0
i_4	i_4	3	N/A	N/A

This outcome violates interim elimination of justified envy, as $\overline{\varphi}(i_3) = (\langle\beta_1, 1\rangle, 0) \succ_{i_2} (\langle\beta_1, 2\rangle, 2) = \overline{\varphi}(i_2)$, with $i_2 \rhd_{\beta_1} i_3$. In a more general sense, Theorem 3 and Theorem 4 together imply that acyclicity characterizes a non-wasteful, strategy-proof, and interim Pareto efficient and free of justified envy mechanism. This insight is formalized in Corollary 2.

Theorem 4 (Necessity of Acyclicity)**.** *A non-wasteful, strategy-proof, and interim Pareto efficient and free of justified envy mechanism exists only if* $\triangle$ *is an acyclic priority space.*

Corollary 2 (Acyclicity Characterization)**.** *A non-wasteful, strategy-proof, and interim Pareto efficient and free of justified envy mechanism exists if and only if* $\triangle$ *is an acyclic priority space.*

4.2 Alternative Time Perspectives for Fairness and Efficiency

4.2.1 Ex-Ante Perspective

One possibility is to evaluate a mechanism from an *ex-ante* perspective—before applicants interact with the mechanism. From an ex-ante perspective, applicants only know the distribution of arrivals, as units have not started arriving yet. Different realized arrival sequences can lead to different mechanism outcomes.[24] Consequently, from an ex-ante perspective, any PHAP induces for applicants a lottery over potential

[24]For instance, assume there are two buildings β_1 and β_2, one applicant i_1 that finds only β_1 acceptable, and one arrival, equally likely to be on either building. In a non-wasteful mechanism, from an ex-ante perspective, i_1 has a 50% chance of being matched with β_1 or remaining unmatched.

matching outcomes, a problem referred to as *fractional matching* or *random matching* in the literature (Roth, Rothblum and Vande Vate, 1993; Kesten and Ünver, 2015). We define a **random matching outcome** $\mathcal{L}(P, \varphi)$ as the lottery over matching outcomes induced by mechanism φ in PHAP $P \in \mathcal{P}$.[25]

The literature documents how random matching ex-ante properties can be tied to the properties of their support. For instance, Vande Vate (1989), Rothblum (1992), Roth, Rothblum and Vande Vate (1993), and Kesten and Ünver (2015) show that a random matching is ex-ante stable only if its support entirely consists of stable matching outcomes. We define our ex-ante notions following a similar idea:

Definition 14 (Ex-Ante Pareto Efficiency)**.** Given a mechanism φ and a PHAP $P \in \mathcal{P}$, $\mathcal{L}(P, \varphi)$ is an **ex-ante Pareto efficient random matching outcome** if it has a support that includes only ex-post Pareto efficient matching outcomes. A public housing allocation mechanism φ is *ex-ante Pareto efficient* if for any $P \in \mathcal{P}$, $\mathcal{L}(P, \varphi)$ is an ex-ante Pareto efficient random matching outcome.

Definition 15 (Ex-Ante Elimination of Justified Envy)**.** Given a mechanism φ and a PHAP $P \in \mathcal{P}$, $\mathcal{L}(P, \varphi)$ is **ex-ante free of justified envy random matching outcome** if it has a support that includes only ex-post free of justified envy matching outcomes. A public housing allocation mechanism φ *eliminates ex-ante justified envy* or is *ex-ante free of justified envy* if for any $P \in \mathcal{P}$, $\mathcal{L}(P, \varphi)$ is an ex-ante free of justified envy random matching outcome.

The next result is a direct consequence of the counter-examples in the proof of Theorem 1. In the first PHAP, any mechanism is bound to deliver, with strictly positive probability, a Pareto inefficient matching outcome. In the second PHAP, any mechanism will lead, with strictly positive probability, to a matching outcome where at least one applicant experiences justified envy.

Theorem 5 (Ex-Ante Impossibility)**.** *From an ex-ante perspective, a non-wasteful public housing allocation mechanism is neither Pareto efficient nor free of justified envy.*

Two remarks are on point. First, Theorem 5 shows that the impossibility issues behind Theorem 1 are not generic because the arrival sequences that lead to them have

[25]Because lotteries over matching outcomes are computed before arrivals realize, $\mathcal{L}(P, \varphi) = \mathcal{L}(P', \varphi)$ whenever $P = (I, B \succ, \rhd, \pi, a)$ and $P' = (I, B \succ, \rhd, \pi, a')$.

probability bounded away from zero. Second, although our ex-ante notions are weaker than some of the analogs in the literature (for instance, Kesten and Ünver, 2015), our result suggests that a mechanism attaining such stronger analogs is impossible in our framework. This is because attaining our ex-ante notions is a necessary condition for attaining their stronger analogs.

4.2.2 Full-Swap Interim Perspective

Recall that we defined applicant i's expected matching outcome in $\varphi(P)$ from time t as follows:

$$\overline{\varphi}_P(i,t) = (\mu_{\varphi(P)}(i), t + \tau_{\varphi(P)}(\mu_{\varphi(P)}(i), t)).$$

In this definition, t can be interpreted as the **reference time** for applicants' preferences and expected times. Namely, $\overline{\varphi}_P(i,t)$ considers the assignment of unit $\mu_{\varphi(P)}(i)$ at time t with an expected waiting time of $\tau_{\varphi(P)}(\mu_{\varphi(P)}(i), t))$. For a given mechanism ϕ, and applicants i and i', our baseline interim notions of Pareto efficiency and elimination of justified envy (Definition 9 and Definition 12) consider the following preference inequalities, respectively:

$$\overline{\phi}_P(i, T_{\varphi(P)}(i)) \succeq_i \overline{\varphi}_P(i, T_{\varphi(P)}(i)), \tag{Base-PE}$$

$$\overline{\varphi}_P(i', T_{\varphi(P)}(i)) \succ_i \overline{\varphi}_P(i, T_{\varphi(P)}(i)). \tag{Base-JE}$$

When comparing two expected matching outcomes to verify for Pareto improvements or justified envy, our baseline interim notions hold the reference time constant at the applicant's assignment time. Our baseline notions, then, compare, *from the time $T_{\varphi(P)}(i)$ applicant i is chosen and assigned by the mechanism*, whether they to wait for their assigned option compared to waiting for a different assignment.

An alternative is to consider the case where an applicant is offered the possibility of swapping their assignment for another applicant's assignment, *including the assignment time.* This perspective amounts to ultimately trading places with another applicant, varying the comparison's reference time with each considered option. We call this perspective **full-swap (FS) interim**, where the inequalities in the baseline interim Pareto efficiency and justified envy are modified as follows:

$$\overline{\varphi}_P(i', T_{\varphi(P)}(i')) \succeq_i \overline{\varphi}_P(i, T_{\varphi(P)}(i)), \tag{FS-PE}$$

$$\overline{\varphi}_P(i', T_{\varphi(P)}(i')) \succ_i \overline{\varphi}_P(i, T_{\varphi(P)}(i)). \tag{FS-JE}$$

Note that in the Pareto efficiency FS-interim notion, we replace $\overline{\phi}_P$ with $\overline{\varphi}_P$. This is because we focus on assignment swaps only, so for any $i \in I$, there exists $i' \in I$ such that $\mu_{\phi(P)}(i) = \mu_{\varphi(P)}(i')$, which implies that $\overline{\phi}_P(i, T_{\varphi(P)}(i')) = \overline{\varphi}_P(i', T_{\varphi(P)}(i'))$.

Due to considering a different reference time on each comparison, the interaction of non-wastefulness and dynamic uncertainty behind Theorem 1 lead to similar issues from the FS-interim perspective:

Theorem 6 (Full-Swap Interim Impossibility)**.** *From an FS-interim perspective, a non-wasteful public housing allocation mechanism cannot be Pareto efficient and free of justified envy.*

5 Welfare Estimation

5.1 Public Housing Policy Context

Over 2 million people in the US reside in public housing, with more than 3,000 state or local government-run Public Housing Authorities (PHAs) administering federal housing assistance programs. To describe the policies and procedures governing the administration of public housing, the Department of Housing and Urban Development (HUD) issues a detailed manual for PHAs known as the Public Housing Occupancy guidebook (HUD, 2003).

The guidebook discusses "two generally accepted approaches to unit offers—Plan A and Plan B." Under Plan A, the PHA offers a unit that becomes available to the applicant with the highest priority, and a refusal results in the applicant's removal from or placement at the bottom of the waiting list. Plan B allows applicants to decline an offer but still remain eligible for an additional offer. Kaplan (1984) refers to these as *k-strike refusal systems*, where $k \in \mathbb{N}$ denotes the maximum total number of units that the PHA offers to a given applicant.

More broadly, the guidebook describes "two waiting list approaches": a single community-wide or centralized waiting list (CWL) or multiple site-based waiting lists (SBWL). About 75 percent of PHAs implement CWL, following Plan A or Plan B (Devine, Rubin and Gray, 1999). The remaining PHAs operate a SBWL system, which requires annual approval from HUD. Under SBWL, applicants indicate a set of locations, potentially with a limit, where they are willing to accept offers and then receive the next available unit from that set. We refer to CWL with a k-strike refusal system as the CWL-k mechanism, and we refer to SBWL with a limit of k location choices as the SBWL-k mechanism. For our theoretical result below, we further distinguish between the *ultimatum* case where the CWL-k penalizes the k-th refusal by removing the applicant from the list and the *demotion* case where the CWL-k penalizes the k-th refusal by placing the applicant at the bottom of the list. We note that a CWL-∞ waiting list system constitutes a First-Come, First-Serve (FCFS) allocation mechanism, in which each applicant may reject as many units as they wish without forfeiting their position on the waiting list. In addition, SBWL-∞ denotes a site-based waiting list system that does not impose any limit on the number of location choices.

The existing waiting list approaches suffer from notable shortcomings. Firstly, any CWL-k mechanism or SBWL-k leads to interim fairness and efficiency violations. Secondly, any SBWL-k and Plan B (i.e., $k \geq 2$) results in violations of non-wastefulness. Theorem 7 below formalizes these claims.

Theorem 7. *The Public Housing Occupancy guidebook's mechanisms have the following theoretical properties:*

i. Every demotion-CWL mechanism violates non-wastefulness, interim Pareto efficiency, and suffers from interim justified envy.

*ii. Every ultimatum-CWL mechanism violates interim Pareto efficiency and suffers from interim justified envy. Moreover, among ultimatum-CWL mechanisms, only ultimatum-CWL-*1 *is non-wasteful.*

iii. Every SBWL-k violates non-wastefulness, interim Pareto efficiency, and suffers from interim justified envy.

The strategy-proofness of the CWL and SBWL mechanisms requires extending our framework beyond the scope of this paper. On the one hand, our strategy-proofness definition is suitable for direct mechanisms where preference revelation is the only interaction of the applicants with the mechanism. On the other hand, the CWL and SBWL are indirect mechanisms for which the applicability of a revelation principle is not apparent. One possibility is to extend our standard definition as proposed by Li (2017), namely, by defining a strategy-proof allocation mechanism as one that induces in any PHAP a game with a Bayesian Nash Equilibrium (BNE) in dominant strategies.[26] A practical issue with this notion of strategy-proofness is that applicants need to compute contingent scenarios whose complexity increases with the number of applicants and intricacies of the mechanism, making it unlikely to identify and play their dominant strategies. In fact, evidence shows that they are more likely to follow through with their dominant strategies in simpler mechanisms.[27] Consequently, the simplicity and (direct) strategy-proofness of the CBWL suggest that it may be superior to the CWL and SBWL mechanisms in a strategic sense, a claim that deserves further study.

The shortcomings of the CWL and SBWL mechanisms raise questions about why policymakers implement such waiting list systems. The Public Housing Occupancy guidebook explains the advantages of each system:

> "Under Plan A: The amount of time spent making offers to any applicant is limited to the time it takes to make one offer.
>
> ...
>
> Under Plan B: Applicants have greater choice of units.
>
> ...
>
> Site-Based Waiting Lists... applicants have full choice of locations.

The desire to achieve non-wastefulness motivates the selection of Plan A, whereas efficiency and envy-freeness considerations make Plan B more appealing. SBWL provides an alternative that increases choice further, though PHAs recognize the significant concerns that arise due to its failure to satisfy strategy-proofness. The New York City Housing Authority, for example, uses SBWL with a limit of two borough choices (SBWL-2) and explicitly advises applicants to "select their first borough choice carefully" due to "longer waiting lists and fewer vacancies in the boroughs of Manhattan and Queens."

The CBWL system proposed in Section 4 establishes a hybrid system that uses both types of waiting lists. In particular, households start in a centralized queue and receive an offer as in CWL but then join a site-specific queue if they refuse an offer. Comparing Theorem 3 and Theorem 7, we see that this combined system of waiting

[26] In a dominant BNE equilibrium, applicants cannot benefit from any additional information about how other players play the game or from playing any strategy other than their dominant strategy. Therefore, a dominant BNE equilibrium eliminates any benefit from "manipulating" the mechanism through deviations from the dominant strategy.

[27] The evidence is nicely framed and summarized by Li (2017) and Pycia and Troyan (2023).

lists offers considerable benefits from an axiomatic perspective. To highlight the value of the CBWL system for public housing allocation, we proceed to examine the welfare consequences of the different waiting list systems.

5.2 Data and Methodology

To evaluate ex-post welfare, we conduct simulations comparing allocations under CBWL with those from existing public housing allocation mechanisms. Estimating the welfare gains from using CBWL requires data on building arrivals as well as applicant preferences. We combine data on public housing developments from HUD with data from a choice experiment by Naik and Thakral (2023) to measure preferences for public housing. We focus on the set of PHAs that operate at least 3,000 units and that do not limit the number of locations where applicants may indicate a willingness to accept offers. This gives us a sample of 10 PHAs that collectively manage over 50,000 housing units: Akron, OH; Boston, MA; Cuyahoga County (Cleveland), OH; Dallas, TX; Detroit, MI; Los Angeles, CA; Louisville, KY; Miami, FL; Richmond, VA; and San Antonio, TX.

5.2.1 Buildings and Arrivals

The HUD data contain information about every public housing development in the United States. We focus on the 251 housing developments operated by the 10 large PHAs in our sample. The address type of each development allows us to infer whether a given development is a townhouse community or a high-rise/low-rise building, and we use a cutoff of 100 housing units to separate high-rise from low-rise buildings. The data also contain information about zip codes, which we map to coarser geographical areas. For example, within the jurisdiction of the Boston housing authority, we divide locations into central areas (Boston and Charlestown), southern areas (Roxbury, Dorchester, and Mattapan), eastern areas (East Boston), western areas (Brighton), and southwestern areas (Jamaica Plain, Roslindale, West Roxbury, and Hyde Park). Based on information about the number of units and the turnover rate in each building, we estimate a multinomial arrival process. A multinomial arrival process implies that the number of periods between successive arrivals follows a geometric distribution, consistent with Kaplan (1984).[28] In addition, we randomly assign priority ratings to applicants and assume that all buildings share a common priority ordering over applicants.

5.2.2 Applicants and Preferences

We operationalize applicant preferences in our simulations by drawing from a discrete choice experiment that elicits households' willingness to pay (WTP) for different characteristics of public housing developments. In particular, we use data on preferences over building type and location collected by Naik and Thakral (2023) using a Bayesian Adaptive Choice Experiment (BACE). BACE measures individual-level preference parameters with high precision by asking respondents to choose between pairs of alternatives optimized based on their previous responses (Drake, Thakral and Tô, 2022; Drake et al., 2024). The dataset consists of a total of 14,675 such choices from 587 respondents who were on a waiting list for public housing in the UK.

We assume that each applicant i maximizes the present-discounted value of utility, where their utility from residing in building j takes the form

$$u_{i,j} = \text{income}_i - \text{rent}_{i,j} + \alpha_i X_j + \varepsilon_{i,j}, \tag{1}$$

where the vector X_j consists of building j's type and location, α_i represents the strength of applicant i's preferences for building characteristics, and $\varepsilon_{i,j}$ captures idiosyncratic factors that affect utility. Rent payments constitute 30 percent of income for public housing residents. The idiosyncratic component of utility follows a type-I extreme-value distribution with an individual-specific scale parameter, which we obtain directly from the Naik and Thakral (2023) data along with individual-level estimates of the discount factor and α_i (households' WTP to reside in a preferred location, a low-rise building, and a house). We model the location of a building by its coarse geographical area, and we endow each applicant with a preferred location at random, in proportion to the housing stock in each location.

Completing the specification of preferences requires information on applicants' incomes and and outside options. Accordingly, we rely on the Naik and Thakral (2023) survey data, which include measures of household income and the characteristics of

[28] Kaplan (1984) argues that the length of time that a household lives in public housing follows the continuous-time analog of the geometric distribution, namely the exponential distribution.

their previous residence, including rental costs, building type, and the desirability of the location. We evaluate Equation (1) with the characteristics of each applicant's previous residence to define their utility from living in private housing.

5.2.3 Simulation Environment

We draw from the arrival process for each city 50 times and simulate the allocation of housing units under each waiting list system: CWL-1 (Plan A), CWL-2 (Plan B), SBWL, and CBWL.

Under CWL-1, each applicant receives only one offer and faces the straightforward decision problem of accepting a unit if and only if the outside option provides lower utility.

CWL-2 gives applicants the opportunity to refuse one offer without any penalty. This decision depends on the applicant's beliefs about higher-priority households on the waiting list. We assume that the applicant knows how many households have higher priority but does not know their preferences. Instead, the applicant knows the distribution of preferences and understands that any household that rejects an offer must either accept their next offer or remain unmatched and take their outside option. To estimate the expected utility of refusing an offer, we use numerical approximations, drawing from the distributions of preferences and future arrivals.

SBWL enables applicants to specify a set of buildings they consider acceptable, ensuring that they only receive an offer of a unit from this list. We assume that households truthfully report all acceptable options and thus match with the next acceptable unit.

CBWL allows applicants to choose which site-specific waiting list to join based on information about the arrival process. Since the arrival time of the r^{th} unit in a building follows a negative binomial distribution, we use closed-form expressions to compute the expected utilities from joining waiting lists.

According to data from HUD, households spend an average of 5.9–8.5 years and a median of 3–4.7 years in public housing (Lubell, Shroder and Steffen, 2003; McClure, 2018). Since a shorter length of stay reduces the benefit of allocation mechanisms that provide applicants with more choice, we conservatively assume a match duration of three years, irrespective of the allocation mechanism.

5.3 Simulation Results

To analyze welfare, we convert the difference in utility between each mechanism and CWL-1 to a monetary value. We do so by computing the equivalent variation (EV), or the transfer that each applicant would have to receive when public housing is assigned by CWL-1 that would give the applicant the same lifetime utility as their assignment under each alternative mechanism.

Replacing the allocation mechanism from CWL-1 to CBWL creates substantial average welfare gains in all 10 cities. In particular, the welfare benefit of CBWL ranges from $2,700 per applicant in Detroit to over $5,000 per applicant in Akron, as the right panel of Figure 1 shows. The left panel of the figure displays characteristics of the building arrival process that correlate with the average EV. The average number of arrivals each week, ranging from about 2 per week in Detroit to nearly 9 per week in Akron, features a strong positive relationship with the EV. Los Angeles and Richmond deviate from this trend, achieving large welfare gains from adopting CBWL despite having fewer weekly arrivals than Louisville and Miami. This occurs in part due to the high quality of public housing options in Los Angeles and Richmond, as these two cities have the highest shares of townhouse units (most desirable on average) and the lowest shares of high-rise units (least desirable on average). In addition, the high geographic concentration of the public housing stock in Louisville and Miami limits the potential for increasing utility by improving the match between applicants' preferred and assigned locations.

Adopting the CWL-2 or SBWL system, on the other hand, leads to minimal, if any, improvements relative to CWL-1. In some cases, we see a negative average EV because increasing choice allows some applicants to receive marginally better allocations while imposing more significant negative externalities on other applicants. For example, since SBWL ensures that applicants only receive acceptable offers, the mechanism may assign a unit to an applicant who would remain unmatched under CWL-1. This can create a small gain for that applicant but a larger loss for the household that the mechanism displaces. CBWL restricts the scope of such losses, as Figure 2 shows, because a household that loses the opportunity to receive a high-value match has the ability to choose a building-specific queue to maximize their utility. The length of each bar represents the average welfare gain in each city among those who benefit from adopting CBWL, as well as the average welfare loss among those who prefer their assignment under CWL-1. The the vertical width of each bar depicts the share of applicants gaining or losing under CBWL relative to CWL-1. We find the highest shares of applicants with negative EVs for CBWL in Los Angeles (14 percent) and Richmond (17 percent), likely stemming from the same factors that contribute to their higher average EVs.

To shed light on which types of applicants gain under each alternative mechanism, we explore the relationship between EVs and outside options (Figure 3). To measure a household's outside option, we compute their flow utility from residing in private housing. We find that adopting CBWL leads to a more equitable distribution of welfare gains from public housing, generating greater benefits for households with worse outside options. This pattern arises despite randomly drawing an independent priority ordering in each simulation. Meanwhile, SBWL (and CWL-2 to a lesser extent) predominantly benefits those with better outside options at the expense of those with worse outside options. We see similar relationships when analyzing each city separately.

6 Discussion

While the primary application in this paper considers the allocation of public housing, our model provides general lessons that extend beyond this domain. We present a parsimonious framework that highlights the fundamental tension between non-wastefulness and uncertainty, leading to the impossibility of attaining efficiency or fairness from an ex-post perspective. This goes beyond previous dynamic models, which often make restrictive assumptions such as homogeneous applicants' waiting costs or Markovian discrete arrival processes. We show that any mechanism can result in ex-post efficiency and fairness failures for at least nearly half of the participants. This underscores the significance of developing alternative perspectives that overcome the impossibility result, which can apply in a variety of dynamic stochastic allocation problems. We make progress on this problem by proposing an interim perspective from which a non-wasteful mechanism can be Pareto efficient and free of justified envy.

Our work identifies a simple, non-wasteful choice-based waiting list (CBWL) mechanism that implements fair and efficient outcomes from an interim perspective if and only if buildings' priorities are acyclic. The CBWL mechanism combines features from the most widely used systems for allocating public housing—centralized and site-based waiting lists—which significantly limit applicants' abilities to wait for offers that better align with their preferences. Our mechanism achieves efficiency and fairness by allowing applicants to wait for more desirable options while respecting buildings' priorities. Initially, the housing authority offers a unit to the household on the centralized waiting list with the highest priority. If the household rejects the offer, they are moved to their chosen site-specific waiting list after verifying that no violations to the priority of other unassigned households occur. Our empirical analysis demonstrates that the CBWL has the potential to deliver substantial welfare gains, given that over 1.2 million households reside in public housing in the US alone.

The simplicity of the CBWL is likely to have benefits beyond the theoretical implementation of interim fair and efficient matching outcomes. Prior work extensively documents that complex mechanisms will likely lead to player deviations from their optimal strategies, and the most widely used systems for allocating public housing allow for rich strategy spaces that may feature difficult strategy computation. In contrast, the CBWL mechanism reduces this complexity by requiring applicants only to report their preferences in a context where truthful reporting is a *strictly* dominant strategy.

Finally, our empirical approach to evaluating welfare gains from adopting CBWL in public housing allocation also provides a template for future work. The new data source on building arrivals creates the opportunity to analyze welfare under any alternative allocation mechanism in an expansive range of environments resembling any US public housing authority, given appropriate preference data. This paves the way for assessing many design-related tradeoffs quantitatively, including the optimal design of priority systems, the partitioning of units into waiting lists, the welfare consequences of heterogeneity in valuations versus waiting costs, and racial disparities in welfare. The approach also provides the means for incorporating other objectives, such as racial or economic integration of housing projects (Kaplan, 1987) or vacancy costs. Moreover, the methodology opens the door for research investigating behavioral biases in choices under complex mechanisms.

A natural direction for follow-up research is to apply the proposed mechanism in practice. Indeed, the mechanism has already been adopted recently in the context of refugee resettlement to match Ukrainian refugees with hosts (Farajzadeh et al., 2023). The success of this mechanism in the refugee resettlement context offers a compelling case for broader implementation, especially in contexts such as disaster response or emergency housing where the urgency of the situation demands appropriately designed systems that deliver timely and effective support.

References

Abdulkadiroğlu, Atila, and Aram Grigoryan. 2021. "Priority-based assignment with reserves and quotas." National Bureau of Economic Research.

Abdulkadiroğlu, Atila, and Aram Grigoryan. 2023. "Efficient and Envy Minimal Matching." Working Paper.

Abdulkadiroğlu, Atila, and Tayfun Sönmez. 1999. "House Allocation with Existing Tenants." *Journal of Economic Theory*, 88(2): 233–260.

Abdulkadiroğlu, Atila, and Tayfun Sönmez. 2003. "School Choice: A Mechanism Design Approach." *American Economic Review*, 729–747.

Abdulkadiroğlu, Atila, Yeon-Koo Che, Parag A Pathak, Alvin E Roth, and Olivier Tercieux. 2020. "Efficiency, justified envy, and incentives in priority-based matching." *American Economic Review: Insights*, 2(4): 425–442.

Akbarpour, Mohammad, Shengwu Li, and Shayan Oveis Gharan. 2020. "Thickness and information in dynamic matching markets." *Journal of Political Economy*, 128(3): 783–815.

Albright, S Christian. 1974. "Optimal sequential assignments with random arrival times." *Management Science*, 21(1): 60–67.

Alcalde, José, and Salvador Barberà. 1994. "Top dominance and the possibility of strategy-proof stable solutions to matching problems." *Economic theory*, 4(3): 417–435.

Andersson, Tommy, Lars Ehlers, and Alessandro Martinello. 2018. "Dynamic refugee matching." *Mimeo.*

Arnosti, Nick, and Peng Shi. 2020. "Design of lotteries and wait-lists for affordable housing allocation." *Management Science*, 66(6): 2291–2307.

Ashlagi, Itai, Afshin Nikzad, and Philipp Strack. 2023. "Matching in dynamic imbalanced markets." *The Review of Economic Studies*, 90(3): 1084–1124.

Ashlagi, Itai, Faidra Monachou, and Afshin Nikzad. forthcoming. "Optimal allocation via waitlists: Simplicity through information design." *The Review of Economic Studies.*

Baccara, Mariagiovanna, SangMok Lee, and Leeat Yariv. 2020. "Optimal dynamic matching." *Theoretical Economics*, 15(3): 1221–1278.

Balinski, Michel, and Tayfun Sönmez. 1999. "A Tale of Two Mechanisms: Student Placement." *Journal of Economic Theory*, 84(1): 73–94.

Bando, Keisuke, and Ryo Kawasaki. 2021. "Stability and Substitutability in Dynamic Matching Markets." *Available at SSRN 3933533*.

Bergemann, Dirk, and Juuso Välimäki. 2019. "Dynamic mechanism design: An introduction." *Journal of Economic Literature*, 57(2): 235–274.

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

Bloch, Francis, and Nicolas Houy. 2012. "Optimal assignment of durable objects to successive agents." *Economic Theory*, 51: 13–33.

Budish, Eric, and Estelle Cantillon. 2012. "The multi-unit assignment problem: Theory and evidence from course allocation at Harvard." *American Economic Review*, 102(5): 2237–2271.

Che, Yeon-Koo, and Olivier Tercieux. 2023. "Optimal queue design." *Mimeo*.

Combe, Julien. 2023. "Reallocation with priorities and minimal envy mechanisms." *Economic Theory*, 76(2): 551–584.

Combe, Julien, Olivier Tercieux, and Camille Terrier. 2022. "The design of teacher assignment: Theory and evidence." *The Review of Economic Studies*, 89(6): 3154–3222.

Combe, Julien, Umut Mert Dur, Olivier Tercieux, Camille Terrier, and M Utku Ünver. 2022. "Market design for distributional objectives in (re) assignment: An application to improve the distribution of teachers in schools." *Mimeo*.

Damiano, Ettore, and Ricky Lam. 2005. "Stability in Dynamic Matching Markets." *Games and Economic Behavior*, 52(1): 34–53. 00008.

David, Israel, and Uri Yechiali. 1990. "Sequential assignment match processes with arrivals of candidates and offers." *Probability in the Engineering and Informational Sciences*, 4(4): 413–430.

Delacrétaz, David, Scott Duke Kominers, and Alexander Teytelboym. 2023. "Matching mechanisms for refugee resettlement." *American Economic Review*, 113(10): 2689–2717.

Department for Communities and Local Government. 2011. "Choice-based lettings, potentially disadvantaged groups and accessible housing registers: A positive practice guide." Queen's Printer and Controller of Her Majesty's Stationery Office Guidance Report, Eland House, Bressenden Place, London, SW1E 5DU. Accessed: insert date here.

Derman, Cyrus, Gerald J Lieberman, and Sheldon M Ross. 1972. "A sequential stochastic assignment problem." *Management Science*, 18(7): 349–355.

Devine, Deborah J., Lester Rubin, and Robert W. Gray. 1999. *The Uses of Discretionary Authority in the Public Housing Program: A Baseline Inventory of Issues, Policy, and Practice.* US Department of Housing and Urban Development, Office of Policy Development and Research.

Doğan, Battal, and Lars Ehlers. 2021. "Minimally unstable Pareto improvements over deferred acceptance." *Theoretical Economics*, 16(4): 1249–1279.

Doğan, Battal, and Lars Ehlers. 2022. "Robust minimal instability of the top trading cycles mechanism." *American Economic Journal: Microeconomics*, 14(4): 556–582.

Doval, Laura. 2022. "Dynamically stable matching." *Theoretical Economics*, 17(2): 687–724.

Drake, Marshall, Fernando Payró, Neil Thakral, and Linh T Tô. 2024. "Bayesian adaptive choice experiments." *Mimeo*.

Drake, Marshall, Neil Thakral, and Linh T Tô. 2022. "Wage differentials and the price of workplace flexibility." *Mimeo*.

Du, Songzi, and Yair Livne. 2014. "Rigidity of transfers and unraveling in matching markets." *Available at SSRN 2203910*.

Ergin, Haluk I. 2002. "Efficient Resource Allocation on the Basis of Priorities." *Econometrica*, 70(6): 2489–2497.

Ergin, Haluk, Tayfun Sönmez, and M Utku Ünver. 2020. "Efficient and incentive-compatible liver exchange." *Econometrica*, 88(3): 965–1005.

Farajzadeh, Fatemeh, Ryan Baylor Killea, Alexander Teytelboym, and Andrew Christopher Trapp. 2023. "Optimizing Sponsored Humanitarian Parole." 1–13.

Galiani, Sebastian, Alvin Murphy, and Juan Pantano. 2015. "Estimating Neighborhood Choice Models: Lessons from a Housing Assistance Experiment." *American Economic Review*, 105(11): 3385–3415.

Gershkov, Alex, and Benny Moldovanu. 2009. "Learning about the Future and Dynamic Efficiency." *American Economic Review*, 99(4): 1576–1587.

Gershkov, Alex, and Benny Moldovanu. 2010. "Efficient sequential assignment with incomplete information." *Games and Economic Behavior*, 68(1): 144–154.

Gershkov, Alex, and Benny Moldovanu. 2012. "Optimal search, learning and implementation." *Journal of Economic Theory*, 147(3): 881–909.

Gershkov, Alex, Benny Moldovanu, and Philipp Strack. 2015. "Efficient dynamic allocation with strategic arrivals." *Available at SSRN 2548740*.

Gibbard, Allan. 1973. "Manipulation of voting schemes: a general result." *Econometrica: journal of the Econometric Society*, 587–601.

HUD. 2003. "Public Housing Occupancy Guidebook." *https://www.hud.gov/sites/documents/DOC_10760.PDF*, U.S. Department of Housing and Urban Development.

Kadam, Sangram V, and Maciej H Kotowski. 2018*a*. "Multiperiod matching." *International Economic Review*, 59(4): 1927–1947.

Kadam, Sangram V, and Maciej H Kotowski. 2018*b*. "Time horizons, lattice structures, and welfare in multi-period matching markets." *Games and Economic Behavior*, 112: 1–20.

Kaplan, Edward H. 1984. "Managing the Demand for Public Housing." PhD diss. Massachusetts Institute of Technology.

Kaplan, Edward H. 1987. "Analyzing Tenant Assignment Policies." *Management Science*, 33(3): 395–408.

Karp, Richard M, Umesh V Vazirani, and Vijay V Vazirani. 1990. "An optimal algorithm for on-line bipartite matching." 352–358.

Kennes, John, Daniel Monte, and Norovsambuu Tumennasan. 2014. "The day care assignment: A dynamic matching problem." *American Economic Journal: Microeconomics*, 6(4): 362–406.

Kesten, Onur, and M Utku Ünver. 2015. "A theory of school-choice lotteries." *Theoretical Economics*, 10(2): 543–595.

Kotowski, Maciej H. 2020. "A perfectly robust approach to multiperiod matching problems." *HKS Working Paper No. RWP19-016*.

Kurino, Morimitsu. 2014. "House allocation with overlapping generations." *American Economic Journal: Microeconomics*, 6(1): 258–289.

Kurino, Morimitsu. 2020. "Credibility, efficiency, and stability: A theory of dynamic matching markets." *The Japanese Economic Review*, 71: 135–165.

Kwon, Hyukjun, and Ran I Shorrer. 2020. "Justified-envy-minimal efficient mechanisms for priority-based matching." *Available at SSRN 3495266*.

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

Li, Shanjun. 2018. "Better lucky than rich? Welfare analysis of automobile licence allocations in Beijing and Shanghai." *The Review of Economic Studies*, 85(4): 2389–2428.

Li, Shengwu. 2017. "Obviously strategy-proof mechanisms." *American Economic Review*, 107(11): 3257–3287.

Liu, Ce. 2023. "Stability in repeated matching markets." *Theoretical Economics*, 18(4): 1711–1757.

Liu, Tracy, Zhixi Wan, and Chenyu Yang. 2019. "The efficiency of a dynamic decentralized two-sided matching market." *Available at SSRN 3339394*.

Loertscher, Simon, Ellen V Muir, and Peter G Taylor. 2022. "Optimal market thickness." *Journal of Economic Theory*, 200: 105383.

Lubell, Jeffrey M, Mark Shroder, and Barry Steffen. 2003. "Work participation and length of stay in HUD-assisted housing." *Cityscape: A Journal of Policy Development and Research*, 6(2): 207–223.

Lui, Hon-Kwong, and Wing Suen. 2011. "The Effects of Public Housing on Internal Mobility in Hong Kong." *Journal of Housing Economics*, 20(1): 15–29.

McClure, Kirk. 2018. "Length of stay in assisted housing." *Cityscape*, 20(1): 11–38.

Moulin, Hervé. 1980. "On strategy-proofness and single peakedness." *Public Choice*, 35(4): 437–455.

Naik, Canishk, and Neil Thakral. 2023. "A welfare analysis of public housing allocation mechanisms." *Mimeo*.

Nikzad, Afshin, and Philipp Strack. 2023. "Equity and Efficiency in Dynamic Matching: Extreme Waitlist Policies." *Management Science*.

Pereyra, Juan Sebastián. 2013. "A dynamic school choice model." *Games and economic behavior*, 80: 100–114.

Productivity Commission. 2017. "Introducing Competition and Informed User Choice into Human Services: Reforms to Human Services." Productivity Commission Report 85, Canberra. Report No. 85.

Pycia, Marek, and Peter Troyan. 2023. "A theory of simplicity in games and mechanism design." *Econometrica*, 91(4): 1495–1526.

Romeo-Beehler, Beverly. 2019. "Opening Doors to Stable Housing: An Effective Waiting List and Reduced Vacancy Rates Will Help More People Access Housing." *Toronto Auditor General's Office*.

Romm, Assaf, Alvin E Roth, and Ran I Shorrer. 2021. "Stability vs. no justified envy." *Mimeo*.

Roth, Alvin E. 1982. "The economics of matching: Stability and incentives." *Mathematics of operations research*, 7(4): 617–628.

Roth, Alvin E, and Elliott Peranson. 1999. "The redesign of the matching market for American physicians: Some engineering aspects of economic design." *American economic review*, 89(4): 748–780.

Roth, Alvin E, Tayfun Sönmez, and M Utku Ünver. 2004. "Kidney exchange." *The Quarterly journal of economics*, 119(2): 457–488.

Roth, Alvin E., Uriel G. Rothblum, and John H. Vande Vate. 1993. "Stable Matchings, Optimal Assignments, and Linear Programming." *Mathematics of Operations Research*, 18(4): 803–828.

Rothblum, Uriel G. 1992. "Characterization of stable matchings as extreme points of a polytope." *Mathematical Programming*, 54: 57–67.

Satterthwaite, Mark Allen. 1975. "Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions." *Journal of economic theory*, 10(2): 187–217.

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

Schummer, James, and Rakesh V Vohra. 2013. "Assignment of Arrival Slots." *American Economic Journal: Microeconomics*, 5(2).

Shi, Peng, and Junxiong Yin. 2022. "Eliminating Waste in Cadaveric Organ

Allocation." *USC Marshall School of Business Research Paper Sponsored by iORB.*

Sönmez, Tayfun. 2023. "Minimalist market design: A framework for economists with policy aspirations." *arXiv preprint arXiv:2401.00307.*

Sönmez, Tayfun, and M Utku Ünver. 2010*a*. "Course bidding at business schools." *International Economic Review*, 51(1): 99–123.

Sönmez, Tayfun, and M Utku Ünver. 2010*b*. "House allocation with existing tenants: A characterization." *Games and Economic Behavior*, 69(2): 425–445.

Su, Xuanming, and Stefanos A Zenios. 2005. "Patient choice in kidney allocation: A sequential stochastic assignment model." *Operations research*, 53(3): 443–455.

Su, Xuanming, and Stefanos A Zenios. 2006. "Recipient choice can address the efficiency-equity trade-off in kidney transplantation: A mechanism design model." *Management science*, 52(11): 1647–1660.

Su, Xuanming, and Stefanos Zenios. 2004. "Patient choice in kidney allocation: The role of the queueing discipline." *Manufacturing & Service Operations Management*, 6(4): 280–301.

Thakral, Neil. 2016. "The Public-Housing Allocation Problem: Theory and evidence from Pittsburgh." *Mimeo.*

Thakral, Neil. 2017. "The Design of Public-Housing Policies." *Mimeo.*

Thompson, Steve. 2022. "1 in 4 public housing units sit vacant during D.C. affordability crisis."
Ünver

Ünver, M. Utku. 2010. "Dynamic Kidney Exchange." *The Review of Economic Studies*, 77(1): 372–414.

Vande Vate, John H. 1989. "Linear programming brings marital bliss." *Operations Research Letters*, 8(3): 147–153.

Verdier, Valentin, and Carson Reeling. 2022. "Welfare effects of dynamic matching: An empirical analysis." *The Review of Economic Studies*, 89(2): 1008–1037.

Waldinger, Daniel. 2021. "Targeting in-kind transfers through market design: A revealed preference analysis of public housing allocation." *American Economic Review*, 111(8): 2660–2696.

Wallack, Todd, and Christine Willmsen. 2023. "Mass. has a huge waitlist for state-funded housing. So why are 2,300 units vacant?"

WBUR Newsroom. 2024. "WBUR and ProPublica investigation of public housing vacancies in Mass." *https://www.wbur.org/news/2023/09/19/massachusetts-state-funded-public-housing-waitlist-vacant.*

Figure 1: Average welfare and characteristics across cities

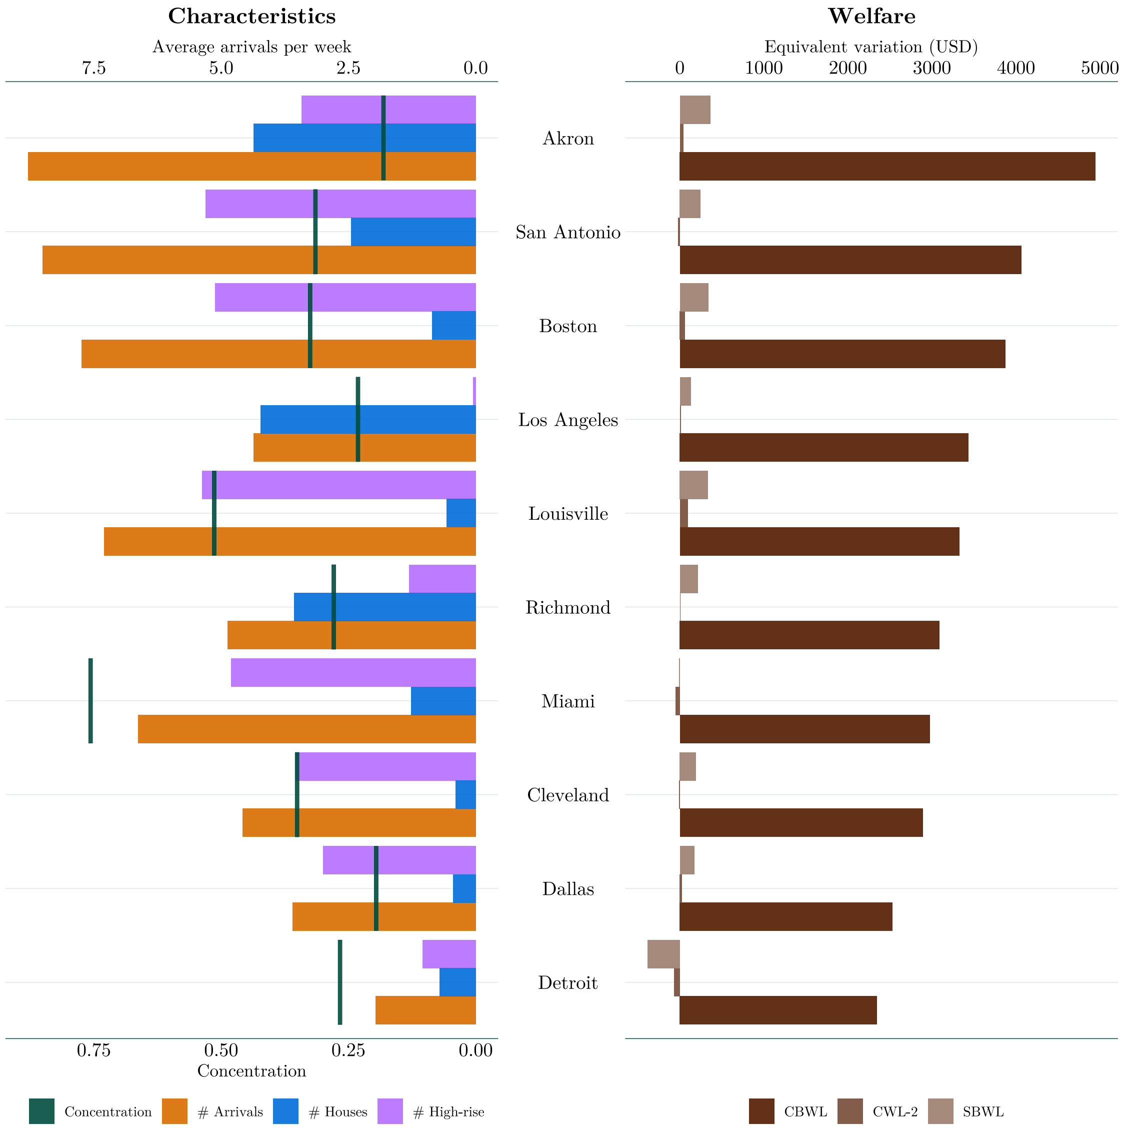

Note: The right panel displays the average equivalent variation (EV), or monetary transfer that an applicant requires under CWL-1 to reach the same level of lifetime utility as their assignment under each alternative mechanism, for each city across 50 simulations involving 587 applicants. The left panel presents characteristics of the building arrival process in each city. The vertical lines, corresponding to the scale on the bottom, denote the geographic concentration of the public housing stock in each city, measured as the sum of squared shares of housing units in each coarse geographic area of the city (Herfindahl-Hirschman index). The bars, corresponding to the scale on the top, denote to the average number of arrivals each week (bottom bar for each city), the average number of units arriving in townhouse communities each week (middle bar for each city), and the average number of units arriving in high-rise buildings (top bar for each city).

Figure 2: Average welfare gains and losses

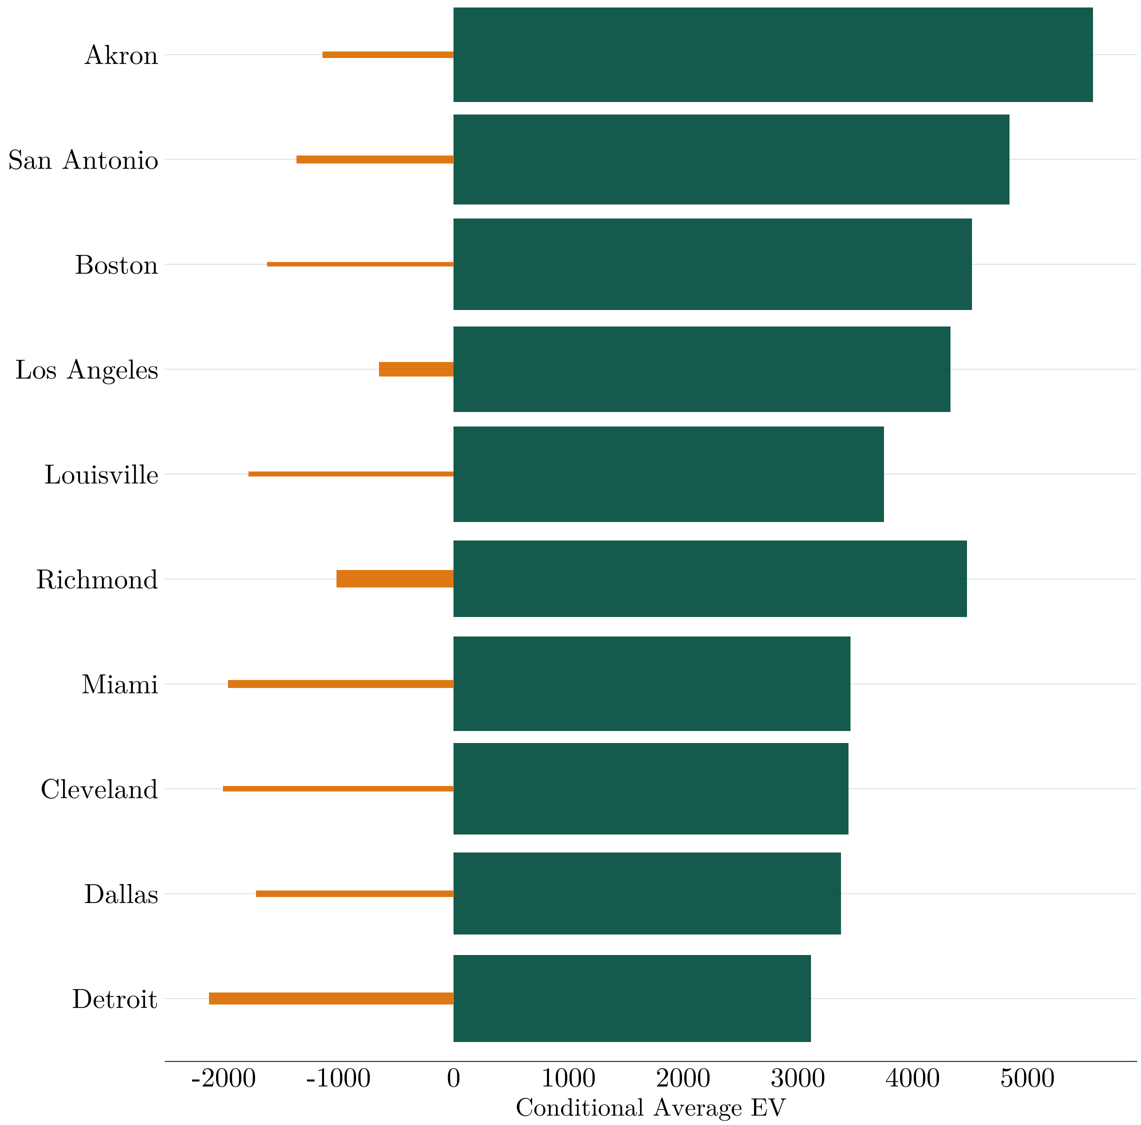

Note: The figure depicts, for each city, average welfare gains and losses, as well as the share of applicants gaining or losing, when changing the allocation mechanism from CWL-1 to CBWL. The horizontal length of each bar extending to the right of the origin corresponds to the average welfare gain (mean equivalent variation conditional on positive values) across 50 simulations, while their vertical widths represent the average share of applicants experiencing gains. The length of each bar extending to the left corresponds to the average welfare loss across 50 simulations, while their vertical widths represent the average share of applicants experiencing losses.

Figure 3: Relationship between outside options and welfare across mechanisms

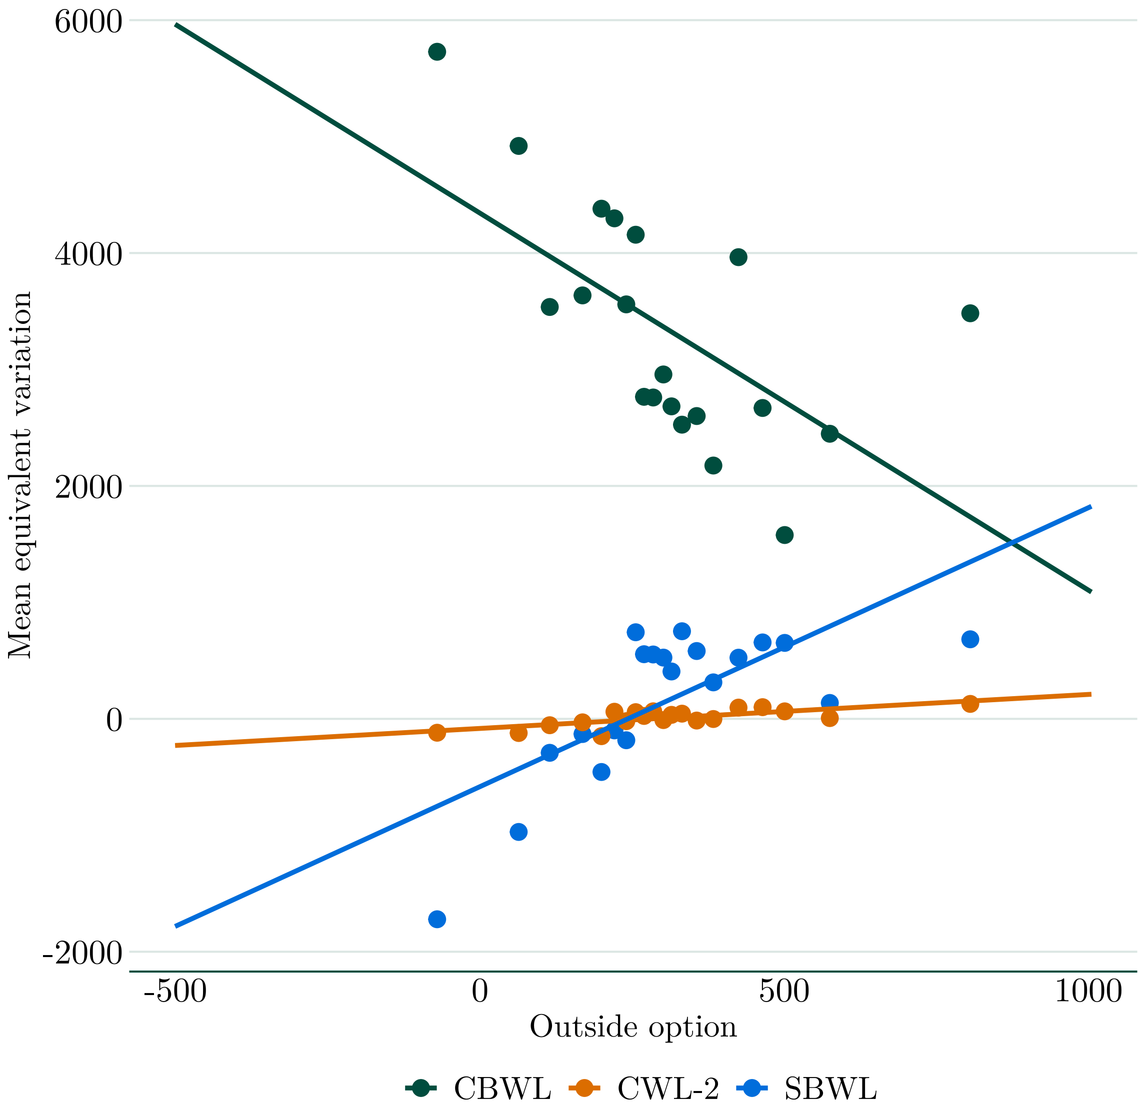

Note: The figure shows the relationship between applicants' outside options and their average welfare under each allocation mechanism across 50 simulations. The welfare measure (vertical axis) corresponds to the equivalent variation, or monetary transfer that an applicant requires under CWL-1 to reach the same level of lifetime utility as their assignment under an alternative mechanism. The outside option measure (horizontal axis) corresponds to each applicant's flow utility from residing in private housing.