

The Public-Housing Allocation Problem: Theory and Evidence from Pittsburgh

Neil Thakral*

August 2016

Abstract

We present a model of public-housing allocation, in which objects that arrive stochastically over time must be matched with applicants on a waiting list. Our framework delivers a strategy-proof, ex-ante efficient and envy-free mechanism which allows applicants to trade off their preferences for different units and waiting times by choosing among a set of waiting lists. A counterfactual change from existing mechanisms, in which applicants receive offers after units become available and can refuse only a limited number of times, to the proposed mechanism improves welfare by a lower bound of \$6,429 per applicant in a sample of Pittsburgh households.

*Department of Economics, Harvard University, Littauer Center, 1805 Cambridge Street, Cambridge, MA 02138 (e-mail: nthakral@fas.harvard.edu). A previous version of this paper circulated under the title “Matching with Stochastic Arrival: Theory and an Application to Public Housing.” I am grateful to Kfir Eliaz for his constant support and encouragement, as well as Scott Duke Kominers for suggestions and advice. I would also like to thank David M. Cutler, Itay P. Fainmesser, Peter Ganong, Edward L. Glaeser, Benjamin Golub, Jerry R. Green, Kevin He, Sonia Jaffe, Lawrence F. Katz, Maciej H. Kotowski, David Laibson, Robin S. Lee, Xiaosheng Mu, Parag A. Pathak, Andrei Shleifer, Ran I. Shorrer, Tayfun Sönmez, and especially Linh T. Tô for helpful comments and discussions.

1 Introduction

The efficacy of the public sector depends not only on the supply of public services but also on the design of systems for provision. While the former receives considerable attention from policymakers, poor design can entail substantial welfare losses. The market for public housing exemplifies this phenomenon. Tenants of public housing throughout the world—from approximately 1.2 million households in the United States to about one-third of all households in Hong Kong and Romania—express little choice in their place of residence, resulting in suboptimal allocations [Lui and Suen, 2011, Schwartz, 2010, Soaita, 2012].

Consistent with the idea of spatial mismatch, public-housing residents are less mobile than their private-housing counterparts, are less likely to work near where they live, and are more likely to move farther away from their original place of residence conditional on moving [Lui and Suen, 2011]. To the extent that housing is misallocated, designing improved allocation mechanisms can lead to substantial welfare gains. In the private-housing market, Glaeser and Luttmer [2003] argue that rent controls contribute to inefficient allocations, but a corresponding analysis of the market for public housing would require a model of how tenants are matched with apartments.

This paper models public-housing allocation as a matching problem in which units arrive stochastically over time and are assigned to applicants on a waiting list. Under this framework, we investigate the design of strategy-proof allocation mechanisms (i.e., those that are not subject to strategic manipulation) and explore whether these mechanisms can satisfy additional efficiency and fairness properties. Although there does not exist a mechanism that achieves these properties ex-post, we introduce a new mechanism—the Multiple-Waitlist Procedure (MWP)—which satisfies these properties ex-ante. Under MWP, applicants begin on a centralized waiting list in order of priority; a unit that arrives is offered to the agent at the top of the centralized waiting list, who can accept the offer or commit to a site-specific waiting list for another unit of their choice.

Mechanisms generally used in practice fail to satisfy desirable properties. Most

public-housing agencies employ a take-it-or-leave-it procedure to allocate units to households on a priority-ordered waiting list.¹ A simple example with two buildings, two applicants, and two periods illustrates the possibility of unfair and inefficient allocations. Assume that household 1 has higher priority and that in period $k \in \{1, 2\}$ a unit in building b_k becomes available with certainty. Further assume that household 1 strictly prefers to wait for b_2 and that household 2 prefers b_1 to b_2 . The take-it-or-leave-it mechanism assigns b_1 to household 1 and b_2 to household 2. The allocation is “unfair” in the sense that the higher-priority applicant prefers the assignment of the lower-priority applicant (i.e., the mechanism fails to eliminate justified envy); additionally, since household 2 also prefers the allocation of household 1, the mechanism is inefficient.² MWP satisfies both of these properties by giving applicants the opportunity to decline an offer and join a First-In/First-Out (FIFO) waiting list for the building of their choice.

To evaluate the performance of MWP in practice, we use a structural model of household preferences for public housing due to [Geyer and Sieg \[2013\]](#) and find that a counterfactual change in the allocation mechanism achieves substantial welfare gains. We estimate an arrival process for housing units using aggregate data from the Housing Authority of the City of Pennsylvania, and we obtain a sample of households eligible for public housing in Pittsburgh with preferences given by the Geyer-Sieg estimates using data from the Survey of Income and Program Participation. We simulate allocations under MWP as well as various alternative mechanisms that are currently being used to allocate public housing, some of which have support in the literature, to compare ex-post welfare. As a lower bound, we find that changing the allocation mechanism in Pittsburgh to MWP improves welfare by an average of \$6,429 per household. Thus, MWP performs well ex-post, attaining 75 percent of an idealized benchmark defined as the maximum possible welfare gain that could be achieved if the realization of the arrival process were known in advance. Moreover,

¹Table 1 in Section 4.1 provides further details about public-housing agencies in the US and the procedures for allocating housing.

²These issues persist under a common variant of the take-it-or-leave-it mechanism that permits a household to decline some number of offers without forfeiting its position in the waiting list, as Section 4.1 explores.

MWP increases the benefit of public housing by almost 20 percent without affecting its cost.

While several studies find evidence of the misallocation of private housing [Glaeser and Luttmer, 2003, Wang, 2011], empirical work related to public-housing allocation is more limited.³ Recent work by Van Ommeren and Van der Vlist [2016] estimates the marginal willingness to pay for public-housing characteristics in Amsterdam and finds little welfare loss due to distortions in housing supply; they argue that deadweight loss mainly results from the inefficient match between households and apartments. Our paper complements these approaches by exploring the design aspect of public-housing allocation.⁴ Geyer and Sieg [2013] develop an equilibrium framework for estimating household preferences for public housing under supply-side restrictions, which our paper uses to analyze welfare.

Our framework for studying the design of public-housing allocation mechanisms consists of a model of matching in a dynamic setting. Standard models from matching theory, such as the canonical “house allocation problem” due to Hylland and Zeckhauser [1979], typically describe static situations; problems such as public-housing allocation, however, involve the feature that units arrive stochastically over time.⁵ We develop a model that is based on existing static frameworks yet is rich enough to shed light on the operation of basic concepts from market design such as stability and efficiency in an environment with uncertainty. Instead of analyzing a fully dynamic problem, which would result in technical complications that are not relevant for the motivating application, we introduce a simple framework that maintains the key aspects of public-housing allocation. Specifically, our model incorporates the following features: (i) units are allocated dynamically as they arrive over time, (ii) there is uncertainty about the availability of the units, and (iii) applicants have preferences over waiting times. The robust policy recommendation that emerges from

³As Geyer and Sieg [2013] note, this is largely because “public housing agencies are not willing to disclose detailed micro-level data on wait lists.”

⁴Recent work by Galiani et al. [2015] also studies the design of housing-assistance programs by using estimates from a structural model of neighborhood choice to simulate the effects of counterfactual housing-voucher policies.

⁵For a survey treatment of the house allocation problem, see Sönmez et al. [2011].

this framework involves households choosing among a set of waiting lists rather than only choosing to accept or reject units after they arrive. [Leshno \[2015\]](#) also studies an allocation problem with stochastic arrival and looks at a utilitarian objective in a simplified setting, while the present paper takes an axiomatic approach to a more general setting. Related work on dynamic assignment problems by [Schummer \[2016\]](#) and [Bloch and Cantala \[2017\]](#) analyzes selectivity thresholds for agents to accept or reject objects. While these papers also mention the allocation of public housing as a motivating example, they abstract away from the problem of eliciting waiting costs. The mechanism-design problem in the present paper, by contrast, permits heterogeneity in preferences over object types as well as over the amount of time spent waiting for an allocation. Furthermore, our results apply for any underlying stochastic process that governs the arrival of units.⁶

A growing literature in market design uses simulations for welfare analysis, though much of this work focuses on the school-choice problem and involves randomly generated preference data.⁷ Our work is among the first to use estimated preferences from data on real-world assignment procedures to quantify welfare gains due to adopting alternative mechanisms.⁸ The counterfactual simulations are based on a model of preferences for public housing by [Geyer and Sieg \[2013\]](#), and we discuss in detail how their assumptions and our implementation all lead to an underestimate of the welfare gains from changing existing public-housing allocation mechanisms to MWP.

The paper is organized as follows. [Section 2](#) describes the dynamic allocation

⁶Other papers in which dynamic considerations arise due to arrival over time include [Gershkov and Moldovanu \[2009\]](#), which investigates an allocation setting in which agents arrive sequentially; [Ünver \[2010\]](#), in which each agent arrives with an object to trade; [Baccara et al. \[2015\]](#), which studies a two-sided matching market with sequential arrivals; [Gershkov et al. \[2015\]](#), in which agents choose when to make themselves available for trade; [Akbarpour et al. \[2016\]](#), in which agents arrive and depart stochastically in a networked market; and [Doval \[2016\]](#), which explores dynamically stable matchings in two-sided markets.

⁷See, for example, [Erdil and Ergin \[2008\]](#), [Dur \[2011\]](#), [Hafalir et al. \[2013\]](#), [Morrill \[2013\]](#), [Abdulkadiroğlu et al. \[2015b\]](#), and [Kesten and Ünver \[2015\]](#).

⁸In the literature on school choice, some recent papers estimate preferences using data from school districts: [He \[2012\]](#) in Beijing, China; [Agarwal and Somaini \[2014\]](#) in Cambridge, MA; [Calsamiglia et al. \[2016\]](#) in Barcelona, Spain; and [Abdulkadiroğlu et al. \[2015a\]](#) in New York City, NY.

problem and introduces various properties. Section 3 characterizes the Multiple-Waitlist Procedure and discusses the main theoretical results and extensions. Section 4 applies the framework to the allocation of public housing and uses estimates from a structural model to evaluate welfare. Section 5 concludes. Proofs can be found in the appendix.

2 A Model of Public-Housing Allocation

Consider the problem that a public-housing agency faces in assigning units from apartment buildings that vary based on location to applicants on a long waiting list. Applicants have heterogeneous preferences over apartment buildings, as they may prefer to live closer to their respective workplaces, and heterogeneous waiting costs.⁹ The public-housing agency ranks applicants based on priority; housing authorities typically rank applicants based on a coarse point system with ties broken based on waiting time, resulting in a strict priority ordering [Greely, 1977]. A unit must be allocated in the period in which it arrives. Although the period in which a given unit becomes available is not known in advance, the distribution of waiting times is known [Kaplan, 1986]. With this example in mind, we now proceed to introduce the components of the model more formally.

A *public-housing allocation problem* is a five-tuple $\langle A, B, \succ_B, \succ_A, \pi \rangle$, where A is a set of applicants, B is a finite set of buildings; $\succ_B = (\succ_b)_{b \in B}$ is a profile of the buildings' strict priority relations over the set of applicants; $\succ_A = (\succ_a)_{a \in A}$ is a profile of applicants' preference relations over building-time pairs $B \times \mathbb{R}_+$; and $\pi = \left(\pi_{b, \hat{t}}(\cdot | h^t) \right)_{b \in B, t, \hat{t} \in \mathbb{N}}$ is an arrival process that specifies, conditional on the history h^t , a probability distribution over the number of units in building b that arrive at time \hat{t} . A *history* is a map $h^t: B \times \{\tilde{t} \in \mathbb{N} : \tilde{t} \leq t\} \rightarrow \mathbb{N}$ that specifies the number of units in each building that have arrived in the previous periods. We make no specific assumptions about the underlying stochastic process which governs the arrival of

⁹As a starting point, assume that units within each apartment building are identical. In practice, apartment buildings contain units of different sizes (i.e., as measured by the number of bedrooms); however, public-housing agencies administer separate waiting lists for units of different sizes.

units but will find it convenient to denote by $\tau_{b,t}(r)$ the expected waiting time of the r^{th} unit in building b to become available (counted from the beginning of time), conditional on the history h^t .

To denote applicant a being matched with the r^{th} unit to arrive in building b , we write $\mu(a) = \langle b, r \rangle$, where we refer to the map $\mu^t: A \rightarrow (B \times \mathbb{N}) \cup \{\emptyset\}$ as an *assignment in period t* and we refer to the collection $\mu = (\mu^1, \mu^2, \dots)$ of assignments in each period as an *assignment*. Agent a can be thought of as being on a waiting list for building b if the r^{th} unit is not yet available: given the history h^t of the arrival process, the unit is expected to arrive in $\tau_{b,t}(r)$ periods. If $\tau_{b,t}(r) = 0$, then this waiting list is degenerate, so the applicant receives the unit immediately in period t . We use the word *allocation* to refer to an agent's realized assignment, i.e., the applicant receives an allocation when the assigned unit arrives or becomes available.

A unit cannot be unmatched after the period in which it becomes available and cannot be reallocated in the future. The assumption that units must be allocated irreversibly upon arrival mirrors the corresponding assumptions in [Gershkov and Moldovanu \[2009\]](#) for the case of agents arriving sequentially. More importantly, this assumption captures a realistic institutional feature of public-housing allocation [[Navarro, 2015](#)].

To simplify the exposition, we assume that preferences are dynamically consistent and that waiting is costly. Preferences satisfy the *dynamic-consistency condition* if $(b, t) \succ_a (b', t')$ implies $(b, t + \hat{t}) \succ_a (b', t' + \hat{t})$ for every $\hat{t} > 0$. Preferences satisfy the *costly-waiting condition* if $t < t'$ implies $(b, t) \succ_a (b, t')$. These assumptions yield for each applicant a constant per-period waiting cost. Equivalently, the preference relation over $B \times \mathbb{R}_+$ reduces to a ranking over buildings and a vector in $\mathbb{R}_+^{|B|}$: for each building b , this representation encodes the greatest number of periods that the applicant is willing to wait before receiving a unit in her most-preferred building rather than receiving a unit in building b immediately. In addition, we assume that applicants are *risk neutral* with respect to preferences over waiting times. After stating the main result, [Section 3.1](#) discusses how these assumptions can be relaxed.

Recall that preferences are defined over building-time pairs: $(b, t) \succ_a (b', t')$ if and only if applicant a prefers to receive a unit in building b in period t over a unit in

building b' in period t' . However, an assignment $\mu^t(a)$ consists of a building $b \in B$ and a unit indexed by $r \in \mathbb{N}$ (i.e., the r^{th} unit that becomes available in building b). An agent a who is assigned $\mu^t(a)$ in period t expects to wait $\tau_{b,t}(r)$ periods for the unit to become available. The assumption of risk neutrality with respect to waiting times implies that applicants evaluate assignments based on expected waiting times: in particular, $\mu^t(a) = \langle b, r \rangle$ is preferred to $(\mu')^{t'}(a) = \langle b', r' \rangle$ if and only if $(b, t + \tau_{b,t}(r)) \succ_a (b', t' + \tau_{b',t'}(r'))$.

An *allocation mechanism* φ is a procedure that uses reported preferences, the exogenous priority orderings, and the history to choose an assignment μ^t in each period t . Let θ'_a denote the reported preferences of applicant $a \in A$, and let θ'_{-a} be the profile of reported preferences of all applicants except a . An allocation mechanism induces a preference-revelation game in which the set of players is A , the strategy space for player a is the set of preferences Θ , and each player $a \in A$ has true preference $\theta_a \in \Theta$.

We say that a mechanism φ is *strategy-proof* if deviation from truthful preference revelation is not profitable along any possible arrival history. Various authors have emphasized the desirability of strategy-proofness because of fairness (agents who lack information or sophistication are not at a disadvantage) and robustness (the equilibrium does not depend on beliefs about other agents' information or preferences). Another justification for strategy-proofness in our dynamic setting is that the social planner may make costly investments (though not modeled here) based on reported preferences.¹⁰

Next we define a property that can be interpreted as a form of fairness. An assignment μ *eliminates justified envy* or is *free of justified envy* if an applicant who prefers an alternate assignment does not have higher priority than the applicant to whom the other unit is assigned. Legal scholars posit this principle in the context of public-housing allocation, noting that “selection... would be by priority of application... to meet certain ‘entitlements’ arising out of a sense of fairness”

¹⁰A housing agency may use reported preferences to determine where to construct a new building. In a separate context, [Abdulkadiroğlu et al. \[2009\]](#) point out that lack of demand as determined by reported student preferences under a strategy-proof mechanism contributed to the closing of an unpopular New York City high school in 2006.

[Greely, 1977]. A literature in market design that studies allocation problems with priorities, starting with the student placement model of Balinski and Sönmez [1999], employs this property as an analogue of stability. In our dynamic setting, whether an applicant prefers an alternate assignment depends on the timing of the match and the information available at the time. We say that an applicant a *envies* another applicant a' if given the information available at the time when a was matched a would have preferred the assignment of a' ; we say that this envy is *justified* if a has higher priority than a' for the building with which a' is matched.

Definition 1. An assignment μ *eliminates justified envy* if whenever a is assigned $\langle b, r \rangle$ in period t and a' is assigned $\langle b', r' \rangle$ in period t' , we have

$$(b', t + \tau_{b',t}(r')) \succ_a (b, t + \tau_{b,t}(r)) \implies a' \succ_{b'} a.$$

An ex-post variant of this no-envy condition would state that an applicant who prefers another unit to her own, given the realized arrival times of their respective units, cannot have higher priority than the applicant to whom the other unit is assigned.

An assignment μ is *efficient* if any reassignment μ' that one agent strictly prefers would make another agent strictly worse off (where, as before, we consider the information available at the time of the match). We refer to this as an ex-ante notion of efficiency because agents only take into account information that is available at the time when they are matched and are evaluating their expected (rather than realized) arrival times.

Definition 2. An assignment μ is *efficient* if for any feasible assignment $\mu' \neq \mu$ for which there exists some a (who is assigned $\langle b_a, r_a \rangle$ in period t_a under μ and is assigned $\langle b'_a, r'_a \rangle$ in period t'_a under μ') such that $(b'_a, t_a + \tau_{b'_a, t_a}(r'_a)) \succ_a (b_a, t_a + \tau_{b_a, t_a}(r_a))$, there exists some $a' \in A$ such that

$$(b_{a'}, t_{a'} + \tau_{b_{a'}, t_{a'}}(r_{a'})) \succ_{a'} (b'_{a'}, t_{a'} + \tau_{b'_{a'}, t_{a'}}(r'_{a'}))$$

An alternate notion of efficiency would be ex-post efficiency: any way of re-

distributing units that have already arrived (among the agents to whom they are assigned) cannot strictly improve the allocation of one applicant without making another strictly worse off.

We say that a mechanism satisfies a given property if the equilibrium allocations in the induced preference-revelation game satisfy that property. The following result shows that no allocation mechanism satisfies the ex-post variants of the properties introduced above.

Proposition 1. *There does not exist an allocation mechanism that is ex-post efficient or ex-post free of justified envy.*

The proof in Appendix A.1 consists of a simple example to demonstrate the impossibility of designing an allocation mechanism that guarantees these ex-post properties in an environment with stochastic arrival because the realization of the arrival process may be such that neither can possibly hold. Motivated by this result, our theoretical analysis in the next section focuses on the ex-ante properties defined above. Nevertheless, the empirical analysis in Section 4 uses ex-post welfare as a criterion for comparing mechanisms.

3 Multiple-Waitlist Procedure

3.1 Common priorities

We begin by considering the case that the buildings share a common priority ordering, i.e., that there exists \succ_* such that $\succ_* = \succ_b$ for all $b \in B$. We introduce the Multiple-Waitlist Procedure (MWP) and characterize the matching that results from this mechanism.

Under MWP, all applicants begin on a centralized waiting list, and associated with each building there is a separate *First-In/First-Out* (FIFO) queue. A unit that becomes available in a given building belongs to the applicant at the top of the queue for that building. If the queue is empty, then the unit is offered to the applicant with the highest priority on the centralized waiting list. Given information about the

distribution of arrival times, the applicant can either accept the offer or opt to join the FIFO queue for the next available unit in a different building of the applicant’s choice. Note that if a' has higher priority than a , then a' receives an assignment (i.e., a unit or a place on some waiting list) before a does. Figure 1 describes MWP more formally.

The following proposition characterizes the main properties of this mechanism.

Proposition 2. *MWP satisfies the following properties: (i) strategy-proofness, (ii) efficiency, and (iii) elimination of justified envy.*

The following arguments summarize the formal proof in Appendix A.2. Note that an applicant receives an assignment after reaching the top of the centralized waiting list once a unit becomes available. The period when an applicant is matched depends on the choices of higher-priority applicants (and on the realization of the arrival process) but not on the agent’s own reported preferences (and not on the choices of lower-priority applicants). In particular, since the priority ordering is fixed and independent of the agents’ strategies, no agent can obtain a match sooner by misreporting preferences. Moreover, since the assignment is chosen to maximize the applicant’s reported preference ordering (given the choices of the agents who have already been assigned, but independent of the choices of the agents who have not yet been assigned), the applicant cannot gain by deviating from truth telling. This argument suggests not only that MWP is strategy-proof but also that it satisfies a stronger condition, namely *obvious strategy-proofness*, which implies that the equilibrium prediction extends to agents with certain cognitive limitations [Li, 2015].¹¹

The observation that the match of each applicant maximizes reported preferences (given the information available at the time of the match) also implies that all units

¹¹A mechanism is obviously strategy-proof if it has an equilibrium in obviously dominant strategies; truth telling is obviously dominant if the best-possible outcome from deviating is no better than the worst-possible outcome from reporting truthfully. Li [2015] notes that these notions also apply to the case that the set of outcomes of the mechanism consists of lotteries, so obvious strategy-proofness applies in our setting where an outcome is a unit in a building (which may consist of a distribution of waiting times). Under this view, MWP continues to implement the desired outcome when agents correctly perceive randomization by nature but use simplified mental representations of the other agents’ strategies.

assigned to lower-priority applicants are available but not chosen. Since priority orderings are common across buildings and no higher-priority applicant prefers the assignment of a lower-priority applicant, the allocation is efficient and free of envy. The results continue to hold after relaxing several assumptions in the framework as follows.

Risk neutrality An applicant who has risk-neutral preferences with respect to waiting times evaluates units based on expected arrival times and would have no use for additional moments of the distribution. A social planner can (i) implement MWP as a direct mechanism by eliciting preferences ex-ante, or (ii) disclose the expected arrival times and implement MWP as an indirect mechanism by making offers as in Figure 1. This holds more generally when the applicants have homogeneous risk preferences that the planner knows, as long as the planner can disclose certainty equivalents under the indirect mechanism. A planner with information about the applicants' risk preferences and the full distribution of waiting times can implement the direct mechanism, and a planner that can disclose the distribution of waiting times can implement the indirect mechanism without knowing the applicants' risk preferences.

Costly waiting The mechanism as described in Figure 1 moves an applicant who declines an offer from the top of the centralized waiting list to the end of a FIFO queue, thereby allocating the applicant with the next available unit in the chosen building. However, MWP can accommodate richer time preferences by allowing an applicant who prefers to wait longer (e.g., if moving costs change over time) to choose any unoccupied position in the queue.¹² In other words, the applicant can select any unassigned unit within the chosen building rather than the next available unassigned unit.

¹²In the context of public-housing allocation, for example, a household with an existing lease elsewhere may prefer a later allocation due to early-termination costs.

Static set of applicants The model captures the idea of an overloaded waiting list.¹³ Under MWP, the centralized waiting list consists of the fixed priority-ordered set of applicants taken as a primitive of the model. Without any modification, the mechanism applies under an alternative formulation that involves an arrival process for applicants. The newly arrived applicants can be of any priority level and correspondingly can be added anywhere in the centralized waiting list. The mechanism likewise allows for departures from the centralized waiting list of applicants at any priority level. Insofar as the applicants’ strategies do not influence such waiting-list dynamics, MWP satisfies all of the same properties since the procedure once an agent reaches the top of the centralized waiting list remains the same.

Static priorities We can interpret the exogenous priority ordering as reflecting the social planner’s objective determined by verifiable information about the applicants. These priorities may change over time, for example, due to changes in the applicants’ circumstances such as disability status or homelessness. The mechanism can incorporate any such changes, as the discussion above suggests for the case of applicants arriving or departing from the waiting list, by modifying the order of the applicants on the centralized waiting list. This approach requires the planner to remain committed to allocating units that have not yet arrived to agents on the FIFO waiting lists. An alternative that also preserves the properties of the mechanism would be to allow applicants to become *critical* exogenously and alter the mechanism so that critical types receive an allocation immediately. The social planner can therefore accommodate situations such as domestic abuse or natural disasters by allowing such applicants to bypass the FIFO queues. The mechanism proceeds by adjusting the expected waiting times accordingly to account for the arrival process of critical agents.

Identical units The mechanism instantiates a separate FIFO queue for each building, where a “building” refers to an exogenously specified collection of identical units. If the classification of units into buildings were too fine, then an applicant who

¹³See [Leshno \[2015\]](#) for a related allocation problem with two kinds of objects and an overloaded waiting list.

is indifferent between the units in two separate buildings would choose the building with the lower expected waiting time. Similarly, if the classification were too coarse (e.g., with unobserved heterogeneity in the applicants’ preferences for units within a building), the applicant at the top of the centralized waiting list could still choose a FIFO queue to maximize expected utility, with the expectation taken over the distribution of unit quality in addition to waiting time.

We have shown that MWP satisfies several desirable properties and continues to do so in a more general setting. MWP satisfies additional regularity conditions discussed more formally in Appendix A.3: non-bossiness (no applicant can change another’s allocation without changing her own allocation) and neutrality (the allocation does not depend on the labeling of the units). The next result complements our characterization of MWP by describing the sense in which MWP is unique. The uniqueness result does not rely on any of the efficiency or fairness properties introduced in Section 2.

Recall that a unit must be allocated in the period when it arrives. Under MWP, applicants at the top of the centralized waiting list receive offers (and move to their chosen FIFO queues if they decline) until no units available in the current period remain unallocated. Proposition 2 also applies to extended versions of this procedure that may place additional applicants from the top of the centralized waiting list in FIFO queues, even after allocating all units available in the current period. The following proposition states that under the regularity conditions, this class of extended multiple-waitlist procedures consists of all strategy-proof allocation mechanisms.

Proposition 3. *Any strategy-proof, non-bossy, and neutral allocation mechanism is an extended multiple-waitlist procedure.*

Appendix A.3 contains a formal proof of this result. Under the alternative procedures, the planner makes commitments before receiving information—even though waiting is feasible—that the planner may not have made after receiving the information. At an extreme lies the serial dictatorship: all applicants on the overloaded waiting list sequentially choose units in the initial period based on expected arrival times. Among this class of strategy-proof mechanisms, MWP uses the most

information from the arrival process and requires the fewest promises about future allocations. Accordingly, MWP is most conducive to the more general setting highlighted above, which allows for time-varying priorities, risk preferences, and changes in the set of agents over time.

3.2 Heterogeneous priorities

Although our empirical analysis in Section 4 focuses on a case with common priority orderings, applications of our framework may not satisfy this restriction. Housing authorities often assign higher priority to elderly, disabled, or homeless applicants, as well as victims of natural disasters or domestic abuse. In some cases, however, priorities may differ across apartment buildings. For example, certain communities may be designated for specific groups such as seniors. For the general case that priority orderings are not common across object types, we provide a necessary and sufficient condition for the existence of a strategy-proof allocation mechanism that satisfies the ex-ante efficiency and fairness properties. The existence of such a mechanism depends on a property that can be interpreted as a measure of similarity between the priority orderings. Suppose a_1 has strong priority over a_3 at building b_1 in the sense that there is some a_2 such that $a_1 \succ_{b_1} a_2 \succ_{b_1} a_3$. We say that the priority orderings are *acyclic* if this implies a_1 has priority over a_3 at all buildings.

By defining the priority ordering over the set of applicants A , the framework implicitly assumes complete priority lists, i.e., that all applicants are acceptable to all buildings. The assumption is without loss of generality for the case of common priority orderings, as an unacceptable applicant can effectively be removed from the set A . With heterogeneous priority orderings, however, there may be applicants who are only acceptable to some buildings. In this case, an acyclic priority ordering satisfies the following condition: if a_1 has strong priority over a_3 at some building where a_3 is acceptable, then a_1 has priority over a_3 at any building where a_1 is acceptable. This property, adapted from Ergin [2002], is formalized as follows.¹⁴

¹⁴Key results in Kojima [2011], Kesten [2006], Dur [2012], Romero-Medina and Triossi [2013], and Doval [2016] also depend on variants of this acyclicity condition.

Definition 3. The collection of priority orderings \succ_B contains a *cycle* if there exist $a_1, a_2, a_3 \in A$ and $b_1, b_2 \in B$ such that $a_1 \succ_{b_1} a_2 \succ_{b_1} a_3 \succ_{b_1} \emptyset$ and $a_3 \succ_{b_2} a_1 \succ_{b_2} \emptyset$. The collection \succ_B is *acyclic* if it does not contain a cycle.

A complete priority ordering is acyclic if and only if the following property holds: for any applicant, there is no more than one other applicant who has higher priority at some buildings but lower priority at other buildings. In the context of public-housing allocation, the acyclicity condition can be satisfied by an agency that administers multiple programs which share a ranking over applicants but have some discretion on final priorities based on interviews.¹⁵ The acyclicity condition is more permissive when priority orderings differ in terms of which applicants are acceptable. For example, the case of a public-housing agency that uses a common point scale to determine priorities but restricts access to certain buildings (e.g., senior housing) satisfies the acyclicity condition. We will show that a generalized version of MWP satisfies the properties in Proposition 2 for acyclic priority orderings.

The generalized Multiple-Waitlist Procedure is similar to MWP except that the order of the applicants' turns may be switched when one applicant prefers a building which gives another applicant higher priority. In particular, an applicant who refuses an offer and prefers to join a waiting list for a different building must wait for the applicant (among those who have not yet been assigned a unit) with the highest priority at that particular building to choose first. In the language of Abdulkadiroğlu and Sönmez [1999], one way to describe this mechanism would be “you request my building—I get your turn.” Figure 2 provides a more formal description of the generalized MWP.

The following result, shown formally in Appendix A.4, characterizes the generalized MWP under acyclic priority orderings.

Proposition 4. *If the priority orderings are acyclic, then the generalized MWP satisfies (i) strategy-proofness, (ii) efficiency, and (iii) elimination of justified envy.*

By exhibiting a mechanism that satisfies the three properties for acyclic priority orderings, this result demonstrates constructively that the acyclicity condition is

¹⁵See, for example, the allocation procedure in the [District of Columbia](#).

sufficient for the existence of a mechanism that satisfies the desired properties. Additionally, the acyclicity condition is necessary for the existence of such a mechanism: to show this, Appendix A.4 provides an example of a deterministic arrival process (in an environment with two buildings that have cyclic priorities over three applicants) such that every possible allocation exhibits inefficiency or justified envy. The following proposition summarizes these findings.

Proposition 5. *There exists a strategy-proof allocation mechanism that is efficient and free of justified envy if and only if the priority orderings are acyclic.*

Recall from the framework that a mechanism assigns a particular unit in a particular building to each applicant. A natural question that arises is whether the characterization holds more generally for *stochastic allocation mechanisms*, i.e., mechanisms that assign a lottery over units to each applicant. Under this more general class of mechanisms, the concept of elimination of justified envy from Definition 1 does not apply since lotteries do not have priority orderings over applicants.

To address this issue, we use the notion of *strong stability* first introduced by Roth et al. [1993] and adapted by Kesten and Ünver [2015] in the context of school-choice lotteries.¹⁶ We say that an applicant a' *strongly envies* another applicant a if there is a unit r in building b such that a can be assigned to $\langle b, r \rangle$ with positive probability while a' can be assigned to a less desirable unit (for her) than $\langle b, r \rangle$ with positive probability, and we say that this strong envy is *justified* if a' has higher priority than a for building b . A stochastic allocation mechanism is *strongly stable* if it eliminates justified strong envy.¹⁷

As the following result demonstrates, the characterization in Proposition 5 extends to stochastic allocation mechanisms.

Proposition 6. *There exists a strategy-proof stochastic allocation mechanism that satisfies efficiency and strong stability if and only if the priority orderings are acyclic.*

¹⁶Kesten and Ünver [2015] refer to this notion as *ex-ante stability* and provide a more general formulation that allows for weak priorities.

¹⁷Under non-stochastic allocation mechanisms, strong stability coincides with the elimination of justified envy because only a single unit is assigned to each agent with positive probability.

Even when considering the more general class of stochastic allocation mechanisms, the proof in Appendix A.5 shows that the presence or absence of cycles in the priority orderings fully characterizes the possibility for a strategy-proof allocation mechanism to satisfy efficiency and the elimination of justified envy.

Due to the incompatibility between efficiency and the elimination of justified envy in the absence of acyclic priority orderings, we suggest strategy-proof allocation mechanisms for arbitrary priority orderings that satisfy each of these criteria separately (see Thakral [2015] for more detail).

A simple variation of MWP that satisfies efficiency would be to choose any ordering of the applicants and apply MWP as if this ordering were the common priority ordering. The ordering can be dynamically constructed, e.g., by choosing the applicant with the highest priority at the building that becomes available. This class of procedures produces efficient allocations: no applicant would benefit from an alternate allocation since each chooses her most-preferred unit at the time of assignment. However, these procedures do not eliminate justified envy since an applicant may choose a place on a waiting list for a building at which her priority is low.

A modified version of MWP can achieve the elimination of justified envy by constructing the priority ordering dynamically (as described above) but restricting the set of waiting lists that an applicant may join. A procedure that only allows an applicant to join a given queue if the applicant has top priority at the associated building satisfies the no-envy condition since any household that envies the applicant necessarily has lower priority. The mechanism permits inefficiencies since applicants who prefer units in buildings at which their respective priorities are low may benefit from switching their assignments.

4 Comparison of Allocation Mechanisms

We begin by evaluating the theoretical properties of mechanisms that are used to assign public housing and then proceed to investigate public-housing allocation mechanisms empirically. Using estimated preferences for public housing from a structural model

due to Geyer and Sieg [2013], we find substantial ex-post welfare gains from changing the existing public-housing allocation mechanism to the Multiple-Waitlist Procedure described in Section 3.

4.1 Existing mechanisms

Public-housing applicants throughout the world express little choice in their place of residence under the most widely used allocation mechanisms. Applicants receive offers after units become available and can refuse only a limited number of times before forfeiting their position on the waiting list. This section provides institutional details on public-housing allocation in the United States (see Thakral [2017] for more details) and characterizes existing mechanisms.

A *Public Housing Authority* (PHA) is a state-run or locally-run entity that administers federal housing assistance programs. There are about 3,300 such agencies in the United States with approximately 1.2 million households living in public housing. The US Department of Housing and Urban Development (HUD) authorizes and funds PHAs and suggests two types of procedures that a PHA may use to allocate units, both of which involve finding the highest-priority applicant who is willing to accept the available unit rather than refuse. Under *Plan A*, the PHA offers a unit that becomes available to the applicant with the highest priority; if the applicant refuses, then the applicant is removed from or placed at the bottom of the waiting list. Under *Plan B*, an applicant who refuses a unit receives another offer, up to a limit of two or three total offers [Devine et al., 1999]. More generally, letting $k \in \mathbb{N}$ denote the maximum total number of units that the PHA offers to a given applicant, we refer to these procedures as PHA- k mechanisms. The top panel of Table 1 shows the number of housing agencies using these procedures. PHA-1 corresponds to Plan A, the take-it-or-leave-it procedure in which no applicant can refuse an offer without losing their position on the waiting list.

As mentioned in the introduction, the take-it-or-leave-it mechanism can lead to unfair and inefficient allocations. This observation applies more generally to the entire class of PHA- k mechanisms.

Proposition 7. *The PHA- k mechanisms are inefficient and do not eliminate justified envy.*

Appendix A.6 provides two separate proofs of this result, each consisting of an example of a deterministic arrival process for which the mechanism results in justified envy and inefficiency. The first example contains only two buildings, and an applicant who refuses an offer from one building may receive another offer from the same building. The second example, in which the number of buildings depends on k , applies even if the mechanism requires that an applicant who refuses an offer does not receive another offer from the same building. Note from the top panel of Table 1, which displays the number of PHAs using the PHA- k mechanism for $k \in \{1, 2\}$ and for $k > 2$, that k is typically small in practice.

The first-come-first-served (FCFS) allocation mechanism can be thought of as a PHA- k mechanism with $k \rightarrow \infty$. Although not widespread in practice, this mechanism appears in several theoretical papers [Su and Zenios, 2004, Bloch and Cantala, 2017, Schummer, 2016]. Declining an offer creates a positive externality by reducing other agents’ waiting time, but the agent at the top of the waiting list prefers to receive an allocation sooner due to private waiting costs.¹⁸ Van Ommeren and Van der Vlist [2016] present empirical evidence from Amsterdam that the FCFS mechanism for allocating public housing produces an inefficient matching.

The properties of the PHA- k mechanisms also depend on the format of the waiting list. Under a *centralized* waiting list system, the applicant with the highest priority can be offered a unit in any building that becomes available. A *non-centralized* waiting list is either *site-specific* or *sub-jurisdictional*, depending on whether an applicant can only be assigned a unit in a particular building or in a group of buildings. The bottom panel of Table 1 shows the prevalence of centralized waiting lists. MWP uses both types of waiting lists: the PHA offers one unit to a household on the centralized waiting list and moves them to their chosen site-specific waiting list if the offer is refused. An alternative mechanism that places applicants on two (or

¹⁸As Su and Zenios [2004] argue, FCFS allocation mechanisms lead to an “inherent inefficiency [because of their] inability...to contain the externalities generated by [applicants’] self-serving behavior.”

more) site-specific waiting lists and then makes take-it-or-leave-it offers would fail to be strategy-proof. Consider an applicant with low waiting costs and a strong preference for a particular building: to maximize the chance of receiving a unit in that building, the applicant may report as her second choice a popular building with a long waiting list (even if the popular building lies at the bottom of her true preference ordering). Details about how the waiting lists are constructed can affect applicants' incentives, and procedures that are generally used in practice for assigning applicants to non-centralized waiting lists tend not to be strategy-proof.¹⁹

A potential justification for the use of the take-it-or-leave-it mechanism and its variants might be that these mechanisms act as screening devices, but several facts about housing policy suggest otherwise. First, the application process itself functions as a screening device: applications are costly to fill out, housing authorities verify the information reported on applications to determine priorities, and many housing authorities conduct in-person interviews of households that approach the top of the waiting list.²⁰ Second, some housing authorities use mechanisms that offer no additional screening benefit: FCFS allocation mechanisms, for example, do not remove applicants from the waiting list after declining an offer.²¹ Third, housing authorities typically aim to minimize the length of time required to fill a vacant unit: this is perhaps the most likely justification for using procedures in which each applicant receives a limited number of offers. MWP requires that the housing authority makes only a single offer to each applicant and thus performs equally as well as the take-it-or-leave-it procedure (and better than its variants) by this measure.

4.2 Estimation of welfare gains

Although the theoretical results establish that MWP is ex-ante efficient, a question that remains is whether the choice of an allocation mechanism affects ex-post welfare

¹⁹The [New York City Housing Authority](#), which uses sub-jurisdictional waiting lists (one for each of the five boroughs), asks applicants to report their first and second borough choice and explicitly advises applicants to “select their first borough choice carefully.”

²⁰See [HUD \[2009\]](#) for more details.

²¹The City of Toronto adopted a FCFS allocation mechanism in July 2014. Similar mechanisms are also used in Britain and the Netherlands.

in real-world settings. We address this question by investigating the allocation of public housing using a sample of 215 eligible households in Pittsburgh, PA from the 2001 Survey of Income and Program Participation (SIPP) collected by the US Census Bureau. The model of household preferences consists of a standard random-utility specification as in [Geyer and Sieg \[2013\]](#). Our goal is to quantify the welfare gains from changing the allocation procedure by simulating arrival processes and matchings under counterfactual mechanisms.

4.2.1 Buildings

Buildings are classified by size (small, fewer than 40 units; medium, between 40 and 100 units; or large, more than 100 units) and community type (family, senior, or mixed). In practice, any applicant can reside in any building. The 34 buildings operated by the Housing Authority of the City of Pittsburgh (HACP) in 2001 fall within six categories based on size and community type: family large, family medium, family small, mixed, senior large, and senior small. Buildings have a common priority ordering over applicants.

Table 2 displays the number of units in each building category that became available over a five-year period. From these data we estimate a binomial arrival process.²² In any given building, a unit arrives with an estimated probability of approximately 0.17 each week.²³

4.2.2 Applicants

The utility of applicant $i \in A$ living in building $j \in B$ at time $t \in \mathbb{N}$ is given by

$$u_{i,j,t} = \beta \log y_{i,t} + \gamma_j + \kappa x_i - c \mathbf{1}_{\{d_{i,t} \neq d_{i,t-1}\}} + \varepsilon_{i,j,t}, \quad (1)$$

²²[Kaplan \[1984\]](#) argues that the length of time that a household lives in public housing follows an exponential distribution; in a discrete-time setting (where the length of time in public housing follows a geometric distribution), the “moveout process” follows a binomial distribution.

²³The results in this section are unaffected if we simulate a daily or monthly arrival process.

where $y_{i,t}$ denotes net income; γ_j is a building-category fixed effect; x_i is a vector of demographic characteristics, namely indicators for female, nonwhite, senior, and children; $d_{i,t} \in B \cup \{0\}$ denotes the residence and c is a moving-cost parameter; and $\varepsilon_{i,j,t}$ captures idiosyncratic tastes for public housing. We normalize the utility of living in private housing ($j = 0$) to be

$$u_{i,0,t} = \log y_{i,t} - c\mathbf{1}_{\{d_{i,t} \neq d_{i,t-1}\}} + \varepsilon_{i,0,t}. \quad (2)$$

Following [McFadden \[1973\]](#), we assume that the idiosyncratic components are independently and identically distributed according to a standard type-I extreme-value distribution.²⁴

A household that lives in private housing does not necessarily *prefer* to live there over public housing. Due to supply-side restrictions, households exhibit preferences for public housing by joining the waiting list.²⁵ A simple logit demand model based on observed choices would therefore fail to capture the reality of strong preferences for public housing. [Geyer and Sieg \[2013\]](#) develop an equilibrium framework that incorporates rationing and excess demand to model public-housing allocation. They identify the structural parameters of the utility function above using household exit behavior.²⁶ [Table 3](#) reproduces estimates of the structural parameters from [Geyer and Sieg \[2013\]](#), which are based on household-level data from the HACP and the SIPP.²⁷

Using a sample of low-income households eligible for public housing in Pittsburgh

²⁴[Geyer and Sieg \[2013\]](#) find that using a nested logit specification designed to account for correlation in unobserved preferences among public-housing communities does not improve the fit of the model.

²⁵A household in Pittsburgh typically waits between 14 and 22 months for a unit [[Geyer and Sieg, 2013](#)].

²⁶Identification relies on the assumption of voluntary exit. The fact that housing authorities do not evict “over-income households” (i.e., those that exceed income thresholds) supports this assumption.

²⁷The estimates suggest that minorities and female-headed households with children exhibit stronger preferences for public housing than other households. The fact that the coefficient β on log income is less than one suggests that an increase in income makes public housing less desirable than the outside option, consistent with the fact that a household residing in public housing pays 30 percent of its income as rent.

from the SIPP, we construct preferences based on the structural model above.²⁸ Since equations (1) and (2) describe the utility of living in a given building in a given period, a complete description of preferences requires assumptions about the discount factor and the match duration. Given our objective of measuring the welfare gains from using a mechanism that provides agents more choice at the cost of additional waiting time, we make conservative assumptions which cause our estimates to understate the true gains. A low discount factor increases the cost of waiting for a more preferred building, and a low match duration decreases the benefit of being matched with a more-preferred building. We assume that applicants maximize a discounted sum of utilities with a weekly discount factor of $\delta \approx 0.99$. Moreover, every match lasts for three years, so the applicant resides in private housing in all periods except for the three years immediately after moving into public housing. Section 4.3 provides further discussion on these assumptions.

4.2.3 Mechanisms

For each simulated arrival process, we determine the allocations that would result under each of the following mechanisms: MWP, PHA-1, PHA-2, and PHA- ∞ . During the sample period, the existing housing-allocation procedure is the take-it-or-leave-it mechanism (PHA-1).

MWP 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.

Under PHA-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.

PHA-2 gives applicants the opportunity to refuse an offer once. This decision depends on the applicant's beliefs about the households on the waiting list that have higher priority. We assume that the applicant knows how many households are ahead

²⁸Annual gross income in an eligible household must fall below 80 percent of the Area Median Income (AMI).

on the waiting list but does not know their preferences; instead, the applicant knows the distribution of preferences.²⁹ We use numerical approximations, drawing from the distributions of preferences and future arrivals, to compute the expected utility of refusing an offer.

PHA- ∞ (FCFS) allows applicants to refuse an unlimited number of offers. We do not attempt to characterize equilibrium behavior under this mechanism; Bloch and Cantala [2017] study a model in which agents have private values and ultimately focus on equilibrium behavior under the two-agent case due to difficulties in providing a general characterization. For a given applicant, a state can be described by the number of households that are ahead in the waiting list and the number of offers that each of them has already refused. Any subset of the households ahead in the waiting list can still be on the waiting list by the time the applicant gets another offer, and each refusal reveals some additional information about the preferences of those households. This yields a number of states that is exponential in the size of the waiting list, which creates difficulties in trying to approximate equilibrium behavior. One way to proceed would be to introduce some heuristics by which applicants make such complex decisions. However, the resulting estimates would only provide a lower bound for welfare under the mechanism, as the heuristics might be too simplistic. Instead of assuming that applicants make suboptimal decisions, we make generous assumptions about the environment to obtain an upper bound for welfare under this mechanism. Specifically, we assume that applicants know the realization of the arrival process in advance and have complete information about the preferences of other households.³⁰ We denote the resulting upper bound on welfare by PHA-max.

²⁹The applicant also knows that any households that are ahead on the waiting list will accept the next offer if and only if the unit is acceptable (since they have already declined an offer).

³⁰This assumption could in theory lead to an underestimate of welfare if lower-priority applicants have stronger preferences for public housing, but this concern does not arise in our setting: first, priorities are randomly drawn in each simulation; second, the model incorporates limited heterogeneity in household preferences, as Section 4.3 discusses in more detail.

4.2.4 Welfare

Our counterfactual simulations provide evidence that the welfare gain from changing the public-housing allocation mechanism to MWP is substantial. We convert the difference in utilities between each mechanism and PHA-1 to monetary values by computing the equivalent variation (EV), i.e., the transfer that the applicant would have to receive when public housing is assigned by PHA-1 that would give the applicant the same lifetime utility as the assignment under the new mechanism. Table 4 reports the average of the present-discounted values of these payments for PHA-2, PHA-max, and MWP.

Our estimates suggest that a change from PHA-1 to MWP improves the welfare of the average applicant who receives a housing assignment by an amount equivalent to a one-time transfer payment of between \$6,100 and \$6,700.³¹ Providing applicants with some choice in the allocation process by using PHA-2 rather than PHA-1 achieves 60 percent of this gain. By allowing applicants to express additional choice, MWP provides further improvement and falls within about 5 percent of PHA-max. Based on these estimates, we conclude that MWP outperforms PHA- ∞ . Recall that PHA-max is an idealized bound which assumes that applicants have complete information about other households' preferences and perfect foresight about the realization of the stochastic arrival process, which cannot be attained in practice. Although a precise prediction about welfare under PHA- ∞ would require assumptions about applicants' strategies, the difficulty in characterizing equilibrium behavior suggests that real-world decision making under PHA- ∞ would be far from optimal.

Note that the welfare measure only accounts for matches and ignores the costs of processing refusals and leaving units vacant. In our simulations, the next agent on the waiting list receives an offer immediately after a refusal. In practice, however, housing authorities typically provide three to five days for applicants to respond to an offer, which leads to longer vacancies. Due to underutilization, a mechanism such as PHA-2 or PHA- ∞ that makes multiple offers to each applicant produces lower

³¹We exclude applicants at the top of the centralized waiting list (before more than one applicant accepts an offer in the present period), as these applicants may benefit from facing empty site-specific waiting lists when our simulation begins.

social welfare than our estimates suggest.

4.3 Interpretation of magnitudes

Given that a typical Pittsburgh household living in public housing earns less than \$15,000 annually, MWP substantially improves ex-post welfare.³²

As one way to assess the economic significance of the welfare gain, we compare our result with estimates of the benefits and costs of public housing. We measure the benefit of living in public housing using the equivalent variation, i.e., the transfer that each applicant would have to receive in private housing to give the applicant the same utility as the public-housing assignment under PHA-1. According to this measure, living in public housing is equivalent to a \$14,412 increase in annual income. The gain from changing the allocation mechanism to MWP corresponds to an annual transfer payment of \$2,572, which represents an 18 percent increase in the benefit of public housing. Moreover, the change in the allocation mechanism increases the cost effectiveness of public housing by 14 percentage points.³³

As another way to evaluate the magnitude of the gain, we express the estimates relative to commuting times. The welfare gain corresponds to a daily transfer payment of \$9, or equivalently 105 minutes at the minimum hourly wage (\$5.15 from 2001 to 2006). According to the 2000 Census, a low-income worker in Pittsburgh spends on each working day an average of 46 minutes commuting by car or 83 minutes commuting by public transit. In that sense, the gain from changing the allocation mechanism exceeds the gain from eliminating commuting times.

Finally, we compare the welfare gain to the maximum possible gain that could be achieved by a social planner that knows the complete realization of the arrival process in advance. As the bottom panel of Table 4 shows, the ex-post optimal allocation

³²The median total annual income (which includes means-tested transfers) of households eligible for public housing in our sample is \$14,184. Households living in public housing earn an average of \$9,082 in the HACP data as reported by [Geyer and Sieg \[2013\]](#).

³³Using data from New York City, [Olsen and Barton \[1983\]](#) estimates a resource cost for providing a public-housing unit—consisting of loan payments for initial development costs, property taxes, operating costs—of \$18,049 (inflated to 2006 dollars using the Consumer Price Index, and adjusted for administrative costs as pointed out by [Olsen \[2003\]](#)).

would result in an average welfare gain of \$8,505, which provides an upper bound for the maximum possible gain that any mechanism can reach. MWP attains 75 percent of this perfect-foresight optimal-allocation benchmark.

We interpret our welfare estimates as providing a lower bound on the gains from changing the allocation mechanism to MWP, as discussed below.

First, the assumptions about timing in the simulations lead to an underestimate of the welfare gains. Note that households in our simulations spend three years in public housing, compared to almost seven years according to the HACP data as reported by [Geyer and Sieg \[2013\]](#). A low match duration decreases the benefit of being matched with a more preferred building. If households with higher utility from public housing spend more time living in public housing, then using a fixed match duration further understates the average gains. Next note that the weekly discount factor of about 0.99 corresponds to an annual discount factor of 0.6, which falls below most estimates in the literature.³⁴ A low discount factor increases the cost of waiting for a more preferred building. Both assumptions about timing lead to smaller estimated gains from a mechanism that allows applicants to wait longer for better matches.

Second, the measure of welfare places equal weight on all households. This might not be the most realistic welfare objective because of the fact that housing authorities assign priorities to each applicant. Under an objective that gives more weight to higher-priority applicants, the welfare gains are even larger.

Third, the structural model understates the extent of heterogeneity in preferences. Note that the model only includes six building-category fixed effects instead of a full set of building-specific fixed effects. Moreover, equation (1) does not contain any interaction between household characteristics and the building-category fixed effects, which implies that building categories are vertically differentiated and households on average only differ in terms of waiting costs. Some forms of heterogeneity depend on persistent idiosyncratic factors: applicants may prefer to live closer to their respective workplaces, allowing them to reduce commuting time and travel costs [[Lui and Suen,](#)

³⁴[Laibson et al. \[2015\]](#) point out that simultaneously fitting data on wealth accumulation and credit-card borrowing yields an annual exponential discount factor of 0.63, in contrast with typical estimates in the literature of around 0.95.

2011]. A more realistic model would consist of more heterogeneity in preferences from serially correlated error terms. Applicant characteristics also contribute to heterogeneity: households with children may prefer units that are located near better schools, and seniors may prefer units that have access to certain amenities. Using a model that incorporates these aspects, the estimated gain from using an allocation mechanism that improves match quality would be even larger.

Finally, the magnitude of our welfare estimate does not arise due to assumptions about unobserved heterogeneity in preferences. The idiosyncratic preference component can only affect the decision to accept an offer versus join a waiting list—but not which waiting list to join—because households can only condition on the realization of the idiosyncratic shock for the present period. Intuitively, the situation when the idiosyncratic term can affect decisions arises only when a household is on the margin of accepting an offer, when the difference in utilities is smallest. Additional simulations confirm that the estimates are not sensitive to assumptions about functional forms. Taking the extreme approach of entirely omitting idiosyncratic preferences, the point estimate for the welfare gain of MWP still lies within the 95 percent confidence interval presented in Table 4 at approximately the lower bound. Therefore we conclude that the estimated welfare gains are not driven by the extent of unobserved heterogeneity imposed by the structural model. To the extent that households’ preferences in the real world do exhibit significant unobserved heterogeneity that the model does not capture, this further suggests that our results underestimate the actual welfare gains.

5 Discussion

We conclude with a discussion of practical matters pertaining to MWP in the context of public-housing allocation.

First, with over 1.2 million households living in public housing in the US alone, the overall gains from improved matching are substantial. Using a sample of households eligible for public housing in Pittsburgh, we find a lower bound to the welfare gains from changing the most commonly used take-it-or-leave-it allocation mechanisms to

MWP of \$6,429 per applicant.³⁵ Although our theoretical analysis does not focus on ex-post welfare maximization, which may require specific assumptions about the arrival process, MWP performs well relative to the ex-post optimal allocation.

Second, MWP allows applicants to express choice without creating additional delays. Expanding choice by allowing applicants to refuse multiple offers creates delays in filling vacancies, which in practice results in welfare losses. The media often publicizes the average number of days that units remain vacant because vacancies lead to losses in rent, which imposes a burden on taxpayers when the federal government has to subsidize housing authorities for repairs and other operating costs. As a recent example, rent losses due to vacancies in New York City exceed \$8 million [Navarro, 2015]. Under MWP, households choose among a set of waiting lists rather than accept or reject units after they arrive. By only requiring housing authorities to make a single offer to each applicant, MWP provides both choice and utilization.

Third, households face a simple decision problem under MWP, which is strategy-proof. In the direct mechanism, an applicant must report (i) a ranking over buildings, and (ii) for each building, the number of periods the applicant would be willing to wait to receive a unit in her most-preferred building rather than receiving a unit in that building immediately. In the indirect mechanism, an applicant must select among a set of units associated with expected waiting times given information about the arrival process, which a housing authority can provide [Kaplan, 1986]. Existing mechanisms, by contrast, require difficult computations for agents to behave optimally.

Finally, although the model involves various stylized assumptions, the general lessons apply in more realistic situations. The main results hold when we relax assumptions about time or risk preferences as well as in the presence of unobserved heterogeneity in applicants' preferences over units. Moreover, if additional applicants can join the waiting list or the priority ordering can change over time, a housing authority can implement MWP without any modification.

³⁵These measures do not account for objectives such as racial or economic integration of housing projects; Kaplan [1987] discusses how priority-assignment policies can achieve such goals.

References

- Atila Abdulkadirođlu and Tayfun Sönmez. House allocation with existing tenants. *Journal of Economic Theory*, 88(2):233–260, October 1999. ISSN 0022-0531. doi: 10.1006/jeth.1999.2553. URL <http://www.sciencedirect.com/science/article/pii/S002205319992553X>.
- Atila Abdulkadirođlu, Parag A. Pathak, and Alvin E. Roth. Strategy-proofness versus efficiency in matching with indifferences: Redesigning the NYC high school match. *American Economic Review*, 99(5):1954–1978, 2009. doi: 10.1257/aer.99.5.1954. URL <http://www.aeaweb.org/articles.php?doi=10.1257/aer.99.5.1954>.
- Atila Abdulkadirođlu, Nikhil Agarwal, and Parag A. Pathak. The welfare effects of coordinated assignment: Evidence from the NYC HS match. Working Paper 21046, National Bureau of Economic Research, March 2015a. URL <http://www.nber.org/papers/w21046>.
- Atila Abdulkadirođlu, Yeon-Koo Che, and Yosuke Yasuda. Expanding “choice” in school choice. *American Economic Journal: Microeconomics*, 7(1):1–42, 2015b.
- Nikhil Agarwal and Paulo Somaini. Demand analysis using strategic reports: An application to a school choice mechanism. Working Paper 20775, National Bureau of Economic Research, December 2014. URL <http://www.nber.org/papers/w20775>.
- Mohammad Akbarpour, Shengwu Li, and Shayan Oveis Gharan. Thickness and information in dynamic matching markets. *Available at SSRN 2394319*, 2016.
- Mariagiovanna Baccara, SangMok Lee, and Leeat Yariv. Optimal dynamic matching. *Mimeo*, 2015.
- Michel Balinski and Tayfun Sönmez. A tale of two mechanisms: Student placement. *Journal of Economic Theory*, 84(1):73–94, 1999.
- Francis Bloch and David Cantala. Dynamic assignment of objects to queuing agents. *American Economic Journal: Microeconomics*, 9(1):88–122, 2017.

- Caterina Calsamiglia, Chao Fu, and Maia Güell. Structural estimation of a model of school choices: the Boston mechanism vs. its alternatives. *Mimeo*, 2016.
- Deborah J. Devine, Lester Rubin, and Robert W. Gray. *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, 1999.
- Laura Doval. A theory of stability in dynamic matching markets. *Mimeo*, 2016.
- Umut Dur. Tie-breaking is still a problem: How to find the pareto efficient and fair matching in the school choice problem. *Mimeo*, 2011.
- Umut Dur. Dynamic school choice problem. *Mimeo*, 2012.
- Aytek Erdil and Haluk Ergin. What’s the matter with tie-breaking? Improving efficiency in school choice. *American Economic Review*, 98(3):669–689, 2008. doi: 10.1257/aer.98.3.669. URL <http://www.aeaweb.org/articles.php?doi=10.1257/aer.98.3.669>.
- Haluk I. Ergin. Efficient resource allocation on the basis of priorities. *Econometrica*, 70(6):2489–2497, November 2002. ISSN 1468-0262. doi: 10.1111/j.1468-0262.2002.00447.x. URL <http://onlinelibrary.wiley.com/doi/10.1111/j.1468-0262.2002.00447.x/abstract>.
- Sebastian Galiani, Alvin Murphy, and Juan Pantano. Estimating neighborhood choice models: Lessons from a housing assistance experiment. *American Economic Review*, 105(11):3385–3415, 2015. doi: 10.1257/aer.20120737. URL <http://www.aeaweb.org/articles.php?doi=10.1257/aer.20120737>.
- Alex Gershkov and Benny Moldovanu. Learning about the future and dynamic efficiency. *American Economic Review*, 99(4):1576–1587, 2009. doi: 10.1257/aer.99.4.1576. URL <http://www.aeaweb.org/articles.php?doi=10.1257/aer.99.4.1576>.

- Alex Gershkov, Benny Moldovanu, and Philipp Strack. Efficient dynamic allocation with strategic arrivals. *Available at SSRN 2548740*, 2015.
- Judy Geyer and Holger Sieg. Estimating a model of excess demand for public housing. *Quantitative Economics*, 4(3):483–513, 2013. ISSN 1759-7331. doi: 10.3982/QE148. URL <http://onlinelibrary.wiley.com/doi/10.3982/QE148/abstract>.
- Edward L. Glaeser and Erzo F. P. Luttmer. The misallocation of housing under rent control. *American Economic Review*, 93(4):1027–1046, 2003. doi: 10.1257/000282803769206188. URL <http://www.aeaweb.org/articles.php?doi=10.1257/000282803769206188>.
- Hank Greely. The equality of allocation by lot. *Harvard Civil Rights-Civil Liberties Law Review*, 12:113, 1977.
- Isa E. Hafalir, M. Bumin Yenmez, and Muhammed A. Yildirim. Effective affirmative action in school choice. *Theoretical Economics*, 8(2):325–363, 2013. ISSN 1555-7561. doi: 10.3982/TE1135. URL <http://onlinelibrary.wiley.com/doi/10.3982/TE1135/abstract>.
- Yinghua He. Gaming the Boston school choice mechanism in Beijing. *Mimeo*, 2012.
- HUD. A guide to interviewing for owners of HUD-subsidized multifamily housing projects. Multifamily Housing Case Studies, August 2009. URL <http://www.hud.gov/offices/hsg/mfh/rhiip/interviewguide.pdf>.
- A. Hylland and R. Zeckhauser. The efficient allocation of individuals to positions. *The Journal of Political Economy*, pages 293–314, 1979.
- Edward H. Kaplan. *Managing the Demand for Public Housing*. PhD thesis, Massachusetts Institute of Technology, 1984.
- Edward H. Kaplan. Tenant assignment models. *Operations Research*, 34(6):832–843, 1986.

- Edward H. Kaplan. Analyzing tenant assignment policies. *Management Science*, 33(3): 395–408, 1987. ISSN 0025-1909. URL <http://www.jstor.org/stable/2631859>.
- Onur Kesten. On two competing mechanisms for priority-based allocation problems. *Journal of Economic Theory*, 127(1):155–171, March 2006. ISSN 0022-0531. doi: 10.1016/j.jet.2004.11.001. URL <http://www.sciencedirect.com/science/article/pii/S0022053104002194>.
- Onur Kesten and M. Utku Ünver. A theory of school choice lotteries. *Theoretical Economics*, 10(2):543–595, 2015.
- Fuhito Kojima. Robust stability in matching markets. *Theoretical Economics*, 6(2):257–267, May 2011. ISSN 1555-7561. doi: 10.3982/TE780. URL <http://onlinelibrary.wiley.com/doi/10.3982/TE780/abstract>.
- David Laibson, Peter Maxted, Andrea Repetto, and Jeremy Tobacman. Estimating discount functions with consumption choices over the lifecycle. *Mimeo*, 2015.
- Jacob D. Leshno. Dynamic matching in overloaded waiting lists. *Mimeo*, 2015.
- Shengwu Li. Obviously strategy-proof mechanisms. SSRN Scholarly Paper ID 2560028, Social Science Research Network, November 2015. URL <http://papers.ssrn.com/abstract=2560028>.
- Hon-Kwong Lui and Wing Suen. The effects of public housing on internal mobility in Hong Kong. *Journal of Housing Economics*, 20(1):15–29, March 2011. ISSN 1051-1377. doi: 10.1016/j.jhe.2010.11.001. URL <http://www.sciencedirect.com/science/article/pii/S1051137710000471>.
- Daniel McFadden. Conditional logit analysis of qualitative choice behavior. In *Frontiers in Econometrics*, pages 105–142. Academic Press: New York, 1973.
- Thayer Morrill. Making efficient school assignment fairer. *Mimeo*, 2013.
- Mireya Navarro. New York City public housing units remain empty unnecessarily, audit finds. *The New York Times*, June 2015. ISSN

- 0362-4331. URL <http://www.nytimes.com/2015/06/25/nyregion/new-york-city-public-housing-units-remain-empty-unnecessarily-audit-finds.html>.
- Edgar O. Olsen. Housing programs for low-income households. *NBER*, pages 365–442, January 2003. URL <http://www.nber.org/chapters/c10259>.
- Edgar O. Olsen and David M. Barton. The benefits and costs of public housing in New York City. *Journal of Public Economics*, 20(3):299–332, April 1983. ISSN 0047-2727. doi: 10.1016/0047-2727(83)90028-2. URL <http://www.sciencedirect.com/science/article/pii/0047272783900282>.
- Antonio Romero-Medina and Matteo Triossi. Games with capacity manipulation: Incentives and Nash equilibria. *Social Choice and Welfare*, 41(3):701–720, September 2013. ISSN 0176-1714, 1432-217X. doi: 10.1007/s00355-012-0703-1. URL <http://link.springer.com/article/10.1007/s00355-012-0703-1>.
- Alvin E. Roth, Uriel G. Rothblum, and John H. Vande Vate. Stable matchings, optimal assignments, and linear programming. *Mathematics of Operations Research*, 18(4):803–828, November 1993. ISSN 0364-765X. doi: 10.1287/moor.18.4.803. URL <http://pubsonline.informs.org/doi/abs/10.1287/moor.18.4.803>.
- James Schummer. Influencing waiting lists. *Mimeo*, 2016.
- Heather L. Schwartz. Housing policy is school policy. *A Century Foundation Report*, 2010. URL <https://www.tcf.org/assets/downloads/tcf-Schwartz.pdf>.
- Adriana Mihaela Soaita. Strategies for *In Situ* home improvement in Romanian large housing estates. *Housing Studies*, 27(7):1008–1030, October 2012. ISSN 0267-3037. doi: 10.1080/02673037.2012.725833. URL <http://dx.doi.org/10.1080/02673037.2012.725833>.
- Tayfun Sönmez, M Utku Ünver, et al. Matching, allocation, and exchange of discrete resources. *Handbook of Social Economics*, 1:781–852, 2011.

- Xuanming Su and Stefanos Zenios. Patient choice in kidney allocation: The role of the queueing discipline. *Manufacturing & Service Operations Management*, 6(4): 280–301, 2004.
- Lars-Gunnar Svensson. Strategy-proof allocation of indivisible goods. *Social Choice and Welfare*, 16(4):557–567, September 1999. ISSN 0176-1714, 1432-217X. doi: 10.1007/s003550050160. URL <http://link.springer.com/article/10.1007/s003550050160>.
- Neil Thakral. Matching with stochastic arrival. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, EC '15, pages 343–343, New York, NY, USA, 2015. ACM. ISBN 978-1-4503-3410-5. doi: 10.1145/2764468.2764483. URL <http://doi.acm.org/10.1145/2764468.2764483>.
- Neil Thakral. The design of public-housing policies. *Mimeo*, 2017.
- M. Utku Ünver. Dynamic kidney exchange. *The Review of Economic Studies*, 77(1): 372–414, January 2010. ISSN 0034-6527, 1467-937X. doi: 10.1111/j.1467-937X.2009.00575.x. URL <http://restud.oxfordjournals.org/content/77/1/372>.
- Jos N. Van Ommeren and Arno J. Van der Vlist. Households' willingness to pay for public housing. *Journal of Urban Economics*, 92:91–105, 2016.
- Shing-Yi Wang. State misallocation and housing prices: Theory and evidence from China. *American Economic Review*, 101(5):2081–2107, 2011. doi: 10.1257/aer.101.5.2081. URL <http://www.aeaweb.org/articles.php?doi=10.1257/aer.101.5.2081>.

A Proofs

A.1 Ex-post impossibility

Proof of Proposition 1. We will show by example that no mechanism can guarantee ex-post efficiency or ex-post elimination of justified envy for every arrival process.

Consider the following example with three periods (0, 1, and 2), three buildings (α , β , and γ), and three applicants (A, B, and C).

Assume that each building gives applicant A the highest priority. The applicants' preferences are given in Appendix Table 1.

Let the arrival process be specified as follows. In each period, a unit becomes available with certainty: the unit that becomes available in period 0 is in building α ; the unit that becomes available in period 1 is equally likely to be in building β or in building γ ; and the unit that becomes available in period 2 is equally likely to be in the building from which no unit has become available yet or in building α .

Suppose applicant A is assigned building α in period 0. With probability 1/2, building β becomes available in period 1. There are two cases to consider. First, suppose β is assigned to applicant B. Then with probability 1/2, building γ arrives in period 2 and is allocated to applicant C. Notice that the allocation is ex-post inefficient because applicants A and B prefer to switch: since $(\beta, 1) \succ_A (\alpha, 1)$ and $(\alpha, 1) \succ_B (\beta, 1)$, we see that A and B prefer to leave their assigned buildings to switch with each other in period 1 (or any subsequent period). Second, suppose building β is assigned to applicant C when it becomes available in period 1. Then with probability 1/2, building γ arrives in period 2 and is allocated to applicant B. Again, the allocation is ex-post inefficient because applicants A and C prefer to switch.

The analysis is similar if applicant B or applicant C is assigned building α in period 0. Regardless of which applicant is assigned building α in period 0, there is always some realization of the arrival process in which a Pareto improvement can be found. The remaining cases of the argument are summarized in Appendix Table 2.

In each case, there is justified envy since applicant A (who has the highest priority) prefers another applicant's allocation. \square

A.2 Properties of MWP

Proof of Proposition 2. (i) Consider the strategy of applicant $a \in A$. We will show that truthful preference revelation is weakly dominant. Since the order in which allocations are made depends only on the priority ordering and *not* on the applicants' preferences, we restrict our attention to the step in which applicant a has the highest priority among those who are unmatched. In this step, we see that a is assigned her most-preferred unit among those that have not yet been assigned. Therefore there is no incentive to misreport preferences.

(ii) Let the allocation resulting from MWP be given by μ and consider a reallocation μ' . It suffices to show that there is some applicant who prefers the original allocation μ . Denote the set of applicants whose allocations differ across μ and μ' by $A' = \{a : \mu'(a) \neq \mu(a)\}$. Let a_0 denote the applicant who is assigned first among the applicants in A' . For any $a' \in A'$ with $a' \neq a_0$, the unit $\langle b_{a'}, r_{a'} \rangle$ is assigned at a later step than $\langle b_{a_0}, r_{a_0} \rangle$ is assigned. Since MWP is strategy-proof by (i), applicant a_0 strictly prefers the original allocation $(b_{a_0}, t_0 + \tau_{b_{a_0}, t_0}(r_{a_0}))$ over $(b_{a'}, t_0 + \tau_{b_{a'}, t_0}(r_{a'}))$ as desired.

(iii) Suppose applicant a' has higher priority than another applicant a . According to MWP, a' receives an assignment (i.e., a unit or a place on some waiting list) before a does. This implies that $\langle b_a, r_a \rangle$ is available when $\langle b_{a'}, r_{a'} \rangle$ is assigned. By strategy-proofness, we have that a' prefers $\langle b_{a'}, r_{a'} \rangle$ to the unit assigned to a . In other words, a' does not envy the lower priority applicant a . \square

A.3 Uniqueness

For a given history h^τ of arrivals up to time τ , a preference profile \succ over building-time pairs induces a profile of rankings $\succ' = (\succ^t)_t$ over units, where we denote the profile of period- t rankings by \succ^t . Given an arrival history and the rankings over units induced by reported preferences, a mechanism φ specifies an allocation μ_{\succ}^t in each period.

A mechanism φ is *non-bossy* if for any agent a and any pair of preference profiles \succ and $\hat{\succ}$, we have $\mu_{\hat{\succ}_a, \succ_{-a}} = \mu_{\succ}$ whenever $\mu_{\hat{\succ}_a, \succ_{-a}}(a) = \mu_{\succ}(a)$. In a non-bossy mechanism, no agent can change the allocation without changing her own assignment.

Let \succ^τ be a ranking over units in period τ , and let π_τ be any permutation over the set of units. Define $\hat{\succ}^\tau = \pi_\tau \succ^\tau$ and $\hat{\succ}^t = \succ^t$ for $t \neq \tau$. We say that φ satisfies *within-period neutrality* if $\mu_{\hat{\succ}^\tau}^\tau(a) \neq \emptyset$ implies $\mu_{\hat{\succ}^\tau}^\tau(a) = \pi_\tau \mu_{\succ^\tau}^\tau(a)$. Now let π be any permutation over the set of units, and define $\tilde{\succ}^t = \pi \succ^t$. We say that φ satisfies *across-period neutrality* if $\mu_{\tilde{\succ}^t} = \pi \mu_{\succ^t}$ whenever \succ^t is independent of t . A mechanism is *neutral* if it satisfies both within- and across-period neutrality.

For any unit x and a history h^t , let $L_{\succ_a^t}^t(x) = \{y : x \succ_a^t y\}$ denote the lower contour set of x under ranking \succ_a^t . We say that x is *relatively better* under \succ'_a than under \succ_a conditional on history h^t if the lower contour set expands when the preference changes from \succ_a to \succ'_a , i.e., if $L_{\succ_a^t}^t(x) \subseteq L_{\succ'_a^t}^t(x)$.

A mechanism φ is *monotonic* if $L_{\succ_a^t}^t(\mu_{\succ}(a)) \subseteq L_{\hat{\succ}_a^t}^t(\mu_{\succ}(a))$ implies $\mu_{\hat{\succ}}^t = \mu_{\succ}^t$. In other words, if for every agent the mechanism yields under preference profile \succ an allocation in period t which is relatively better under $\hat{\succ}$, then the mechanism yields the same allocation under $\hat{\succ}$.

Proposition 3 states that any strategy-proof, non-bossy, and neutral mechanism is an extended multiple-waitlist procedure. The proof is related to that of an analogous result due to Svensson [1999] for a setting in which objects do not arrive stochastically.

Proof of Proposition 3. Denote the set of all units by X . Fix a history of arrivals, and let X_t denote the set of units that arrive in period t . The proof uses the following lemmas, which will be proven afterward.

Lemma 1. *If φ is strategy-proof and non-bossy, then φ is monotonic.*

Lemma 2. *When all agents' preferences are identical, the allocation under a strategy-proof, non-bossy, and neutral mechanism is ex-ante Pareto efficient.*

First suppose the preferences are such that the induced rankings over units are common (i.e., identical across agents) and persistent (i.e., identical across time): $\succ_a^t = \succ_*$ for all a and all t . By Lemma 2, the assignment μ_{\succ} is efficient. Without loss of generality, label the agents so that $i < j$ whenever (i) $T(a_i) < T(a_j)$, i.e., whenever a_i is matched before a_j ; or (ii) $T(a_i) = T(a_j) =: T$ and $\mu_{\succ}^T(a_i) \succ_* \mu_{\succ}^T(a_j)$, i.e., a_i and a_j are matched in the same period but the unit assigned to a_i is preferable. Since

φ is neutral, for any profile \succ' in which preferences are common and persistent, the order is preserved: $\mu_{\succ'}(a_1) \succ_* \mu_{\succ'}(a_2) \succ_* \dots$.

Now consider an arbitrary preference profile \triangleright . Given a sequence $t(j)$, let y_j be the most-preferred unit among $X \setminus \{y_{j'}\}_{j' < j}$ under ranking $\triangleright_{a_j}^{t(j)}$. Choose the mapping t so that (i) $t(j) \geq t(j-1)$; and (ii) $t(j) > t(j-1)$ only if $\bigcup_{t < t(j)} X_t \subseteq \{y_{j'}\}_{j' < j}$. Define the ranking $\hat{\succ}_*$ by $y_j \hat{\succ}_* y_{j+1}$ for all j , and let $\hat{\succ}'$ be the corresponding profile of common and persistent rankings. By the result above from Lemma 2, we have that $\mu_{\hat{\succ}'}(a_j) = y_j$ for all j .

The above shows that under common and persistent rankings, a non-bossy and strategy-proof allocation mechanism is an extended multiple-waitlist procedure. The conditions on $t(\cdot)$ require that all elements of X_t (and perhaps more units) are assigned in period t . We will proceed to show that by imposing neutrality, this characterization extends to the arbitrary profile of rankings.

Suppose x satisfies $\mu_{\hat{\succ}'}(a_j) \hat{\succ}_* x$. Then by construction $x \in X \setminus \{y_{j'}\}_{j' < j}$, so we have $\mu_{\hat{\succ}'}(a_j) \succeq_{a_j}^{t(j)} x$. Hence, by Lemma 1, we have $\mu_{\triangleright'} = \mu_{\hat{\succ}'}$ as desired. \square

Proof of Lemma 1. Consider a preference profile \succ' such that for each agent a , the allocation $\mu_{\succ}(a)$ is relatively better under \succ' than under \succ . To prove that φ is monotonic, it suffices to show that $\mu_{\succ'} = \mu_{\succ}$.

We start by showing that $\mu_{\succ'_a, \succ_{-a}} = \mu_{\succ}$ for any agent a . Since φ is strategy-proof, agent a reports her preference truthfully, so we have $\mu_{\succ}(a) \succeq_a \mu_{\succ'_a, \succ_{-a}}(a)$. Now since $\mu_{\succ}(a)$ is relatively better under \succ' ,

$$\mu_{\succ}(a) \succeq'_a \mu_{\succ'_a, \succ_{-a}}(a).$$

Again by strategy-proofness, we have

$$\mu_{\succ'_a, \succ_{-a}}(a) \succeq'_a \mu_{\succ}(a).$$

The strategy-proofness condition thus implies that a is indifferent between $\mu_{\succ'_a, \succ_{-a}}$ and μ_{\succ} . From non-bossiness we obtain the result that $\mu_{\succ'_a, \succ_{-a}} = \mu_{\succ}$.

Repeating the argument for each of the remaining agents gives the desired result

that the mechanism yields the same allocation under \succ' . \square

Proof of Lemma 2. Let \succ_* denote the common preference. Without loss of generality, label the agents so that $i < j$ whenever (i) $T(a_i) < T(a_j)$, i.e., whenever a_i is matched before a_j ; or (ii) $T(a_i) = T(a_j) =: T$ and $\mu_\succ^T(a_i) \succ_*^T \mu_\succ^T(a_j)$, i.e., a_i and a_j are matched in the same period but the unit assigned to a_i is preferable.

Suppose on the contrary that μ_\succ is not efficient. This implies that there exists a unit $\xi = \langle \xi_b, \xi_r \rangle$ such that (i) ξ is ranked higher in period t than the unit that a_i is assigned for some i , i.e., $\xi \succ_*^t \mu_\succ(a_i)$, where $t := T(a_i)$; and (ii) ξ is not assigned to any agent in period t , i.e., $\xi \neq \mu_\succ^t(a)$ for any a .

Define $n = \min\{i : \xi \succ_*^t \mu_\succ^t(a_i) \neq \emptyset\}$ to be the agent who receives the most-preferred unit among those which are ranked below ξ but still allocated under μ_\succ . Let $\hat{\succ}_*^t$ be a ranking over units that induces a ranking over units which coincides with that of \succ_*^t except that $\mu_\succ(a_n) \hat{\succ}_*^t \xi$ (i.e., the order of ξ and $\mu_\succ(a_n)$ is switched). Denote by π_t the permutation that switches the rankings of ξ and $\mu_\succ(a_n)$ so that $\hat{\succ}^t = \pi_t \succ_*^t$.

The argument consists of two steps. First we will use Lemma 1 to show that the assignments under \succ and $\hat{\succ}$ must be identical. Next, using the neutrality property, we will obtain a contradiction.

We begin by showing that $\mu_\succ^t(a)$ is relatively better under $\hat{\succ}^t$ than under \succ^t in period t for every a . Let x satisfy $\mu_\succ(a) \succeq_* x$. If $x = \mu_\succ(a_n)$ and $a = a_n$, then trivially we have $\mu_\succ(a_i) \hat{\succeq}_* x$, so suppose otherwise. By construction, note that (i) \succ^t and $\hat{\succ}^t$ produce identical rankings over the units excluding ξ and $\mu_\succ(a_n)$, and (ii) $\mu_\succ(a)$ is either ranked higher than both ξ and $\mu_\succ(a_n)$ or neither of them. This implies $\mu_\succ(a) \hat{\succeq}_*^t x$. Since φ is monotonic by Lemma 1, we have that $\mu_\succ = \mu_{\hat{\succ}}$.

Now since φ satisfies within-period neutrality, we have $\mu_{\hat{\succ}}^t(a_n) = \xi$, which contradicts the result from Lemma 1. Therefore we conclude that μ_\succ is efficient. \square

A.4 Acyclicity and Generalized MWP

Proof of Proposition 4. (i) We will show that truthful preference revelation is weakly dominant for each applicant $a \in A$. Since the priority ordering is acyclic, there is

at most one applicant \hat{a} who has higher priority than a at some buildings but lower priority at the others. If there is no higher-priority applicant, then applicant a receives her most-preferred unit (among those that are available) when she reaches the top of the centralized waiting list. Otherwise, if there is a higher-priority applicant, then a receives her most-preferred unit unless \hat{a} prefers the same unit and has higher priority for the building. In that case, a is assigned her most-preferred unit among those that remain after \hat{a} selects a unit. Since the event that a receives her most-preferred unit among those that are available depends only on the priorities and the other applicants' stated preferences, there is no incentive for a to misreport preferences.

(ii) For any reallocation μ' , it suffices to show that there is some applicant who prefers the original allocation μ resulting from the generalized MWP. Let a_0 denote the applicant who is offered a unit first among the set of applicants $A' = \{a : \mu'(a) \neq \mu(a)\}$ whose allocations differ across μ and μ' . There is at most one applicant \hat{a}_0 who has higher priority than a_0 at some buildings but lower priority at the others. If such \hat{a}_0 does not exist, or if a prefers a unit in a building at which \hat{a}_0 does not have higher priority, then the proof proceeds as in Proposition 2. Otherwise, \hat{a}_0 receives her most-preferred unit among those that are available at the time of assignment, which implies that \hat{a}_0 prefers the original allocation $(b_{\hat{a}_0}, t_0 + \tau_{b_{\hat{a}_0}, t_0}(r_{\hat{a}_0}))$ over $(b_{a'}, t_0 + \tau_{b_{a'}, t_0}(r_{a'}))$ for any $a' \in A'$ as desired.

(iii) Suppose applicant a' has higher priority than another applicant a at the building b_a (where a is assigned), and consider the time at which a' receives an assignment. If a' prefers a unit in a building at which she has the highest priority, then the proof proceeds as in Proposition 2: a' receives an assignment under generalized MWP before a does, so $\langle b_a, r_a \rangle$ is available when $\langle b_{a'}, r_{a'} \rangle$ is assigned, which means that a' prefers $\langle b_{a'}, r_{a'} \rangle$ to the unit assigned to a . Now suppose that a' prefers a unit in a building at which another applicant \hat{a}' has higher priority.³⁶ In the case that $\hat{a}' \neq a$, the argument is the same as before because a' receives an assignment before a does. Otherwise, we have $\hat{a}' = a$, i.e., that a has higher priority than a' at some building. Since a' prefers a unit in a building at which a has higher priority, a receives an assignment before a' does. However, $\langle b_a, r_a \rangle$ was available to a'

³⁶As noted earlier, acyclicity implies that there is at most one such applicant.

(since a' has higher priority at b_a by assumption) but not chosen, which implies that a' prefers $\langle b_{a'}, r_{a'} \rangle$. In all cases, a' does not envy the lower priority applicant a . \square

Proof of Proposition 5. Proposition 4 implies that the acyclicity condition is sufficient for the existence of an allocation mechanism that is strategy-proof, efficient, and eliminates justified envy. We now show that if the priority orderings violate the acyclicity condition, then there does not exist a strategy-proof allocation mechanism that satisfies both efficiency and the elimination of justified envy.

If \succ_B does not satisfy the acyclicity condition, then there exist buildings α and β such that the priority orderings form a cycle; that is, there exist applicants A, B, and C such that:

$$\begin{aligned} A \succ_{\alpha} B \succ_{\alpha} C, \\ \text{and } C \succ_{\beta} A. \end{aligned}$$

Let the applicants' preferences be as given in Appendix Table 3. Consider the following deterministic arrival process: from building α , a unit becomes available in period 0 and another unit becomes available in period 2; from building β , a unit becomes available in period 1.

Suppose applicant A is assigned building α in period 0. If B is assigned β in period 1, then the allocation is inefficient because A and B prefer to switch. Likewise, if C is assigned β in period 1, then the allocation is inefficient because A and C prefer to switch.

The analysis is similar if applicant B or applicant C is assigned building α in period 0. Regardless of which applicant is assigned building α in period 0, either there is a Pareto improvement or some applicant justifiably envies another. The remaining cases of the argument are summarized in Appendix Table 4.

This demonstrates the necessity of the acyclicity condition. \square

A.5 Stochastic mechanisms

Proof of Proposition 6. Necessity can be demonstrated using the same preferences and arrival process as in the proof of Proposition 5. Assume \succ_B does not satisfy acyclicity so that there exist buildings α and β , and applicants A, B, and C such that $A \succ_\alpha B \succ_\alpha C$ and $C \succ_\beta A$. The applicants' preferences are given in Appendix Table 3. As before, the arrival process is deterministic: a unit in building α arrives in period 0; a unit in building β arrives in period 1; and another unit in building α arrives in period 2.

We begin by determining the assignment of applicant A. If A is assigned a positive probability of $(\alpha, 2)$, then A would justifiably lottery-envy any applicant who is assigned a positive probability of $(\alpha, 0)$. If A is assigned a positive probability of $(\beta, 1)$, then C must assigned positive probability of either $(\alpha, 0)$ or $(\alpha, 2)$: in the former case, B has justified lottery-envy towards C; and in the latter case, C has justified lottery-envy towards A. This leaves us with the conclusion that A must be assigned $(\alpha, 0)$ with certainty.

Now either B is assigned $(\alpha, 2)$ with certainty or B is assigned a positive probability of $(\beta, 1)$: in the former case, the assignment is inefficient since A and C would prefer to switch; in the latter case, B has justified lottery-envy towards C.

This establishes the necessity of acyclicity. Sufficiency follows from the same construction as in Proposition 4. \square

A.6 PHA- k mechanisms

Proof of Proposition 7. Consider the following example with 2 buildings, $B = \{b_0, b_1\}$, each with one unit; $k + 2$ applicants, $A = \{a_i\}_{i=0}^{k+1}$; and $k + 2$ periods.

Assume that the buildings' common priority list ranks applicant a_i higher than applicant a_{i+1} for all i .

Units arrive deterministically: in periods 0 through k , a unit in building b_1 becomes available with certainty; in period $k + 1$, a unit in building b_0 becomes available with certainty.

Applicant a_0 prefers the unit in building b_0 that arrives in period $k + 1$, and all other applicants prefer units that arrive earlier.

Since applicant a_0 cannot attain any unit outside building b_1 , she refuses no offers and receives a unit in building b_1 immediately. Applicants a_1 through a_{k+1} also do not refuse any offers, as they prefer units that arrive earlier.

Applicant a_0 prefers the unit in building b_1 and therefore justifiably envies a_{k+1} . Matching applicant a_0 with the unit in building b_0 and matching each of the remaining applicants with the room that arrives one period earlier would be a Pareto improvement. \square

Proof of Proposition 7. Consider the following example with $k + 2$ buildings, $B = \{b_i\}_{i=0}^{k+1}$, each with one unit; $k + 2$ applicants, $A = \{a_i\}_{i=0}^{k+1}$; and $k + 2$ periods.

Assume that the buildings' common priority list ranks applicant a_i higher than applicant a_{i+1} for all i . The applicants' preferences are given in Appendix Table 5.

Units arrive deterministically: in period i , the unit in building i becomes available with certainty.

Applicant a_0 refuses k offers and receives a unit in building b_k , since all units that become available sooner are less desirable. For $i = 1, \dots, k - 1$, applicant a_i refuses $k - i$ offers; her first-choice unit (in building $k - i + 1$) will be taken by applicant a_{i-1} , who has higher priority, so applicant a_i receives her second-choice unit (in building b_{k-i}). Likewise, applicant a_k accepts the offer of a unit in building b_0 , since a_{k-1} will accept the unit in building 1. This leaves the unit in building b_{k+1} for applicant a_{k+1} . The allocation procedure concludes with each applicant receiving her second-choice unit.

Since it is possible to redistribute the units so that each applicant receives her most-preferred unit, the assignment is inefficient. Furthermore, the procedure fails to eliminate justified envy since applicant a_0 has the highest priority but prefers the unit assigned to applicant a_{k+1} . \square

Figure 1: Multiple-Waitlist Procedure

- Step 0** Proceed to the next period.
- Step 1** If there are no more available units this period: return to Step 0. Otherwise: choose a unit randomly from the set of units that are available this period.
- Step 2** If the building in which the unit belongs has a non-empty FIFO queue: allocate the unit to the applicant at the top of the queue, and return to Step 1. Otherwise: offer the unit to the applicant at the top of the centralized waiting list.
- Step 3** If the applicant accepts: allocate the unit to the applicant, and return to Step 1. Otherwise: the applicant chooses a different building and joins the associated FIFO queue to receive the next available unit in that building; return to Step 1.

Figure 2: Generalized Multiple-Waitlist Procedure

- Step 0** Proceed to the next period.
- Step 1** If there are no more available units this period: return to Step 0. Otherwise: choose a unit randomly from the set of units that are available this period.
- Step 2** If the building in which the unit belongs has a non-empty FIFO queue: allocate the unit to the applicant at the top of the queue, and return to Step 1. Otherwise: offer the unit to the applicant (among those who have not yet been assigned a unit) with the highest priority for that building.
- Step 3** If the applicant accepts: allocate the unit to the applicant, and return to Step 1. Otherwise: the applicant specifies a different building and requests to join the associated FIFO queue.
- Step 4** If another applicant has the highest priority at the specified building: move that applicant to the top of the priority list for every building, and return to Step 2. Otherwise: the applicant joins the specified FIFO queue to receive the next available unit in that building; return to Step 1.

Table 1: Distribution of allocation procedures by size of housing agency

	Small	Medium	Large	Extra-large	Total
PHAs making 1 or 2 offers	430	114	71	3	618
PHAs making more than 2 offers	446	96	40	7	589
	Small	Medium	Large	Extra-large	Total
Centralized waiting list	929	179	117	9	1,234
Non-centralized waiting list	293	70	14	7	384

Note: Data are from the Division of Program Monitoring and Research, US Department of Housing and Urban Development, 1998. The top frame shows the relationship between housing-agency size and the maximum number of offers an applicant can receive before being removed from the waiting list. The bottom frame shows the relationship between housing-agency size and waiting-list method. Small: between 100 and 500 units. Medium: between 500 and 1,250 units. Large: between 1,250 and 6,600 units. Extra-large: 6,600 units or more. From the universe of over 3,100 housing agencies, those that operate fewer than 100 units are excluded, leaving a set of agencies that accounts for 94 percent of all public-housing units.

Table 2: Arrivals by building type

Building category	Number of buildings	Number of arrivals
Family large	12	677
Family medium	6	144
Family small	1	24
Mixed	4	300
Senior large	3	59
Senior small	8	191

Note: This table shows the number of units that arrive for each building category from June 2001 to June 2006. Data on the number of arrivals are from the HACP, as reported in Table 3 of [Geyer and Sieg \[2013\]](#).

Table 3: Parameter estimates for public-housing preferences

Parameter	Mean	Standard error
Income	0.329	(0.028)
Moving cost	3.186	(0.017)
Demographics		
Nonwhite, nonsenior	1.222	(0.071)
White, senior	0.209	(0.113)
Nonwhite, senior	1.000	(0.101)
Children	-0.315	(0.123)
Female	0.053	(0.061)
Female, senior	-0.174	(0.094)
Female, children	0.426	(0.130)
Fixed effects		
Family large	4.217	(0.254)
Family medium	4.848	(0.261)
Family small	4.604	(0.277)
Mixed	4.394	(0.260)
Senior large	4.626	(0.263)
Senior small	4.907	(0.258)

Note: This table reports estimates of the parameters from the utility function in equations (1) and (2). Parameter estimates are from the model with supply-side restrictions in Table 10 of [Geyer and Sieg \[2013\]](#), estimated using household-level data from the HACP and the SIPP.

Table 4: Welfare gain relative to PHA-1 mechanism

Mechanism	Mean \overline{EV}	95% Confidence Interval
PHA-2	\$3,866	[\$3,589, \$4,158]
MWP	\$6,429	[\$6,153, \$6,705]
PHA-max	\$6,779	[\$6,503, \$7,053]
Ex-post optimum	\$8,505	[\$8,212, \$8,800]

Note: This table contains the results of 100 counterfactual simulations. In each simulation, we compute the average across all applicants of the equivalent variation (\overline{EV}) of changing the allocation mechanism from PHA-1. The second column reports the mean of present-discounted values (with an annual discount factor of 0.6). The final column provides lower and upper bounds of the bootstrapped 95-percent confidence interval (based on 10,000 replications). PHA-max denotes an upper bound for welfare under PHA- ∞ in which applicants know the realization of the arrival process in advance and have complete information about the preferences of other households. The bottom row displays the maximum possible gain that could be achieved by a social planner that knows the complete realization of the arrival process in advance.

Appendix Table 1: Preferences for applicants A, B, and C in Proposition 1

γ_A	γ_B	γ_C
$(\beta, 0)$	$(\alpha, 0)$	$(\alpha, 0)$
$(\beta, 1)$	$(\alpha, 1)$	$(\alpha, 1)$
$(\gamma, 0)$	$(\gamma, 0)$	$(\gamma, 0)$
$(\alpha, 0)$	$(\alpha, 2)$	$(\alpha, 2)$
$(\beta, 2)$	$(\beta, 0)$	$(\beta, 0)$
$(\gamma, 1)$	$(\gamma, 1)$	$(\gamma, 1)$
$(\alpha, 1)$	$(\beta, 1)$	$(\beta, 1)$
$(\gamma, 2)$	$(\gamma, 2)$	$(\gamma, 2)$
$(\alpha, 2)$	$(\beta, 2)$	$(\beta, 2)$

Note: Preferences for applicants A, B, and C listed in order from most-preferred to least-preferred. Agent A's preference can be generated by the utility function $u_A(b, t) = f(b) - 3t$, where $f(\alpha) = 1$, $f(\beta) = 6$, $f(\gamma) = 2$; applicant B's and applicant C's preference can be generated by the utility function $u_B(b, t) = u_C(b, t) = g(b) - 3t$, where $g(\alpha) = 8$, $g(\beta) = 1$, $g(\gamma) = 3$.

Appendix Table 2: Example used in proof of Proposition 1

Period 0		Period 1		Period 2		Switch	Envy
α	A	β	B	γ	C	A, B	A, B
α	A	β	C	γ	B	A, C	A, C
α	B	γ	A	β	C	A, C	A, C
α	B	γ	C	α	A	A, C	A, C
α	C	γ	A	α	B	A, B	A, B
α	C	γ	B	β	A	A, B	A, B

Note: An example in which any assignment rule can lead to a violation of ex-post efficiency and ex-post elimination of justified envy. Buildings are denoted α , β , and γ . Agents are denoted A, B, and C, and their preferences are given in Appendix Table 1. Each building gives applicant A the highest priority. The penultimate column specifies which applicants could trade to obtain a Pareto improvement in each setting, and the last column specifies whether some applicant justifiably envies another.

Appendix Table 3: Preferences for applicants A, B, and C in Proposition 5

\succ_A	\succ_B	\succ_C
$(\beta, 0)$	$(\alpha, 0)$	$(\beta, 0)$
$(\beta, 1)$	$(\alpha, 1)$	$(\alpha, 0)$
$(\alpha, 0)$	$(\beta, 0)$	$(\beta, 1)$
$(\beta, 2)$	$(\alpha, 2)$	$(\alpha, 1)$
$(\alpha, 1)$	$(\beta, 1)$	$(\beta, 2)$
$(\alpha, 2)$	$(\beta, 2)$	$(\alpha, 2)$

Note: Preferences for applicants A, B, and C listed in order from most-preferred to least-preferred. Agent A's preference can be generated by the utility function $u_A(b, t) = f(b) - 2t$, where $f(\alpha) = 1$ and $f(\beta) = 4$; applicant B's preference can be generated by the utility function $u_B(b, t) = g(b) - 2t$, where $g(\alpha) = 4$ and $g(\beta) = 1$; applicant C's preference can be generated by the utility function $u_C(b, t) = h(b) - 2t$, where $h(\alpha) = 1$ and $h(\beta) = 2$.

Appendix Table 4: Example used in necessity proof of Proposition 5

Period 0		Period 1		Period 2		Switch	Envy
α	A	β	B	α	C	A, B	—
α	A	β	C	α	B	A, C	—
α	B	β	A	α	C	—	C, A
α	B	β	C	α	A	—	A, B
α	C	β	A	α	B	—	B, C
α	C	β	B	α	A	—	A, C

Note: An example in which no assignment rule can satisfy Pareto efficiency and the elimination of justified envy when priority orderings violate acyclicity. Buildings are denoted α and β . Agents are denoted A, B, and C, and their preferences are given in Appendix Table 3. Building α ranks applicant A above C, with B in between; but building β ranks C above A. The penultimate column specifies which applicants could trade to obtain a Pareto improvement in each setting, and the last column specifies whether some applicant justifiably envies another.

Appendix Table 5: Preferences for each applicant a_i in Proposition 7

\succ_{a_0}	\succ_{a_1}	\cdots	$\succ_{a_{k-1}}$	\succ_{a_k}	$\succ_{a_{k+1}}$
$(b_{k+1}, k+1)$	(b_k, k)	\cdots	$(b_2, 2)$	$(b_1, 1)$	$(b_0, 0)$
(b_k, k)	$(b_{k-1}, k-1)$	\cdots	$(b_1, 1)$	$(b_0, 0)$	$(b_{k+1}, k+1)$
$(b_{k-1}, k-1)$	$(b_{k-2}, k-2)$	\cdots	$(b_0, 0)$	$(b_{k+1}, k+1)$	(b_k, k)
\vdots	\vdots	\ddots	\vdots	\vdots	\vdots
$(b_2, 2)$	$(b_1, 1)$	\cdots	$(b_5, 5)$	$(b_4, 4)$	$(b_3, 3)$
$(b_1, 1)$	$(b_0, 0)$	\cdots	$(b_4, 4)$	$(b_3, 3)$	$(b_2, 2)$
$(b_0, 0)$	$(b_{k+1}, k+1)$	\cdots	$(b_3, 3)$	$(b_2, 2)$	$(b_1, 1)$

Note: Preferences for applicants $\{a_i\}_{i=0}^{k+1}$ listed in order from most-preferred to least-preferred. The most-preferred unit for applicant a_i is in building b_{k-i+1} which arrives in period $k-i+1$. Each applicant a_i prefers the unit in building b_r over the unit in building b_{r-1} for all $r \not\equiv -i \pmod{k+2}$.