Matching with Stochastic Arrival

Neil Thakral*

20 January 2015

Abstract

We study matching in a dynamic setting, with applications to the allocation of public housing. In our model, objects of different types that arrive stochastically over time must be allocated to agents in a queue. For the case that the objects share a common priority ordering over agents, we introduce a strategy-proof mechanism that satisfies certain fairness and efficiency properties. More generally, we show that the mechanism continues to satisfy these properties if and only if the priority relations satisfy an acyclicity condition. We then turn to an application of the framework by evaluating the procedures that are currently being used to allocate public housing. The estimated welfare gains from adopting the new mechanism are substantial, exceeding $5,000 per applicant.

1 Introduction

A typical matching model consists of a set of objects, a list of capacities, a set of agents, and a list of preferences. This framework applies to problems in which units of each object are allocated to agents based on a priority ordering. The objective is to design

*E-mail: `nthakral@fas.harvard.edu`. I am grateful to Kfir Eliaz for his constant support and encouragement. I would also like to thank Itay P. Fainmesser, Scott Duke Kominers, David Laibson, Xiaosheng Mu, Parag Pathak, Tayfun Sönmez, and especially Linh T. Tô for helpful comments and discussion.

allocation procedures that satisfy desirable properties such as strategy-proofness, efficiency, and stability. For example, the deferred acceptance algorithm introduced by Gale and Shapley (1962) results in a matching that is stable.

While these procedures may be appropriate for settings in which units are allocated simultaneously, various real-world problems are dynamic in the sense that units arrive stochastically. Consider the example of a company with offices in different parts of the world and relocation opportunities that arise occasionally. For a job opening at a given location, the company may have preferences over workers based on seniority, work ability, language skills, etc. Workers have preferences over job locations as well as the amount of time they spend waiting for a position. A worker may, for instance, settle on an offer from a less preferred location if she believes that she would have to wait for too long before the more preferred location becomes available to her.[1]

As another example, consider the problem of allocating public housing to applicants on a waiting list. The public housing agency is in charge of allocating rooms in apartment buildings that vary by location. A room becomes available when the former tenants vacate the property. Agents have heterogeneous preferences over apartment buildings, as they may prefer to live closer to their respective workplaces. Waiting is costly, so each applicant prefers to receive a housing assignment earlier rather than later. The public housing agency ranks applicants on the waiting list by priority, and in some cases, priorities may differ across apartment buildings. For example, some buildings may only be available to high-priority applicants such as the elderly, disabled, homeless, victims of natural disasters, or victims of domestic abuse; and buildings that have elevators may give even higher priority to the disabled or elderly.

The following aspects of these situations are not accounted for by the standard matching framework described earlier: (i) objects are allocated dynamically as they arrive over time, (ii) there is uncertainty about the availability of the objects, and (iii) applicants have preferences over waiting times. We introduce a model that incorporates these features. Our aim is to design dynamic allocation mechanisms

[1] Hylland and Zeckhauser (1979) analyze a static version of the "job assignment problem" discussed in this example.

that are strategy-proof (i.e., not subject to strategic manipulation) and to explore whether these mechanisms can satisfy additional fairness and efficiency properties. We find that the appropriate perspective to take in this dynamic setting is to evaluate whether these properties can be satisfied *ex-ante*. Whereas the standard framework is relevant in settings where objects have static capacities, our model consists of objects that are characterized by stochastic arrival of units.

When such situations arise in practice, the mechanisms that are employed typically fail to satisfy desirable properties. Again consider the example of public-housing allocation. Many public housing agencies use a variation of the following procedure to allocate rooms. Eligible agents submit an application which includes information that is used to determine priorities. The applicant with the highest priority is assigned the first room that becomes available.[2] A simple example with two buildings, two applicants, and two periods demonstrates that this take-it-or-leave-it procedure can lead to allocations that are inefficient and unstable. Suppose that applicant 1 strictly prefers to wait for b_2 and that applicant 2 prefers b_1 to b_2. Assume that applicant 1 has high priority and that in period $i \in \{1, 2\}$ a unit in building b_i becomes available with certainty. The take-it-or-leave-it mechanism assigns b_1 to applicant 1 and b_2 to applicant 2. Since the high priority applicant prefers the assignment of the low priority applicant, the mechanism is unfair (i.e., it fails to eliminate justified envy). Additionally, the mechanism is inefficient since applicant 2 also prefers applicant 1's allocation.

We propose an alternative allocation mechanism — the Multiple Waitlist Procedure (MWP) — that essentially modifies the take-it-or-leave-it mechanism by giving applicants the opportunity to decline an offer and opt to be placed on a First-In/First-Out (FIFO) waiting list for a different object. We compare the new allocation mechanism with the procedures currently used by public housing agencies to allocate rooms to applicants. When objects share a common priority ordering over applicants, MWP is strategy-proof, efficient, and fair; existing housing allocation mechanisms,

[2]If she refuses, then she does not receive any allocation and her application is withdrawn or she is moved to the bottom of the waiting list. Such mechanisms are used by public housing agencies in Houston, TX and Providence, RI, for example.

however, fail to satisfy these desirable properties. This motivates our empirical investigation of public-housing allocation, which uses estimated preferences to quantify welfare gains from adopting our proposed mechanism.

Although our empirical analysis focuses on the case of common priority orderings, applications of our framework (including public-housing allocation) may not satisfy this restriction. In the general case that priority orderings are heterogeneous across object types, we provide a necessary and sufficient condition for the existence of a strategy-proof mechanism that satisfies the fairness and efficiency properties. If the priority relations satisfy an *acyclicity* condition, then a generalized version of MWP continues to satisfy strategy-proofness, efficiency, and the elimination of justified envy. Conversely, if the priority relations fail to satisfy the *acyclicity* condition, then there does not exist a strategy-proof mechanism that satisfies these properties simultaneously. We show that these results continue to hold for a more general class of allocation mechanisms, namely lottery mechanisms.

Due to the tension between efficiency and the elimination of justified envy in the absense of acyclicity, we suggest strategy-proof mechanisms that satisfy each of these properties separately. A simple extension of MWP that ignores the objects' priority orderings retains the strategy-proofness and efficiency properties. By modifying MWP so that the waiting list for a given object respects the object's priority ordering (instead of the FIFO property), we obtain a mechanism that satisfies strategy-proofness and the elimination of justified envy.

A central feature of our matching model is that objects become available over time.[3] Related work by Doval (2014) explores dynamically stable matchings in two-sided markets with deterministic arrivals on one side of the market; Dimakopoulos and Heller (2014) consider a model of matching with contracts in which object capacities are deterministically time-dependent; and Gershkov and Moldovanu (2009b) investigate an allocation setting with deterministic arrival of agents. The present paper, by contrast, studies allocation problems in which objects arrive stochastically.

[3]Other models in which dynamic considerations arise due to arrival over time include Ünver (2010), in which each agent arrives with an object to trade; Gershkov et al. (2014), in which agents choose when to make themselves available for trade; and Akbarpour et al. (2014), in which agents arrive and depart stochastically in a networked market.

Leshno (2012) also attempts to incorporate this type of uncertainty in arrival by introducing a model in which an object randomly drawn from one of two types becomes available in each period and must be allocated to an agent on the waiting list. In his model, the goal of the social planner is simply to maximize the fraction of agents who are matched with their most-preferred type of object, irrespective of any particular individual's waiting time. This is justified by a homogeneity assumption: agents are identical in terms of waiting costs as well as values for their most-preferred and least-preferred objects. By contrast, our analysis allows for heterogeneity in preferences over not only object types but also the amount of time spent waiting for an allocation. An advantage of our approach is that, by enriching the domain over which preferences are defined to include the time dimension and by explicitly addressing priorities in the waiting list, we are able to explore the concepts of strategy-proofness and stability, neither of which is considered in Leshno (2012). Furthermore, while Leshno (2012) assumes that a unit becomes available each period with a fixed probability that the unit is of a given type, we make no assumptions about the underlying stochastic process that governs the arrival rate of units.

Some recent papers study matching problems that are dynamic in a different sense. Specifically, their dynamics arise from manipulable priorities (agents can affect the priorities by acting strategically) or reallocation (the same set of objects is allocated among the agents in multiple periods). Often motivated by specific real-world institutions, both of these considerations are generally present in matching models with overlapping generations.[4] These considerations are also relevant in Abdulkadiroğlu and Loerscher (2007), which studies the problem of allocating a continuum of homogeneous goods among a set of agents in one period and reallocating the same goods among the agents in the second period; in their model, priorities in the second period are higher for agents who opt out in the first period. Damiano and Lam (2005) analyze stability in repeated matching markets, and Kurino (2009)

[4]Examples include Dur (2012), which reformulates the school choice problem; Pereyra (2013), which considers the allocation of teaching positions in Mexico; Kennes et al. (2014a), which studies day care assignment in Denmark; Kurino (2014), which introduces a dynamic house allocation problem in the context of on-campus housing for college students; and Kennes et al. (2014b), which generally analyzes such overlapping generations models in a large market setting.

generalizes this to a setting in which preferences can change with time. Kadam and Kotowski (2014) study a model in which matches last for multiple periods and can be revised over time, allowing for the possibility of intertemporal complementarities in preferences. In contrast to these approaches, dynamics in our model arise neither from strategic actions nor from reallocation; we instead consider a situation that is dynamic in that there is uncertainty about the objects' availability.

A growing literature in market design uses simulations for welfare analysis, though much of this work focuses on the school choice problem.[5] Along with Abdulkadiroğlu et al. (2014) and Agarwal and Somaini (2014) (who study school choice in New York City and Cambridge, respectively) our work is among the first to use preferences estimated using data from real-world assignment procedures to quantify welfare gains due to adopting alternative mechanisms. Our counterfactual simulations suggest that changing existing public-housing allocation mechanisms to MWP would lead to welfare gains of about $6,400 for each applicant who is assigned public housing.

Glaeser and Luttmer (2003) and Wang (2011) find evidence of the misallocation of private housing under rent control. The present paper complements these approaches by exploring the design aspect of public-housing allocation. Empirical work related to public-housing allocation mechanisms is more limited.[6] Van Ommeren and Van der Vlist (2014) addresses questions related to the efficiency of certain queueing procedures used for public-housing allocation in Amsterdam by using estimates of marginal willingness to pay. 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.

The paper is organized as follows. Section 2 describes the dynamic matching problem with a discussion of *ex-ante* and *ex-post* properties. Section 3 introduces a mechanism that satisfies various desirable properties, including strategy-proofness, efficiency, and the elimination of justified envy; explores the possibility of designing allocation procedures that satisfy these properties for arbitrary priority orderings,

[5]See, for example, Erdil and Ergin (2008), Dur (2011), Abdulkadiroğlu et al. (2012), Hafalir et al. (2013), Morrill (2013), and Kesten and Ünver (forthcoming).

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

with an emphasis on the role of acyclicity; and extends the analysis to a more general class of allocation mechanisms. Section 4 compares the mechanism we propose with alternative suggestions and provides an application to public-housing allocation by using estimates from a structural model to evaluate welfare. Section 5 concludes.

2 Model

A dynamic matching problem is a five-tuple $\langle A, B, \succ_B, \succ_A, \pi\rangle$, where A is a finite set of agents (applicants), B is a finite set of objects (buildings); $\succ_B = (\succ_b)_{b\in B}$ is a profile of strict priority relations for the buildings in each period; $\succ_A = (\succ_a)_{a\in A}$ is a profile of applicants' preference relations over building-time pairs $B \times \mathbb{R}_+$; and $\pi = (\pi_{b,\tau}(\cdot \mid h^t))_{b\in B, t, \tau\in\mathbb{N}}$ is an arrival process that specifies, conditional on the history h^t, a probability distribution over the number of units (rooms) in building b that arrive at time τ.[7] A *history* is a map $h^t \colon B \times \{\tau \in \mathbb{N} : \tau \leq t\} \to \mathbb{N}$ that specifies the number of rooms 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 rooms but will find it convenient to denote by $X_b(r \mid h^t)$ the expected waiting time of the r^{th} room in building b to become available, conditional on the history h^t.

An *allocation in period* t is a map $\mu^t \colon A \to (B \times \mathbb{N}) \cup \{\emptyset\}$, and an *allocation* $\mu = (\mu^1, \mu^2, \dots)$ is a collection of allocations in each period. We write $\mu(a) = \langle b, r\rangle$ and interpret this as applicant a being assigned the r^{th} room that becomes available in building b.[8] A room cannot be unmatched after the period in which it becomes available.[9] Let $\eta^t(\cdot)$ summarize the history of the allocations at time t as follows: the number of rooms in building b that have already been assigned is given by $\eta^t(b) = \sum_a \mathbf{1}_{\{\mu^t(a)=\langle b,\cdot\rangle\}}$. Agent a can be thought of as being on a waiting list for

[7] We will typically use the terms "buildings," "rooms," and "applicants" to refer to the generic "objects," "units," and "agents." Applicants do not have preferences over rooms in the same building.

[8] Rooms are not reassigned: if $\mu^t(a) \neq \emptyset$, then $\mu^{t+1}(a) = \mu^t(a)$.

[9] The assumption that units must be allocated upon arrival and cannot be reallocated in the future mirrors the corresponding assumptions in Gershkov and Moldovanu (2009b) and Gershkov and Moldovanu (2009a) for the case of agents arriving stochastically.

building b if the r^{th} room is not yet available: given the history h^t of the arrival process, the room is expected to arrive in $X_b(r \mid h^t)$ periods.[10]

We assume that preferences satisfy *dynamic consistency* and *costly waiting* conditions.[11] These assumptions ensure that the preference relation on $B \times \mathbb{R}_+$ can constructed by eliciting an element of $\mathbb{R}_+^{|B|}$. For each building b, this representation encodes the maximal number of periods that the applicant is willing to wait before receiving a room in her most-preferred building rather than receiving a room in building b immediately.

In addition, we assume that applicants are *risk-neutral*.[12] 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 room in building b in period t over a room in building b' in period t'. However, an allocation $\mu^t(a)$ consists not only of a period t and a room r in building b but also an expected waiting time $X_b(r \mid h^t)$. Our assumptions guarantee that applicants evaluate assignments based on expected waiting time in the following way: $\mu^t(a) = \langle b, r\rangle$ is preferred to $(\mu')^{t'}(a) = \langle b', r'\rangle$ if and only if $(b, t + X_b(r \mid h^t)) \succ_a \left(b', t' + X_b\left(r' \,\middle|\, h^{t'}\right)\right)$.

A *mechanism* φ is a procedure that uses priority orderings, reported preferences, and the history to choose an allocation 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 θ_a.

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 that strategy-proofness is desirable because of fairness (agents who lack information or sophistication are not at a disadvantage), simplicity (agents can easily

[10]If $X_b(r \mid h^t) = 0$, then this 'waiting list' is degenerate, so the applicant receives the room immediately in period t.

[11]Preferences are dynamically consistent if $(b,t) \succ_a (b',t')$ implies $(b, t+\tau) \succ_a (b', t'+\tau)$ for every $\tau > 0$. Preferences satisfy costly waiting if $(b,t) \succ_a (b,t')$ is equivalent to $t < t'$. The assumptions are realistic for applications such as public-housing allocation but can be relaxed for our results.

[12]Our results extend to the case that the social planner knows the applicants' attitudes towards risk. However, we maintain the assumption of risk-neutrality throughout for simplicity.

understand the strategies and the equilibrium), and robustness (the equilibrium does not depend on beliefs about other agents' preferences or information).[13] Another justification for strategy-proofness in our dynamic setting is that the social planner (though not modeled here) may make costly investments based on reported preferences; for example, a public housing agency may use reported preferences to determine where to construct a new building.[14]

Next we define a property that can be interpreted as a form of fairness. An allocation μ *eliminates justified envy* if an applicant who prefers an alternate assignment does not have higher priority than the applicant to whom the other room is assigned.[15] In our dynamic setting, whether an applicant prefers an alternate assignment depends on the timing of the allocation and the information available at the time. We say that an applicant a' *envies* another applicant a if (i) a is assigned a room before a', and a' prefers the assignment of a; or (ii) a is assigned a room after a', and given the information available at the time when a' was matched, a' would have preferred the room allocated to a.

Definition 1 (elimination of justified envy)**.** Let $t^* = \min\{t, t'\}$. If a is assigned $\langle b_a, r_a\rangle$ in period t and a' is assigned $\langle b_{a'}, r_{a'}\rangle$ in period t', then

$$\big(b_a, t^* + X_{b_a}\big(r_a \,\big|\, h^{t^*}\big)\big) \succ_{a'} \big(b_{a'}, t' + X_{b_{a'}}\big(r_{a'} \,\big|\, h^{t'}\big)\big) \implies a \succ_{b_a} a'.$$

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

[13] See Azevedo and Budish (2013) and the references therein for further discussion.

[14] Abdulkadiroğlu et al. (2009) point out that one of the reasons for the closing of an unpopular New York City high school in 2006 was lack of demand as determined by reported student preferences under a strategy-proof mechanism.

[15] Balinski and Sönmez (1999) introduce this property the context of an allocation problem with priorities, namely the student placement model, as an analogue of stability. Indeed if we interpret the priority orderings as the buildings' "preferences," then the property requires that no applicant and building would strictly prefer to be matched with each other rather than their respective matches from the mechanism.

An allocation μ is *ex-ante efficient* if any reallocation μ' makes some agent strictly worse off *ex-ante*. We refer to this as an *ex-ante* notion of efficiency because agents only take into account information that is available when they are matched and are evaluating their expected (rather than realized) arrival times.

Definition 2 (efficiency)**.** For any feasible allocation $\mu' \neq \mu$, there exists some a (who is assigned $\langle b_a, r_a \rangle$ in period t under μ and is assigned $\langle b'_a, r'_a \rangle$ in period t' under μ') such that $\left(b_a, t + X_{b_a}\left(r_a \,\middle|\, h^t\right)\right) \succ_a \left(b'_a, t^* + X_{b'_a}\left(r'_a \,\middle|\, h^{t^*}\right)\right)$, where $t^* = \min\{t, t'\}$.

An alternative notion of efficiency would be *ex-post* efficiency: any way of redistributing rooms that have already arrived (among the agents to whom they are assigned) would make at least one agent strictly worse off.

We say that a mechanism satisfies a given property if the allocation resulting from the equilibrium of the induced preference revelation game satisfies that property. Our analysis focuses on *ex-ante* properties, as motivated by the following result.

Proposition 1. *There does not exist a mechanism that satisfies ex-post efficiency or ex-post elimination of justified envy.*

Proof. We will provide an example to show that it is not possible to design a mechanism that guarantees ex-post efficiency or ex-post elimination of justified envy because the realization of the arrival process may be such that neither can possibly hold.

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 Table 1.

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

Table 1: Preferences for applicants A, B, and C

$\succ_{\mathrm{A}}$	$\succ_{\mathrm{B}}$	$\succ_{\mathrm{C}}$
$(\beta, 0)$	$(\alpha, 0)$	$(\alpha, 0)$
$(\beta, 1)$	$(\alpha, 1)$	$(\alpha, 1)$
$(\gamma, 0)$	$(\gamma, 0)$	$(\gamma, 0)$
$(\alpha, 0)$	$(\alpha, 2)$	$(\alpha, 2)$
$(\beta, 2)$	$(\beta, 0)$	$(\beta, 0)$
$(\gamma, 1)$	$(\gamma, 1)$	$(\gamma, 1)$
$(\alpha, 1)$	$(\beta, 1)$	$(\beta, 1)$
$(\gamma, 2)$	$(\gamma, 2)$	$(\gamma, 2)$
$(\alpha, 2)$	$(\beta, 2)$	$(\beta, 2)$

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 $u_{\mathrm{A}}(b,t) = f(b) - 3t$, 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 $u_{\mathrm{B}}(b,t) = u_{\mathrm{C}}(b,t) = g(b) - 3t$, where $g(\alpha) = 8$, $g(\beta) = 1$, $g(\gamma) = 3$.

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 Table 2.

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

Henceforth we omit the term *ex-ante* with the understanding that all properties

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 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.

pertaining to the allocation refer to expected arrival times rather than realized arrival times.

3 Multiple Waitlist Procedure

3.1 Common Priorities

Throughout this section, we consider 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 and characterize the matching that results from this mechanism.

Under MWP, there is a FIFO queue associated with each building. A room 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 room is offered to the applicant with the highest priority. The applicant can either accept the offer or opt to be placed in the queue for the next available room in a different building. Note that if a' has higher priority than a, then a' receives an assignment (i.e., a room or a place on some waiting list) before a does.

Step $(1,1)$. Let the first period in which some room becomes available be given by t_1, and let b_1 denote the building in which a room becomes available.[16] Denote by $a_{1,1}$ the applicant with the highest priority. Offer the room $\langle b_1, 1\rangle$ to applicant $a_{1,1}$.

If $a_{1,1}$ accepts the offer: then set the allocation for $a_{1,1}$ to be $\mu^{t_1}(a_{1,1}) = \langle b_1, 1\rangle$. Proceed to step $(2,1)$.

Otherwise: $a_{1,1}$ chooses a building, is placed at the end of a *First-In/First-Out* (FIFO) waiting list for that building, and will receive the next unassigned room in that building when it arrives.[17] Proceed to step $(1,2)$.

Step (i,j). Let the i^{th} period in which some room becomes available be given by t_i, and let b_i denote the building in which a room becomes available. Denote by $a_{i,j}$ the applicant with the highest priority among those who are unmatched. Offer the room $\langle b_i, \eta^{t_i}(b_i)+1\rangle$ to applicant $a_{i,j}$.

If $a_{i,j}$ accepts the offer: then set the allocation for $a_{i,j}$ to be $\mu^{t_i}(a_{i,j}) = \langle b_i, \eta^{t_i}(b_i)+1\rangle$. Proceed to step $(i+1,1)$.

Otherwise: $a_{i,j}$ chooses a building, is placed at the end of a FIFO waiting list for that building, and will receive the next unassigned room in that building when it arrives. Proceed to step $(i,j+1)$.

The following proposition characterizes the major properties of this mechanism.

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

Proof. (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

[16] If rooms become available in more than one building, choose one randomly; the outcome of the procedure will not be affected by this choice.

[17] In other words, set $\mu^{t_1}(a_{1,1}) = \langle b^*, r^*\rangle$, where b^* is chosen such that $\big(b^*, t_1 + X_{b^*}\big(1 \,\big|\, h^{t_1}\big)\big) \succ_{a_{1,1}} \big(b, t_1 + X_b\big(1 \,\big|\, h^{t_1}\big)\big)$ for all b. Almost surely this b^* is unique; in the event that it is not, the procedure can be extended by temporarily allocating no such b^* to the applicant and waiting until all but one of them are matched.

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 room 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 room $\langle b_{a'}, r_{a'} \rangle$ 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 $\big(b_{a_0}, t_0 + X_{b_{a_0}}\big(r_{a_0} \,\big|\, h^{t_0}\big)\big)$ over $\big(b_{a'}, t_0 + X_{b_{a'}}\big(r_{a'} \,\big|\, h^{t_0}\big)\big)$ as desired.

(iii) Suppose applicant a' has higher priority than another applicant a. According to MWP, a' receives an assignment 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 room assigned to a. In other words, a' does not envy the lower priority applicant a. $\square$

We have shown that MWP satisfies several desirable properties when there is a common priority ordering across objects. The next section explores the extent to which this assumption can be relaxed.

3.2 Heterogeneous priorities

In the general case when priority orderings are not common across objects, the existence of a strategy-proof mechanism that satisfies efficiency and the elimination of justified envy depends on whether the priority orderings are "similar." More precisely, the existence of such a mechanism depends on whether rankings between applicants that are nonadjacent in one priority ordering are preserved across all priority orderings. This property is formalized in the following acyclicity condition.

Definition 3 (acyclicity)**.** For any $a_1, a_2, a_3 \in A$ and $b_1, b_2 \in B$, we have

$$a_1 \succ_{b_1} a_2 \succ_{b_1} a_3 \implies a_1 \succ_{b_2} a_3.$$

In other words, the priority orderings do not contain *cycles* of the form $a_1 \succ_{b_1} a_2 \succ_{b_1} a_3 \succ_{b_2} a_1$. This condition was first introduced by Ergin (2002) and has since appeared in various static settings.[18] The following result demonstrates that acyclicity also plays an important role in our dynamic setting.

Proposition 3. *There exists a strategy-proof mechanism that satisfies efficiency and the elimination of justified envy if and only if the priority orderings are acyclic.*

Proof. We postpone the proof of sufficiency and show here that if the priority orderings violate acyclicity, then there does not exist a strategy-proof mechanism that satisfies both efficiency and the elimination of justified envy.

If $\succ_B$ does not satisfy acyclicity, 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:

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

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

[18]Key results in Kojima (2011), Kesten (2006), Dur (2012), and Romero-Medina and Triossi (2013), for example, depend on variants of acyclicity.

Table 3: Preferences for applicants A, B, and C

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

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 $h(\beta) = 2$.

is a Pareto improvement or some applicant justifiably envies another. The remaining cases of the argument are summarized in Table 4.

This demonstrates the necessity of the acyclicity condition. Sufficiency will be shown constructively in the remainder of this section. □

We now describe the generalized Multiple Waitlist Procedure for acyclic priority orderings. The 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, before an applicant can refuse an offer and join the waiting list for some building, if another applicant has higher priority at that building, then that applicant's turn comes first.[19]

Step $(1, 1)$. Let the first period in which some room becomes available be given by t_1, and let b_1 denote the building in which a room becomes available.[20] Denote by

[19] Another way of describing this using language similar to that of Abdulkadiroğlu and Sönmez (1999) would be "you request my building—I get your turn."

[20] If rooms become available in more than one building, choose one randomly; the outcome of the procedure will not be affected by this choice.

Table 4: Example used in necessity proof of Proposition 3

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 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.

$a_{1,1}$ the applicant with the highest priority for building b_1. Offer the room $\langle b_1, 1\rangle$ to applicant $a_{1,1}$.

If $a_{1,1}$ accepts the offer: then set the allocation for $a_{1,1}$ to be $\mu^{t_1}(a_{1,1}) = \langle b_1, 1\rangle$. Proceed to step $(2, 1)$.

Otherwise: $a_{1,1}$ requests to requests to join the FIFO waiting list of building $\hat{b}$, where $\hat{b}$ is chosen such that $\big(\hat{b}, t_1 + X_{\hat{b}}\big(1 \,\big|\, h^{t_1}\big)\big) \succ_{a_{1,1}} \big(b, t_1 + X_b\big(1 \,\big|\, h^{t_1}\big)\big)$ for all $b \neq \hat{b}$.

- If $a_{1,1}$ has the highest priority at the chosen building $\hat{b}$, then the request is accepted, so we set $\mu^{t_1}(a_{1,1}) = \langle \hat{b}, 1\rangle$. Proceed to step $(1, 2)$.

- If another applicant $\hat{a}_{1,1}$ has higher priority than $a_{1,1}$ at $\hat{b}$, then $\hat{a}_{1,1}$ chooses between (i) taking the available room in building b_1, (ii) joining the waiting list for building $\hat{b}$, and (iii) joining the waiting list for a different building.

  - In case (i), set $\mu^{t_1}(\hat{a}_{1,1}) = \langle b_1, 1\rangle$ and $\mu^{t_1}(a_{1,1}) = \langle \hat{b}, 1\rangle$. Proceed to step $(2, 1)$.
  - In case (ii), set $\mu^{t_1}(\hat{a}_{1,1}) = \langle \hat{b}, 1\rangle$. Repeat step $(1, 1)$.
  - In case (iii), place $\hat{a}_{1,1}$ at the end of the FIFO waiting list for her chosen

building,[21] and set $\mu^{t_1}(a_{1,1}) = \langle \hat{b}, 1 \rangle$. Proceed to step $(1, 2)$.

Step (i, j). Let the i^{th} period in which some room becomes available be given by t_i, and let b_i denote the building in which a room becomes available. Denote by $a_{i,j}$ the applicant with the highest priority for building b_i among those who are unmatched. Offer the room $\langle b_i, \eta^{t_i}(b) + 1 \rangle$ to applicant $a_{i,j}$.

If $a_{i,j}$ accepts the offer: then set the allocation for $a_{i,j}$ to be $\mu^{t_i}(a_{i,j}) = \langle b_i, \eta^{t_i}(b_i) + 1 \rangle$. Proceed to step $(i + 1, 1)$.

Otherwise: $a_{i,j}$ requests to requests to join the FIFO waiting list of building $\hat{b}$, where $\hat{b}$ is chosen such that $\big(\hat{b}, t_i + X_{\hat{b}}\big(\eta^{t_i}(\hat{b}) + 1 \,\big|\, h^{t_i}\big)\big) \succ_{a_{i,j}} \big(b, t_i + X_b\big(\eta^{t_i}(b) + 1 \,\big|\, h^{t_i}\big)\big)$ for all $b \neq \hat{b}$.

- If $a_{i,j}$ has the highest priority at the chosen building $\hat{b}$, then the request is accepted, so we set $\mu^{t_i}(a_{i,j}) = \langle \hat{b}, \eta^{t_i}(\hat{b}) + 1 \rangle$. Proceed to step $(i, j + 1)$.
- If another applicant $\hat{a}_{i,j}$ has higher priority than $a_{i,j}$ at $\hat{b}$, then $\hat{a}_{i,j}$ chooses between (i) taking the available room in building b_i, (ii) joining the waiting list for building $\hat{b}$, and (iii) joining the waiting list for a different building.
  - In case (i), set $\mu^{t_i}(\hat{a}_{i,j}) = \langle b_i, \eta^{t_i}(b) + 1 \rangle$ and $\mu^{t_i}(a_{i,j}) = \langle \hat{b}, \eta^{t_i}(b) + 1 \rangle$. Proceed to step $(i + 1, 1)$.
  - In case (ii), set $\mu^{t_i}(\hat{a}_{i,j}) = \langle \hat{b}, \eta^{t_i}(\hat{b}) + 1 \rangle$. Repeat step (i, j).
  - In case (iii), place $\hat{a}_{i,j}$ at the end of the FIFO waiting list for her chosen building, and set $\mu^{t_i}(a_{i,j}) = \langle \hat{b}, \eta^{t_i}(\hat{b}) + 1 \rangle$. Proceed to step $(i, j + 1)$.

The following result characterizes the generalized MWP under acyclic priority orderings and completes the sufficiency proof for Proposition 3.

[21] In other words, set $\mu^{t_1}(\hat{a}_{1,1}) = \langle b^*, r^* \rangle$, where b^* is chosen such that $\big(b^*, t_1 + X_{b^*}\big(1 \,\big|\, h^{t_1}\big)\big) \succ_{\hat{a}_{1,1}} \big(b, t_1 + X_b\big(1 \,\big|\, h^{t_1}\big)\big)$ for all b. Almost surely this b^* is unique; in the event that it is not, the procedure can be extended by temporarily allocating no such b^* to the applicant and waiting until all but one of them are matched.

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

Proof. (i) We will show that truthful preference revelation is weakly dominant for each applicant $a \in A$. By acyclicity, there is at most one applicant $\hat{a}$ who has higher priority than a at some buildings but lower priority at the others. During her turn, applicant a receives her most-preferred room (among those that are available), unless $\hat{a}$ prefers the same room and has higher priority for the building. Whether a receives her most-preferred room among those that are available depends only on the priorities and the stated preferences of the other applicants; and in the case that the most-preferred room of a is assigned to $\hat{a}$ (who has higher priority), a is assigned her most-preferred room among those that remain. Therefore there is no incentive 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 room 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 room 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 room among those that are available at the time of assignment, which implies that $\hat{a}_0$ prefers the original allocation $\left(b_{\hat{a}_0}, t_0 + X_{b_{\hat{a}_0}}\left(r_{\hat{a}_0} \,\middle|\, h^{t_0}\right)\right)$ over $\left(b_{a'}, t_0 + X_{b_{a'}}\left(r_{a'} \,\middle|\, h^{t_0}\right)\right)$ 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 room in a building at which she has 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 room assigned to a. Now suppose that a' prefers a room in a building at which another applicant $\hat{a}'$ has higher priority.[22]

[22] As noted earlier, acyclicity implies that there is at most one such applicant.

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 room 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' (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. □

3.3 Lottery mechanisms

Note that the preceding results are stated for a class of mechanisms that assign a particular room in a particular building to each applicant. A natural question that arises is that of whether the characterization holds more generally for a class of mechanisms that assign a lottery over rooms to each applicant. Under this more general class of mechanisms, which we refer to as *lottery mechanisms*, the notion of elimination of justified envy from definition 1 does not apply since it is unclear what it means for an applicant to have higher or lower priority for a lottery (as opposed to a particular building).

To address this issue, we use the notion of *strong stability* first introduced by Roth et al. (1993) and adapted by Kesten and Ünver (forthcoming) in the context of school-choice lotteries.[23] We say that an applicant a' *lottery-envies* another applicant a if there is a room 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 room (for her) than $\langle b, r \rangle$ with positive probability, and we say that this lottery-envy is *justified* if a' has higher priority than a for building b. A lottery mechanism is *strongly stable* if it eliminates justified lottery-envy.[24]

As the following proposition demonstrates, acyclicity remains a necessary and sufficient condition for the characterization in Proposition 3 extended to lottery

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

[24] Under non-lottery mechanisms, strong stability coincides with the elimination of justified envy because only a single room is assigned to each agent with positive probability.

mechanisms.[25]

Proposition 5. *There exists a strategy-proof lottery mechanism that satisfies efficiency and strong stability if and only if the priority orderings are acyclic.*

Proof. Necessity can be demonstrated using the same example as in Proposition 3. 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 Table 3. As before, the arrival process is deterministic: a room in building α arrives in period 0; a room in building β arrives in period 1; and another room 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. □

Even when considering the more general class of lottery mechanisms, the presence or absence of cycles in the priority orderings fully determines whether it is possible for a strategy-proof mechanism to satisfy efficiency and the elimination of justified envy. For the case that it is possible to satisfy both conditions, we provide in Proposition 4 a mechanism that does so. In a setting such as public-housing allocation, acyclicity would be satisfied if there are eligibility restrictions whereby some buildings are only available to applicants with sufficiently high priority. Additionally, acyclicity can be

[25] Although the sufficiency direction of Proposition 3 is stronger than that of Proposition 5, the necessity direction of the latter proposition is stronger.

satisfied in a system with multiple programs that share a ranking over applicants but have some discretion on final priorities based on interviews.[26]

4 Comparison of mechanisms

4.1 Characterization of public-housing allocation mechanisms

The most prominent application in which objects arrive stochastically and must be assigned to agents dynamically is that of public-housing allocation. This section begins by evaluating the theoretical properties of mechanisms that are used to assign public housing and then proceeds to investigate public-housing allocation mechanisms empirically. Using estimated preferences from a structural model of preferences for public housing due to Geyer and Sieg (2013), we find that the welfare gains from changing existing public-housing allocation mechanisms to the Multiple Waitlist Procedure introduced in Section 3 are substantial.

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 rooms. Under *Plan A*, the PHA offers a room that becomes available to the applicant with the highest priority; if the offer is refused, then the applicant is removed or placed at the bottom of the waiting list. Under *Plan B*, an applicant who refuses a room receives another offer, up to a limit of two or three total offers (Devine et al., 1999). These procedures involve finding the highest priority applicant who is willing to "accept" the available room rather than "pass." Letting k denote the upper bound on the number of rooms that an applicant is permitted to "pass," we refer to these procedures as PHA-k mechanisms.[27] Another criterion that distinguishes

[26] See, for example, the allocation procedure in the District of Columbia.

[27] PHA-0 corresponds to Plan A, the take-it-or-leave-it procedure in which no applicant can "pass" on an offer without losing their priority; PHA-1 and PHA-2 correspond to Plan B.

Table 5: Distribution of allocation procedures by size of housing agency

	Small	Medium	Large	Extra-large	Total
PHAs offering 1 or 2 units	430	114	71	3	618
PHAs offering more than 2 units	446	96	40	7	589
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 the number of units that may be offered to applicants and housing agency size. The bottom frame shows the relationship between waiting list method and housing agency size. 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.

housing allocation mechanisms is the waiting list method. Under a *centralized* waiting list system, the applicant with the highest priority can be offered a room 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 room in a particular building or in a group of buildings. Table 5 shows the number of housing agencies using each type of procedure described by HUD.

The observation that neither larger nor smaller housing authorities appear to be systematically associated with procedures that offer more choice provides suggestive evidence that administrative costs are not likely to be a significant barrier in implementing alternative allocation procedures. The procedure that we propose (i.e., MWP) can be described in the language used by HUD as offering one unit to a household on the centralized waiting list and moving it to a site-specific waiting list if the offer is refused. As shown in Section 3, MWP is strategy-proof and satisfies fairness and efficiency properties. Although PHA-k may satisfy strategy-proofness,[28] the following result illustrates that PHA-k fails to satisfy other desirable properties.

[28] For centralized waiting lists, it is clear that PHA-k is strategy-proof. However, although there exist non-manipulable procedures for assigning applicants to non-centralized waiting lists (such as random assignment), the procedures that are generally used in practice tend not to be strategy-proof. The public housing agency in New York City, for example, explicitly advises applicants to "select their first borough choice carefully."

Table 6: Preferences for each applicant a_i

$\succ_{a_0}$	$\succ_{a_1}$	$\cdots$	$\succ_{a_{k-1}}$	$\succ_{a_k}$	$\succ_{a_{k+1}}$
$(b_{k+1}, k+1)$	(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 room for applicant a_i is in building b_{k-i+1} which arrives in period $k-i+1$. Each applicant a_i prefers the room in building b_r over the room in building b_{r-1} for all $r \not\equiv -i \pmod{k+2}$.

Proposition 6. *PHA-k does not satisfy efficiency or the elimination of justified envy.*

Proof. Consider the following example with $k+2$ buildings, $B = \{b_i\}_{i=0}^{k+1}$, each with one room; $k+2$ 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 Table 6.

Rooms arrive deterministically: in period i, the room in building i becomes available with certainty.

Applicant a_0 refuses k offers and receives a room in building b_k, since all rooms that become available sooner are less desirable. For $i = 1, \ldots, k-1$, applicant a_i refuses $k-i$ offers; her first-choice room (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 room (in building b_{k-i}). Likewise, applicant a_k accepts the offer of a room in building 0, since a_{k-1} will accept the room in building 1. This leaves the room in building $k+1$ for applicant a_{k+1}. The allocation procedure concludes with each applicant receiving her second-choice room.

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

Table 5 displays the number of PHAs using the PHA-k mechanism for $k \in \{0, 1\}$ and for $k \geq 2$. Although our construction in Proposition 6 depends on k, in practice k is typically small.

4.2 Estimation of welfare gains

A question that remains in spite of the theoretical results in the preceding section is that of whether the selection of an allocation mechanism significantly affects welfare in real-world applications such as public-housing allocation. We address this question by using data on a sample of 94 households eligible for public housing in Pittsburgh, PA from the Survey of Income and Program Participation (SIPP) collected by the US Census Bureau. The model of household preferences is a standard random utility specification used 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.

We begin with a brief discussion of the market for public housing in Pittsburgh. Buildings are classified as "family," "senior," and "mixed" communities. The age distribution of residents within each community is broadly consistent with the classification, but in practice any applicant can be assigned to any building. A building can be "large" (more than 100 rooms), "medium" (between 40 and 100 rooms), or "small" (fewer than 40 rooms). The 34 buildings operated by the Housing Authority of the City of Pittsburgh (HACP) in 2001 can be placed in six categories based on classification and size: family large, family medium, family small, mixed, senior large, and senior small.

The utility of applicant i living in building j at time t is given by

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

where γ_j is a building-specific fixed effect; $y_{i,t}$ denotes net income; x_i is a vector of demographic characteristics, namely indicators for female, nonwhite, senior, and children; $d_{i,t}$ denotes the residence and c is a moving-cost parameter; and $\varepsilon_{i,j}$ captures idiosyncratic tastes for public housing.[29] The utility of living in private housing ($j = 0$) is normalized to be

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

Following McFadden (1974), we assume that the idiosyncratic components are independently and identically distributed according to a standard type-I extreme-value distribution.[30] Heterogeneity in preferences can be due to various factors: applicants may prefer to live closer to their respective workplaces, which would allow them to reduce commuting time and travel costs; unobservable characteristics of children may lead households to exhibit stronger preferences for units that are located near better schools; seniors may differ in the extent to which they prefer units that have access to amenities.

A key insight due to Geyer and Sieg (2013) is that a simple logit demand model would fail to capture the reality of strong preferences for public housing. Indeed households that live in private housing exhibit their preferences for public housing by joining waiting lists, and typical waiting times for an available unit are between 14 and 22 months. Geyer and Sieg (2013) develop an equilibrium framework that incorporates these supply-side restrictions. They identify the structural parameters of the utility function by using households' decisions to exit public housing. Table 7 report estimates of the structural parameters. These estimates suggest that minorities and female-headed households with children exhibit the strongest preferences for public housing. The fact that the coefficient β on income is less than one suggests that a higher income makes public housing less desirable. This is consistent with the fact that a household residing in public housing typically pays 30% of its income as

[29] In our empirical analysis, applicants do not have risk-neutral preferences. As footnote 12 mentions, our results do not depend on the risk-neutrality assumption.

[30] Geyer and Sieg (2013) find that nested logit specifications designed to account for correlation in unobserved preferences among public housing communities does not increase the likelihood function, and their statistical tests suggest that the logit specification is reasonable.

Table 7: Parameter estimates

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: Parameter estimates are from the model with supply-side restrictions in Table 10 of Geyer and Sieg (2013).

Table 8: Offer probabilities

Building category	Number of buildings	Empirical transition probability	Offer probability
Family large	13	0.0373	0.0708
Family medium	4	0.0258	0.0471
Family small	2	0.0086	0.0159
Mixed	4	0.0538	0.1008
Senior large	3	0.0141	0.0261
Senior small	8	0.0171	0.0311

Note: The empirical transition probability $\Pr(d_t = j \mid d_{t-1} = 0)$ is from the Housing Authority of the City of Pittsburgh, as reported in Table 3 of Geyer and Sieg (2013). The offer probability Π_j is computed using the empirical transition probabilities and the parameter estimates in Table 7.

rent each month.

The distributional assumptions on the idiosyncratic component of the utility function enable us to express the probability $\Pi_{j'}$ of receiving an offer to move from private housing into building j' in terms of structural parameters and observed data as

$$\Pr(d_t = j' \mid d_{t-1} = 0) = \frac{\exp(u_{j'})}{\exp(u_0) + \exp(u_{j'})}\Pi_{j'},$$

where u_j denotes the non-idiosyncratic component of the utility of being matched with building j for the average applicant. Table 8 reports observed transitions from private housing to public housing and the implied offer probabilities.

Applicants maximize a discounted sum of utilities with a monthly discount factor of $\delta = 0.96$. This implies an annual discount factor of about 0.61. We assume that every match lasts for two years ($T = 24$ months), which is below the average length of time that a typical household spends in public housing.[31] Due to a lack of applicant-level data on match duration, we use a conservative estimate of two years to avoid bias due to possible correlation between match duration and unobserved

[31]Based on the Pittsburgh data as reported by Geyer and Sieg (2013), the average length of time in public housing is almost seven years.

characteristics.[32] A lower match duration decreases the benefit of being matched with a more-preferred building, and a lower discount factor increases the cost of waiting for a more-preferred building. Since we are interested in estimating the welfare gains from using a mechanism that provides agents more choice at the cost of additional waiting time, we interpret our estimates as a lower bound on the actual welfare gains.

For an applicant who is assigned to a room in building j' that becomes available in period t', the assigned building as a function of time is given by $j(t) = j'\mathbf{1}_{\{t' \leq t < t+T\}}$. In other words, the applicant resides in private housing in all periods except for the two years immediately after moving into public housing. The lifetime utility for applicant i is thus given by

$$
\begin{aligned}
U_i(j', t') &= \sum_{t=0}^{\infty} \delta^t u_{i,j(t),t} \\
&= \frac{u_{i,0,0}}{1-\delta} + \delta^{t'} \left(\left(u_{i,j',0} - u_{i,0,0}\right) \frac{1-\delta^T}{1-\delta} + c\mathbf{1}_{\{j' \neq 0\}} \left(1 + \delta^T\right) \right).
\end{aligned}
$$

Under MWP, since the arrival time of the r^{th} room in a building follows a negative binomial distribution, we use closed-form expressions to compute the expected utilities from joining waiting lists. Under the PHA-k mechanisms, we use numerical approximations to determine applicants' expected utilities from passing on an offer. Given these expected utilities, we use the offer probabilities in Table 8 to simulate arrival processes and determine the allocations that would result under each mechanism.

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 PHA-k and MWP for each applicant by computing the equivalent variation, i.e., the additional amount of income that the applicant would have to receive each month when public housing is assigned by the PHA-k mechanism that would give the applicant the same lifetime utility as the assignment under MWP. Table 9 reports the average of the present-discounted values of these

[32]For example, if a household with short-term needs for public housing has larger utility gains from being matched with its most-preferred building, then using the average match duration may overstate the benefit of changing the allocation mechanism.

Table 9: Comparing welfare under MWP and PHA-k mechanisms

Mechanism	Mean $\overline{\text{EV}}$	Min $\overline{\text{EV}}$	Max $\overline{\text{EV}}$	EV > 0
PHA-0	$7,543	$4,563	$11,747	66–78%
PHA-1	$5,957	$3,482	$8,647	64–77%
PHA-2	$5,641	$3,369	$8,500	61–80%

Note: This table contains the results of 15 counterfactual simulations. In each simulation, we compute the average across all applicants of the equivalent variation ($\overline{\text{EV}}$) of changing the allocation mechanism to MWP. The second, third, and fourth columns report the mean and range of present-discounted values ($\delta = 0.96$) across simulations. The final column provides the range across simulations of the fraction of applicants who prefer the allocation under MWP.

payments for PHA-0, PHA-1, and PHA-2. Our estimates suggest that a change from PHA-k to MWP improves the welfare of the average applicant who receives a housing assignment by an amount that is equivalent to a transfer payment of between $5,600 and $7,600. As discussed earlier, this understates the actual gains from changing the allocation mechanism due to our assumptions of a low discount rate and a short match duration. With 1.2 million households living in public housing in the US, the overall welfare gains from improved matching are substantial.

5 Conclusion

We depart from the usual matching problem with priorities by studying situations that involve stochastic arrival. Instead of using static capacities, our model captures uncertainty regarding the availability of objects. In the case of common priority orderings, we introduce a mechanism (MWP) that is strategy-proof, efficient, and eliminates justified envy. We then construct a general version of MWP for which the same properties hold with acyclic priority orderings. These mechanisms do not rely on any particular stochastic process that governs the arrival rate of units.

We compare these mechanisms with the procedures that public housing agencies currently use (PHA-k mechanisms) to allocate rooms. Our empirical analysis uses estimated preferences for public housing from a structural model to simulate housing assignments under counterfactual allocation mechanisms. Using a sample of households eligible for public housing in Pittsburgh, we find that the welfare gains from changing the most commonly used PHA-0 and PHA-1 allocation mechanisms to MWP are about \$6,000 and \$7,500 per applicant, respectively. This empirical finding complements our characterization of the PHA-k mechanisms, which reveals that these existing procedures may result in inefficiency and justified envy.

Our main characterization result shows that it is impossible to design a mechanism that satisfies all three properties in general: the existence of such a mechanism depends on whether the priority orderings satisfy acyclicity. Since acyclicity imposes a restriction on the priority orderings that are consistent with both efficiency and the elimination of justified envy, we conclude by suggesting strategy-proof mechanisms for arbitrary priority orderings that satisfy each of these desirable criteria separately.

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 at each step the applicant with the highest priority at the building that becomes available, since only the applicant with the highest priority receives an assignment at each step. This class of procedures produces efficient allocations: no applicant would benefit from an alternate allocation since each chooses her most-preferred room 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 abandoning the FIFO property of the waiting lists. In particular, consider an alternative in which an applicant who opts to be placed on a waiting list for a different building receives a room only after every applicant with higher priority for that building receives some assignment. If applicant a is assigned a room in building b, then any applicant a' with higher priority receives either an earlier room in building b or a more-preferred allocation, so there is no justified envy. However, there can be inefficiencies since applicants who prefer rooms in buildings at which their priorities are low may benefit from switching their assignments.

References

Abdulkadiroğlu, A., N. Agarwal, and P. A. Pathak (2014): "The Welfare Effects of Congestion in Uncoordinated Assignment: Evidence from the NYC HS Match," *Mimeo*. 6

Abdulkadiroğlu, A., Y.-K. Che, and Y. Yasuda (2012): "Expanding "Choice" in School Choice," *Mimeo*. 6

Abdulkadiroğlu, A. and S. Loerscher (2007): "Dynamic house allocations," *Mimeo*. 5

Abdulkadiroğlu, A., P. A. Pathak, and A. E. Roth (2009): "Strategy-Proofness versus Efficiency in Matching with Indifferences: Redesigning the NYC High School Match," *American Economic Review*, 99, 1954–78. 9

Abdulkadiroğlu, A. and T. Sönmez (1999): "House Allocation with Existing Tenants," *Journal of Economic Theory*, 88, 233–260. 16

Agarwal, N. and P. Somaini (2014): "Demand Analysis using Strategic Reports: An application to a school choice mechanism," Working Paper 20775, National Bureau of Economic Research. 6

Akbarpour, M., S. Li, O. Gharan, and Shayan (2014): "Dynamic matching market design," *Mimeo*. 4

Azevedo, E. M. and E. B. Budish (2013): "Strategy-Proofness in the Large," *Mimeo*. 9

Balinski, M. and T. Sönmez (1999): "A tale of two mechanisms: student placement," *Journal of Economic Theory*, 84, 73–94. 9

Damiano, E. and R. Lam (2005): "Stability in dynamic matching markets," *Games and Economic Behavior*, 52, 34–53, 00008. 5

Devine, D. J., L. Rubin, and R. W. Gray (1999): *The uses of discretionary authority in the public housing program: A baseline inventory of issues, policy, and practice*, US Department of Housing and Urban Development, Office of Policy Development and Research. 22

Dimakopoulos, P. D. and C.-P. Heller (2014): "Matching with Waiting Times: The German Entry-Level Labour Market for Lawyers," *Mimeo*. 4

Doval, L. (2014): "A theory of stability in dynamic matching markets," *Mimeo*. 4

Dur, U. (2011): "Tie-Breaking is Still a Problem: How to Find the Pareto Efficient and Fair Matching in the School Choice Problem," *Mimeo*. 6

——— (2012): "Dynamic school choice problem," *Mimeo*. 5, 15

Erdil, A. and H. Ergin (2008): "What's the Matter with Tie-Breaking? Improving Efficiency in School Choice," *American Economic Review*, 98, 669–89. 6

Ergin, H. I. (2002): "Efficient resource allocation on the basis of priorities," *Econometrica*, 70, 2489–2497. 15

Gale, D. and L. Shapley (1962): "College admissions and the stability of marriage," *The American Mathematical Monthly*, 69, 9–15. 2

Gershkov, A. and B. Moldovanu (2009a): "Dynamic Revenue Maximization with Heterogeneous Objects: A Mechanism Design Approach," *American Economic Journal: Microeconomics*, 1, 168–98. 7

——— (2009b): "Learning about the Future and Dynamic Efficiency," *American Economic Review*, 99, 1576–87. 4, 7

Gershkov, A., B. Moldovanu, and P. Strack (2014): "Dynamic allocation and learning with strategic arrivals," *Mimeo*. 4

Geyer, J. and H. Sieg (2013): "Estimating a model of excess demand for public housing," *Quantitative Economics*, 4, 483–513. 6, 22, 25, 26, 27, 28

Glaeser, E. L. and E. F. P. Luttmer (2003): "The Misallocation of Housing Under Rent Control," *American Economic Review*, 93, 1027–1046. 6

Hafalir, I. E., M. B. Yenmez, and M. A. Yildirim (2013): "Effective affirmative action in school choice," *Theoretical Economics*, 8, 325–363. 6

Hylland, A. and R. Zeckhauser (1979): "The efficient allocation of individuals to positions," *The Journal of Political Economy*, 293–314. 2

Kadam, S. and M. Kotowski (2014): "Multi-period matching," *Mimeo*. 6

Kennes, J., D. Monte, and N. Tumennasan (2014a): "The day care assignment: a dynamic matching problem," *American Economic Journal: Microeconomics*, 6, 362–406. 5

——— (2014b): "Dynamic Matching Markets and the Deferred Acceptance Mechanism," *Mimeo*. 5

Kesten, O. (2006): "On two competing mechanisms for priority-based allocation problems," *Journal of Economic Theory*, 127, 155–171. 15

Kesten, O. and M. U. Ünver (forthcoming): "A theory of school-choice lotteries," *Theoretical Economics*. 6, 20

Kojima, F. (2011): "Robust stability in matching markets," *Theoretical Economics*, 6, 257–267. 15

Kurino, M. (2009): "Credibility, efficiency and stability: a theory of dynamic matching markets," *Mimeo*. 5

——— (2014): "House allocation with overlapping generations," *American Economic Journal: Microeconomics*, 6, 258–289. 5

Leshno, J. (2012): "Dynamic matching in overloaded systems," *Mimeo*. 4, 5

McFadden, D. (1974): "Conditional logit analysis of qualitative choice behavior," in *Frontiers in Econometrics*, Academic Press: New York, 105–142. 26

Morrill, T. (2013): "Making Efficient School Assignment Fairer," *Mimeo.* 6

Pereyra, J. S. (2013): "A dynamic school choice model," *Games and Economic Behavior*, 80, 100–114. 5

Romero-Medina, A. and M. Triossi (2013): "Games with capacity manipulation: incentives and Nash equilibria," *Social Choice and Welfare*, 41, 701–720. 15

Roth, A. E., U. G. Rothblum, and J. H. Vande Vate (1993): "Stable Matchings, Optimal Assignments, and Linear Programming," *Mathematics of Operations Research*, 18, 803–828. 20

Ünver, M. U. (2010): "Dynamic Kidney Exchange," *The Review of Economic Studies*, 77, 372–414. 4

Van Ommeren, J. N. and A. J. Van der Vlist (2014): "Households' willingness to pay for public housing," *Mimeo.* 6

Wang, S.-Y. (2011): "State Misallocation and Housing Prices: Theory and Evidence from China," *American Economic Review*, 101, 2081–2107. 6