Design of Lotteries and Waitlists for Affordable Housing Allocation

Nick Arnosti
Columbia University
nicholas.arnosti@gmail.com

Peng Shi
University of Southern California
pengshi@marshall.usc.edu

July 27, 2018

Abstract

We study a setting in which dynamically arriving items are assigned to waiting agents, who have heterogeneous values for distinct items and heterogeneous outside options. An ideal match would both *target* items to agents with the worst outside options, and *match* them to items for which they have high value.

Our first finding is that two common approaches – using independent lotteries for each item, and using a waitlist in which agents lose priority when they reject an offer – lead to identical outcomes in equilibrium. Both approaches encourage agents to accept items that are marginal fits. We show that the quality of the match can be improved by using a common lottery for all items. If participation costs are negligible, a common lottery is equivalent to several other mechanisms, such as limiting participants to a single lottery, using a waitlist in which offers can be rejected without punishment, or using artificial currency.

However, when there are many agents with low need, there is an unavoidable tradeoff between matching and targeting. In this case, utilitarian welfare may be maximized by focusing on good matching (if the outside option distribution is light-tailed) or good targeting (if it is heavy-tailed). Using a common lottery achieves near-optimal matching, while introducing participation costs achieves near-optimal targeting.

1 Introduction

Lotteries and waitlists are commonly used to ration items for which demand exceeds supply. For example, New York City allocates public housing using a waitlist, and allocates newly-built affordable housing by lottery. Many Broadway shows, musicians, and sports teams offer lotteries for discounted tickets. Organs from deceased donors are typically allocated using a waitlist. Occasionally, more complex allocation systems are employed – for example, Feeding America allows food banks to bid for donations using a virtual currency (Prendergast 2017).

Designers of these systems face many questions, such as

a) Is it better to use lotteries, waitlists, or an artificial currency system?

b) When using lotteries, should there be a limit on how many times each agent can apply?

c) In a waitlist, should agents who reject an offer keep their spot in line, or lose it?

We address these questions by studying how different allocation systems perform according to the following objectives:[1]

- **Targeting** individuals with the highest need. Food banks and housing assistance programs target low-income individuals, and organs are preferentially allocated to sicker patients.
- **Matching** individuals with items that are well-suited to their needs. Food bank populations differ in their diet, housing units are in different locations of the city, and organs have different biological markers.

Targeting may be achieved *explicitly* through eligibility and priority rules based on observable characteristics, or *implicitly* due to the fact that agents with different levels of need make different choices about where to apply and what to accept. In this paper, we focus on implicit targeting by studying anonymous mechanisms – that is, mechanisms that treat all eligible applicants identically. In many settings, anonymity is a reasonable approximation of current practice. In settings where agents are given priority based on observable characteristics, our study can be interpreted as analyzing the allocation within each priority group.

We reach several high-level conclusions. First, form of the system – lottery, waitlist, or virtual currency – matters much less than design details that influence whether individuals will be selective when applying. In fact, seemingly disparate approaches may yield identical outcomes in equilibrium. Second, using a common lottery to determine priority for all items results in better matching than several systems found in practice, and near-optimal matching if agents remain eligible for many periods.

[1] These objectives appear to be fairly universal. For example, a policy document released by the government of the United Kingdom lists "Support for those in greatest housing need" as an "outcome which allocation policies must achieve", and lists "Greater choice and wider options" as an "outcome which the Government believes allocation policies should achieve." See `http://webarchive.nationalarchives.gov.uk/20120919214909/http://www.communities.gov.uk/documents/housing/pdf/1403131.pdf`

Third, when there are many eligible agents with low need, there is a tradeoff between matching and targeting: improving one comes at the expense of the other. In these cases, we give simple rules on which objective to prioritize based on the shape and support of the outside option distribution, and discuss ways to target effectively.

We now discuss each of these conclusions in more detail.

Equivalence of Common Allocation Procedures

We use as a leading example the allocation of affordable housing in New York, where developers receive a tax break if they offer a fraction of their newly-built units to low- and mid-income residents. These units are awarded by lottery when the development is completed, and lotteries are independent across developments. Similar systems are used in Toronto and many cities in India.

We capture the main features of this setting using a stylized model in which developments arrive over time, and upon arrival, are allocated to agents who are waiting for them. Each agent has an outside option, which is heterogeneous across the population, and agents have different values for each development. To limit ourselves to anonymous mechanisms, we assume that agent values and outside options are private information. We model current practice in New York as follows:

- *Independent Lotteries*: In each period, agents may enter a lottery for a unit at the current development. Tickets are drawn until the development fills or all tickets have been selected.

In Providence, public housing is allocated using a waitlist in which people who reject an offer lose their position on the list. Minneapolis uses a similar system, but waits until a second rejection before removing an applicant from their list. These practices motivate the following definition:

- *Waitlist without Deferral.* Entering agents are placed at the bottom of a waitlist. In each period, the current development is offered to agents at the top of the waitlist until it fills or is offered to every agent. Agents who reject an offer lose their spot and must reapply.

Our first finding, Theorem 1 a), states that despite their very different descriptions, these two systems lead to identical outcomes in equilibrium.

Match Quality

Independent lotteries and the waitlist without deferral often fail to match agents to developments that fit their needs. The reason is that when lottery odds are low or waits are long, agents are willing to accept almost any development, and are therefore assigned nearly at random. One way to improve match quality is to determine the winners and losers using a common lottery.

- *Common Lottery*: Agents are assigned a random priority number upon entering the system. In each period, agents may apply for a unit at the current development. Units are offered to agents in order of their priority number.

In a common lottery, winners (those with good priority numbers) can get anything they choose, and will therefore be selective in where they apply. They are protected from competing with the lottery losers, who never get matched. Hence, everyone who matches is matched to a good fit. Theorem 2 a) shows that a common lottery always achieves better matching than the previous two mechanisms. In fact, when agents remain eligible for many periods, a common lottery approximately maximizes match quality over the space of all anonymous mechanisms.

Tradeoff between Matching and Targeting

Although a common lottery improves matching, it might perform poorly in terms of targeting. The intuition, formalized in Theorem 2 b), is that in a common lottery, agents are selected at random and therefore all agents match at similar rates. With independent lotteries, meanwhile, agents with worse outside options enter more lotteries and match at higher rates.

When there are many agents with low need, Theorem 3 establishes that the tradeoff between matching and targeting is not unique to the mechanisms above, but holds for any pair of anonymous mechanisms. This tradeoff does not hold when all agents have high need, and apply to all developments: in that case, a common lottery outperforms independent lotteries on both matching and targeting.

Maximizing Utilitarian Welfare

When there are many agents with low need – so that there is a tradeoff between matching and targeting – Theorem 4 gives guidance on which objective to prioritize, in order to maximize utilitarian welfare. If the distribution of outside options is light tailed, it is more important to match well, for example by using a common lottery. If the distribution of outside options is heavy tailed, it is most important to target effectively. In this case, Theorem 6 shows that it is approximately optimal to increase participation costs until only the most needy agents apply. This resembles the approach commonly used to allocate discounted tickets to popular sports events or shows, where agents engage in a costly competition by physically waiting in long queues.

When all agents have high need, it is not worthwhile to try to achieve targeting endogenously. Theorem 5 shows that if all applicants have sufficiently poor outside options, then a common lottery is approximately welfare optimal, *regardless* of the shape of the outside option distribution. In the context of affordable housing in New York, we believe that most individuals who meet the income restrictions would prefer many of the new buildings to their current living situation. This conjecture

is consistent with the findings of Waldinger (2018), who estimates that in Cambridge, Massachusetts, 31% of applicants (and 56% of those in the very low-income category) prefer every development to their outside option. If a similar pattern holds in New York, our results suggest that a common lottery would achieve high utilitarian welfare.

Regardless of whether matching or targeting is more important in a given context, our findings suggest that the mechanisms most commonly implemented for allocating affordable housing – independent lotteries and the waitlist without deferral – yield strictly sub-optimal welfare. In the extreme case where agents prefer every development to their outside options, supply is scarce, and participation costs are negligible, these mechanisms match random people to random items, which is the worst possible outcome for both matching and targeting (see Proposition 5 in Appendix E).

Other Ways to Achieve Good Matching

A broader message of our paper is that what is most important in designing an allocation system is not the form of the mechanism, but whether it incentivizes agents to match to marginal fits. In fact, just as independent lotteries and the waitlist without deferral are equivalent, Theorem 1 b) shows that the following mechanisms are equivalent:

- *Waitlist with Deferral*: A waitlist in which agents keep their spot after rejecting an offer.[2]
- *Ticket-Saving Lottery*: Hold a lottery for every development. All agents receive a ticket each period, which can be used for the current development or for any future development. Agents can enter multiple tickets in a given lottery, and are allocated if any of their tickets wins.

If participation costs are negligible, Theorem 1 c) shows that both of the above are equivalent to the common lottery, and all three are equivalent to the following:

- *Single-Entry Lottery*: allocate by a separate lottery for every development, but restrict each agent to enter at most one lottery in his or her lifetime.

These results suggest that there are multiple ways to improve upon the inefficient matching induced by independent lotteries, and that cities and housing authorities can select among them on the basis of other criteria, such as the ease of implementing the system and explaining it to participants.

[2]Such a mechanism is used to allocate public housing in the Amsterdam metropolitan area (Van Ommeren and Van der Vlist 2016).

2 Related Work

This paper contributes to the growing literature on dynamic matching markets. For reviews of the literature on static matching markets, see Roth and Sotomayor (1992) or Sönmez and Ünver (2011).

One strand of the dynamic matching literature focuses on generalizing the concept of stable matchings in static two-sided markets to dynamic settings. Papers that fall into this category include Damiano and Lam (2005), Kurino (2009), Pereyra (2013), Kennes et al. (2014), and Doval (2018). Contrasted with this work, the markets we study are one-sided as items have no preferences. Thus, the concept of stability is not meaningful.

Another set of papers assume that the social planner has all relevant information about the quality of each match. Much of this literature focuses on the application to kidney exchange, which started with the seminal paper of Roth et al. (2004). Representative recent works include Dickerson et al. (2012), Gurvich and Ward (2014), Akbarpour et al. (2014), Baccara et al. (2018), and Ashlagi et al. (2018). In earlier work, Kaplan (1987a,b, 1988) formulates the allocation of affordable housing as a queueing problem, and studies waiting times and development diversity under various priority rules. In contrast with the papers above, we assume that most of the relevant match information is privately known and revealed strategically by agents.

Our paper falls into the category of dynamic matching with private, one-sided preferences. A series of papers in this category is motivated by the allocation of cadaver organs (Su and Zenios 2004, 2005, 2006, Schummer 2016, Agarwal et al. 2017). In this setting, items (organs) are perishable and thus can only be offered to a limited number of individuals, and agents agree on their relative preferences across organs. Su and Zenios (2004) advocate for switching from a first-come-first-serve queue to something resembling a last-come-first-serve queue, in order to make agents less picky and increase the utilization of less desirable organs. Schummer (2016) notes that preventing agents from rejecting offers may decrease wastage, at the expense of reducing match quality for agents at the top of the queue. Agarwal et al. (2017) study the organ wastage problem from an empirical perspective and estimate agent preferences from data and simulate counterfactuals. In our setting, wastage is not a concern, and preference heterogeneity is horizontal rather than vertical. As a result, it is generally preferable to induce agents to be more (rather than less) selective.

Closer to our work are the papers of Bloch and Cantala (2017) and Leshno (2015), which are motivated by the allocation of subsidized housing units, and focus on how to match people to the right places. Bloch and Cantala (2017) find, as we do, that the waitlist with deferral induces agents to be pickier than under independent lotteries, resulting in higher match quality. Leshno (2015) notes that agents who have a middling position in a waitlist with deferral would be more selective under a hybrid mechanism, which makes offers randomly among all agents with sufficiently high positions

on the waitlist. The biggest difference between our work and these papers is that our agents have heterogeneous outside options, and thus the efficiency of a matching depends crucially on *which* agents match. Additionally, by studying a continuum model, we are able to consider richer environments (rather than assuming that values for a development are binary), and develop new insights about the equivalence of various mechanisms.

The problem of targeting aid to certain sub-populations has been considered in the public finance literature on the design of subsidies. Nichols and Zeckhauser (1982) and Blackorby and Donaldson (1988) use a simple model with two agent types to illustrate that one can target the type with higher need by restricting the flexibility of the subsidies or by adding friction. A similar idea appears in a series of papers on "money-burning auctions"(Hartline and Roughgarden 2008, Hoppe et al. 2009, Condorelli 2012, Chakravarty and Kaplan 2013), in which a social planner allocating a homogeneous good cannot use monetary payments to determine who value it the most, but may screen agents based on how much wasteful effort they are willing to incur. Several of these papers have results resembling our Theorem 4: when the valuation distribution is heavy-tailed, the designer should use wasteful effort to improve targeting; when it is light-tailed, it is more efficient to allocate randomly. We extend this insight to a setting where agents care about *which* good they receive, and illustrate that reducing match quality – instead of requiring wasteful effort – is an alternative way of targeting agents with greater need.

Finally, there is a growing empirical literature on the allocation of affordable housing. Glaeser and Luttmer (2003) provide evidence on the misallocation of rent controlled housing in New York City, and argue that it is caused by the random matching that arises from rationing. We show that in spite of the reality of scarce supply, there are mechanisms that can improve the matching. Geyer and Sieg (2013), Sieg and Yoon (2017), and Waldinger (2018) estimate random utility models of development choice using public housing data from Pittsburgh, New York, and Cambridge respectively. All three papers assume a certain parametric form for the outside option of agents, and use this to separately identify agent values for various developments and agent outside options; these two entities respectively correspond to our F and G distributions in Section 3.1, except that the empirical papers allow for a richer correlation structure through the use of agent and development characteristics. Thakral (2016) simulates the demand model of Geyer and Sieg (2013) and reports significant welfare gains by switching from the waitlist without deferral to alternatives that encourage greater selectivity. Waldinger (2018) performs simulations using Cambridge data, and reports that increasing choice leads to better matching but worse targeting (similar to our Theorem 3), but the net benefit in social welfare is positive (similar to our Theorem 4 a)). Both papers estimate economically significant welfare gains from switching to a mechanism that improves matching: Thakral (2016) estimates gains equivalent to a cash transfer of \$2,572 per unit per year in Pittsburgh, and Waldinger (2018) estimates gains equivalent to a transfer

of $1,557 per unit per year in Cambridge.

Our theoretical analysis complements the empirical works by showing that the insights of poor match quality from the waitlist without deferal, matching-targeting trade-off, and positive net benefit of better matching on welfare are not particular to the data from these cities, but also hold for a wide variety of distributions and allocation mechanisms. Furthermore, our theory can also guide empirical researchers on what functional forms to explore for robustness checks. For example, Geyer and Sieg (2013), Sieg and Yoon (2017), Waldinger (2018) all parameterize outside options within each demographic group as a linear function of the logarithm of household income, which makes it likely to be approximately light-tailed as household income is known to be approximately log-normal distributed for low-middle income groups (Clementi and Gallegati 2005). It is possible that the simulation result from Thakral (2016) and Waldinger (2018) that mechanisms that encourage selectivity have better utilitarian welfare is an artifact of the parametric form. One robustness check suggested by our Theorem 4 for these researchers is to also explore heavy-tailed parameterizations of outside options, such as using square root of income instead of the logarithm.[3]

3 Model

Section 3.1 describes the timing of agent arrivals, and our assumptions about agent utilities. Section 3.2 discusses the dynamic decision problem facing each agent. Section 3.5 defines our equilibrium concept, which builds on a definition of optimal agent strategies (Section 3.3) and match outcomes (Section 3.4). Section 3.6 introduces the metrics that we use to evaluate equilibria. For clarity of exposition, we refer to all agents using female pronouns.

3.1 Agents, Outcomes, Utilities, Timing

Time is discrete. In every period j, a continuum of agents of unit mass arrives, as does a new development, which can house a mass μ of agents and must be allocated immediately. We refer to μ as the *supply-demand* ratio.

Each agent will eventually either be matched to a single development, or depart from the system unmatched. Before being matched, agent i receives payoff α_i in each period, and after being matched to development j, the agent receives payoff v_{ij} in each period. We refer to α_i as the agent's *outside option*, and v_{ij} as her *value* for development j. We sometimes refer to α_i as the *type* of agent i.

Each agent (matched or not) has a life event with probability $1 - \delta$ in each period, after which she becomes ineligible for future allocations and leaves her affordable unit if she has been allocated one.[4]

[3] If income Y is log-normal distributed, then using Y^a for any parameter $a \in (0, 1)$ results in a heavy-tailed distribution of outside options that also exhibits diminishing returns to scale with respect to income.

[4] Life events can be thought of as capturing scenarios such as marrying and moving in with a partner, receiving a big

We normalize an agent's utility after her life event to be zero. We assume that the timing of the life event is independent of the agent's past actions. The expected number of periods before an agent's life event is $\frac{1}{1-\delta}$: we refer to this quantity as her expected eligibility time.

The timing within each period is as follows:

1. **Arrival and Participation Choice**: A unit mass of new agents arrives, and each is given a "state" (typically, a lottery number or waitlist position). Unmatched agents who remain choose whether to continue to participate in the mechanism, or exit forever. Those who participate incur participation cost $c \geq 0$. For convenience, we assume that agents who are exactly indifferent between participating or exiting will choose to participate.[5]

2. **Life Event** Every agent (matched or not) has a major life event with probability $1-\delta$, in which case she leaves her current housing. Her utility after this point is normalized to zero.

3. **New Development and Matching:** A development of mass μ arrives, labeled j. Each agent i observes v_{ij}. Agents participate in a matching rule as described in Section 3.2. Those who are matched become ineligible for future matches.

4. **Payoff**: Every agent that remains receives a payout that depends on her current housing (v_{ik} if in development k, and α_i if unmatched).

The outside options α_i are distributed according to CDF F, and the values v_{ij} are drawn iid (across agents and developments) from CDF G. We refer to F as the *outside option distribution*, and G as the *value distribution*. For convenience, we assume that distributions F and G are continuous, with strictly positive density on their domains $(\underline{\alpha}, \overline{\alpha})$ and $(\underline{v}, \overline{v})$ respectively.[6] We allow for the possibility that $\underline{\alpha}$ or $\underline{v}$ may be $-\infty$, or that $\overline{\alpha}$ or $\overline{v}$ may be ∞, and assume without loss of generality that $\overline{\alpha} \leq \overline{v}$ (agents with outside options exceeding $\overline{v}$ will never choose to participate, so we can exclude them and normalize F appropriately). We denote the density of F by the function $f(\alpha)$, and define $\overline{G}(v) = 1 - G(v)$.

3.2 Actions

The values v_{ij} and the outside option α_i are privately known to agent i. Thus, they cannot be directly used to determine an allocation. Instead, agents participate in a *matching rule*, which asks them

promotion and moving to a nicer apartment close to work, or relocating to another city to take care of an elderly family member. Mathematically, the presence of these life events ensures that the system is stable – that is, the number of unmatched agents does not grow indefinitely. All of our results continue to apply in a model where the probability of a life event in a given period depends on whether the agent is currently matched.

[5]The assumption that indifferent agents will participate is not substantial, since the set of such agents has measure zero in our model. This assumption allows us to rule out mixed strategies for these agents, simplifying the characterization of equilibrium outcomes (see Propositions 2 and 6 in the Appendix).

[6]The assumptions of continuity and positive density for F and G are not necessary for any of our results in the main body. They are only included to simplify the notation in the proofs and remove uninteresting technicalities, such as allowing mixed strategies for agents who are exactly indifferent in order to ensure that the market clears exactly.

to take an action in each period, and uses the actions to determine who will match to the current development.

Before giving our formal definition of a matching rule, we motivate this definition: Although agents are in principle playing a dynamic game, we restrict attention to designs in which agents are affected only by the *aggregate* profile of actions selected by others, and assume that agents respond to this aggregate (rather than to actions of specific other agents). This implies that no single agent can directly influence the market, or the future behavior of others. Therefore, each agent perceives herself not as playing a dynamic game, but rather as facing a Markov decision process (MDP). She begins each period in some state, which determines the set of actions available to her.[7] Her action, in turn, influences whether she matches, and which state she transitions to in the event that she does not match.

Definition 1 (Matching rule)**.** A matching rule $R = (S, D, A, T)$ *specifies a countable set of states* S*, and a distribution* D *over* S *specifying the probability of assigning each state* $s \in S$ *as the initial state. There is a countable set of actions* $A = \bigcup_{s \in S} A_s$*, where* A_s *is a finite set of actions for state* $s \in S$*. For each action* $a \in A_s$*, there is a transition function* $T_a : S \times (S \cup \{m\}) \to [0, 1]$*, where* $T_a(s, s')$ *is the probability of transitioning to state* s' *after taking action* a *in state* s*, and* m *corresponds to being matched to the current development.*

Implicit in the above definition is the assumption that the mechanism is **anonymous**: it can differentiate agents based on the history of actions taken, but not based on the identity of the agent. Another implicit assumption is that the mechanism is **stationary**, meaning that in our continuum model, the aggregate profile of agent types and actions is deterministic and constant across periods. For this reason, the transition function T_a is not indexed by the period j.

Below, we describe how the allocation systems described in the introduction can be encoded as matching rules. The *lottery-based* rules are fully characterized by a success probability $p \in (0, 1]$, which is the chance that any given ticket will win a lottery. The *waitlist-based* rules are characterized by an average idle time $\tau \geq 0$, which is the expected number of periods a newly arrived agent must wait before receiving an offer.

I. **Lottery Matching Rules:**

- **Independent Lotteries**. S consists of a single state. In it, the agent chooses from the action set $\{Enter, Abstain\}$. If she abstains, she is not matched. If she enters, she is matched with probability p. When $p = 1$, we refer to this as the **guaranteed choice** matching rule.

[7] For example, her state may be the number of periods that she has waited. If she has just arrived, she can only continue to wait, whereas if she has waited for a long time, she may be offered the current development and asked to accept or reject the offer. Alternately, her state may represent the number of lotteries that she has entered so far, or her priority as determined by a common lottery.

- **Common Lottery**. $S = \{0, 1\}$. The initial state of an agent is 1 with probability p, and an agent's state remains the same in every period. In both states, agents choose from the action set $\{Enter, Abstain\}$. An agent is matched if and only if she is in state 1 and chooses to enter (agents in state 0 will never match).
- **Single-Entry Lottery**. $S = \{0, 1\}$. All agents start in state 1, from which they can choose from actions $\{Enter, Abstain\}$. If an agent abstains, she does not match and remains in state 1. If she enters, she matches with probability p, and otherwise transitions to state 0, from which she will never be matched.
- **Ticket-Saving Lottery**. $S = \mathbb{N}$ represents the number of tickets possessed by the agent. The agent starts in state 1. From state s, the agent chooses an action $j \in \{0, \ldots, s\}$ (the number of tickets to use this period). An agent in state s choosing action j matches with probability $1 - (1 - p)^j$, and otherwise transitions to state $s - j + 1$.

II. **Waitlist Matching Rules**:

The state space is $\mathbb{N}$, representing the number of periods that the agent has waited. The initial state is zero. In states $s < \lfloor \tau \rfloor$, the agent has a single action $\{Wait\}$, and transitions deterministically to state $s + 1$. In states $s > \lfloor \tau \rfloor$, the agent selects an action from $\{Accept, Reject\}$. From state $s = \lfloor \tau \rfloor$, the agent is offered the action $\{Wait\}$ with probability $\tau - \lfloor \tau \rfloor$, and otherwise offered the actions $\{Accept, Reject\}$.[8] An agent matches if and only if she chooses *Accept*. The two variants are as follows:

- **Waitlist with Deferral**. Agents who reject retain their position (increment their state).
- **Waitlist without Deferral**. Agents who reject lose their position (go back to state 0).

We define a **mechanism** M to be a class of matching rules. For example, the *independent lotteries mechanism* is the set of all independent lotteries matching rules, parameterized by all possible values of success probability p. Analogously define the mechanisms for the common lottery, single-entry lottery, ticket-saving lottery and waitlists with and without deferral.

3.3 Strategies

A matching rule $R = (S, D, A, T)$ (along with the probability δ of remaining eligible, a value α for going unmatched, the value distribution G, and the participation cost c) induces an MDP for each agent. A **strategy profile** Σ consists of a Markovian **strategy** $\Sigma(\alpha)$ for every agent type α. This strategy specifies, for each state s, whether to exit and what action to take as a function of the value

[8] While our formal definition does not allow for the action set A_s to be randomized, it is straightforward to encode an equivalent matching rule with a deterministic action set: in state $s = \lfloor \tau \rfloor$, the agent is always offered the actions $\{Accept, Reject\}$. With probability $\tau - \lfloor \tau \rfloor$, the agent's action is ignored and her state incremented, and otherwise she transitions as defined above. We choose the description with a randomized action set for its conceptual clarity.

v for the current development. A strategy is optimal if the continuation value of being in each state satisfies the following Bellman equation:

$$V(s) = \max\left(0, \delta\mathbb{E}_{v\sim G}\left[\max_{a\in A_s}\left\{T_a(s,m)\left(\frac{v-\alpha}{1-\delta}\right) + \sum_{s'\in S} T_a(s,s')V(s')\right\}\right] - c\right). \tag{1}$$

The above equation can be interpreted as follows: the value of being in state $s \in S$ is the maximum of the value of exiting (normalized to zero), and staying. The value of staying is based on choosing the best action $a \in A_s$, and each action determines a probability $T_a(s,m)$ of matching in state s to the current development, as well as a probability $T_a(s,s')$ of transitioning to state s'. If the agent matches with the current development with value $v \sim G$, her value is $\frac{v-\alpha}{1-\delta}$ because she would receive a net benefit (over her current situation) of $v - \alpha$ in each period, and she is expected to be able to enjoy this benefit for $\frac{1}{1-\delta}$ periods. We multiply the term within the expectation by δ and subtract c because agents only receive value if they pay the participation cost and do not have a life event in the current period. Appendix B formally defines the MDP facing each agent, and shows that for any $\delta < 1$, a solution $V(s)$ to the above Bellman Equation exists and is unique.

3.4 Outcomes

An *outcome* specifies a distribution of payoffs for each agent type. Formally, an **outcome** E specifies an **outcome function** $P^E : [\underline{\alpha}, \overline{\alpha}] \times [\underline{v}, \overline{v}] \to [0,1]$ and a **waiting time function** $t^E : [\underline{\alpha}, \overline{\alpha}] \to [0, \infty)$, where $P^E(\alpha, v)$ specifies the probability that an agent with outside value α matches to a development for which her value is at most v, and $t^E(\alpha)$ specifies the expected number of periods she participates in the mechanism before leaving or being matched. For a given matching rule R, every strategy profile Σ induces an unique outcome, which refer to as the **outcome corresponding to** R and Σ.[9] Proposition 1 implies that for outcomes that correspond to *optimal* strategy profiles, the waiting time function t^E can be expressed in terms of the outcome function.[10] For any such outcome E, it suffices to specify only the outcome function, and our proofs in the Appendix abuse notation and use $E(\alpha, v)$ to refer to the outcome function instead of $P^E(\alpha, v)$.

Given outcome $E = (P^E, t^E)$, define the corresponding

- **allocation function** to be

$$\pi^E(\alpha) = P^E(\alpha, \overline{v}) \tag{2}$$

This specifies the probability that an agent of outside option α matches to *some* development.

[9] See Appendix B for more details about how to compute the outcome induced by the matching rule R, strategy profile Σ, and market primitives.

[10] Specifically, it holds that $t^E(\alpha) = \frac{1}{c(1-\delta)}\left(\int_\alpha^\infty P^E(x,\overline{v})\,dF(x) - \int_{\underline{v}}^{\overline{v}}(v-\alpha)\,dP^E(\alpha,v)\right)$.

- **match rate** to be the fraction of agents who match:

$$\overline{\pi}^E = \mathbb{E}_{\alpha\sim F}[\pi^E(\alpha)]. \tag{3}$$

- **expected utility** function to be

$$u^E(\alpha) = \int_{\underline{v}}^{\overline{v}} (v-\alpha)\, dP^E(\alpha, v) - (1-\delta)ct^E(\alpha). \tag{4}$$

  This specifies each agent's expected total benefit from participation, multiplied by $(1-\delta)$. This scaling factor keeps the utility in the same scale as each period's value and cost, as the expected eligibility time of an agent is $\frac{1}{1-\delta}$ periods.

3.5 Equilibrium

For simplicity, we first rule out the trivial case in which supply is so abundant that the market designer can offer every development to every agent – if this were feasible, then it would be clearly optimal to do so. Define GC to be the outcome when agents play optimally in the MDP induced by the guaranteed choice matching rule.[11] We assume the following for the remainder of the paper.

Assumption 1. *It is infeasible to offer guaranteed choice to all agents ($\mu < \overline{\pi}^{GC}$).*

Under this assumption, an equilibrium can be defined as a matching rule and an optimal strategy profile that exactly clears the market.

Definition 2 (Equilibrium Outcome)**.** *An outcome E is an* **equilibrium outcome** *of matching rule R if it can be expressed as the outcome corresponding to R and a strategy profile Σ, such that*

a) For every α, the strategy $\Sigma(\alpha)$ is optimal for an agent with outside option α.

b) The average match rate of E equals supply-demand ratio: $\overline{\pi}^E = \mu$.[12]

We sometimes refer to E simply as an equilibrium outcome, without mentioning the matching rule R. If E satisfies a) but not b), we refer to it as a **partial equilibrium** *outcome.*

In our model, the exogenous parameters are the market primitives F, G, δ, c, μ, and the designer's choice of mechanism M. The matching rule $R \in M$ is endogenously determined in equilibrium, as are the strategy profile Σ and the outcome E. The pair (R, Σ) corresponds to an equilibrium outcome E if and only if aggregate demand is equal to aggregate supply.[13] When M is a lottery-based matching rule,

[11] Proposition 2 in Appendix C shows that this outcome is unique.

[12] In ruling out the case $\overline{\pi}^E < \mu$, we eliminate mechanisms that intentionally withhold supply. Nevertheless, one can study the effect of withholding supply by performing comparative statics with respect to μ.

[13] This is analogous to the definition of a Walrasian equilibrium. A price vector can arise in Walrasian equilibrium if when agents respond optimally, the market clears. In our model, a matching rule $R \in M$ can arise in equilibrium if when agents respond optimally, the average match rate equals to μ.

the endogenous quantities are characterized by the success probability p, which uniquely determines the matching rule R and the corresponding strategies and outcome. When M is waitlist-based, the endogenous quantities are characterized by the average idle time τ.

3.6 Metrics for Evaluating Equilibria

Below, we define several metrics used to evaluate outcomes. Given outcome E, define the

- **match distribution** F^E to be the distribution of outside options conditional on matching:

$$F^E(\alpha) = \frac{\int_{-\infty}^{\alpha} \pi^E(x) dF(x)}{\int_{-\infty}^{\infty} \pi^E(x)\, dF(x)}. \tag{5}$$

  Thus, $F^E(\alpha)$ is the fraction of matched agents who have outside options no better than α.

- **value per match** $\nu^E(\alpha)$ for each type α to be the expected benefit per unit of housing allocated to type α,

$$\nu^E(\alpha) = \frac{u^E(\alpha)}{\pi^E(\alpha)} \tag{6}$$

  For convenience, define the value per match to be zero when the denominator is zero.

- **utilitarian welfare** W^E to be the aggregate benefit per allocated housing unit over all types:

$$W^E = \mathbb{E}_{\alpha \sim F}[u^E(\alpha)]/\overline{\pi}^E. \tag{7}$$

In the introduction, we discussed two objectives: ensuring that matched individuals receive a desirable development with minimal participation cost, and targeting the most needy individuals. The first of these objectives – which we refer to as *matching* – is captured by the value per match $\nu^E(\alpha)$, while the second – which we refer to as *targeting* – is captured by the match distribution F^E. We now define what it means for one outcome to result in better matching or targeting than another. The definitions have a strong requirement of point-wise or stochastic dominance, but we will show that such relationships exist among the mechanisms we study.

Definition 3. *Let E and E' be arbitrary outcomes. We say that*

- *E* **match dominates** *E' if $\nu^E(\alpha) \geq \nu^{E'}(\alpha)$ for all α.*
- *E'* **targeting dominates** *E if the match distribution of E first-order stochastic dominates the match distribution of E'. That is, $F^{E'}(\alpha) \geq F^E(\alpha)$ for all $\alpha \in (\underline{\alpha}, \overline{\alpha})$.*

4 Results

4.1 Equivalence of Mechanisms

We first show that mechanisms that look very different can achieve equivalent outcomes. In fact, when participation costs are negligible compared to the values of being allocated, all six mechanisms defined in Section 3.2 are equivalent to either independent lotteries or the common lottery.

To state the result formally, we say that mechanisms M and M' are **outcome equivalent** if the set of equilibrium outcomes are equal: $\mathcal{E}^M = \mathcal{E}^{M'}$, where

$$\mathcal{E}^M = \{(P^E, t^E) : E \text{ is an equilibrium outcome of some matching rule } R \in M\}. \tag{8}$$

In other words, there is a one-to-one correspondence between the equilibrium outcomes of the two mechanisms, such that in each pair of equilibrium outcomes, the distribution of matches and the expected waiting times are equal for every agent type.

Theorem 1 (Equivalence of Mechanisms)**.**

a) Independent lotteries is outcome equivalent to the waitlist without deferral.

b) The ticket-saving lottery is outcome equivalent to the waitlist with deferral.

c) When $c = 0$, the ticket-saving lottery, waitlist with deferral and the single-entry lottery are outcome equivalent to the common lottery.

The proof is in Appendix F.1. We give the intuition below.

For part a), think of the following implementation of independent lotteries: instead of asking agents to enter the lottery and *then* selecting winners, select winners among *all* eligible agents and offer these winners the opportunity to match to the development. This procedure is equivalent because the agents who choose to enter the lottery in the first description are exactly those who will accept the development in the second. Therefore, in both independent lotteries and the waitlist without deferral, agents are periodically offered the chance to match to the current development. In the waitlist without deferral, this occurs approximately every τ periods. Under independent lotteries, this occurs independently in each period, with some probability p. However, what matters to each agent is not the *distribution* of when she will next receive an offer, but rather the *probability* that she will receive at least one more offer (call this q).[14] This probability determines which developments she will accept, and thus her probability of matching. Because both mechanisms match the same number of agents, it follows that any value of q that arises in equilibrium of independent lotteries must also be an equilibrium of

[14] In both mechanisms, the number of offers received by an agent who decides to reject all offers follows a geometric distribution on $\{0, 1, 2, \ldots\}$ with a mean of $\frac{q}{1-q}$.

the waitlist without deferral – and that in these equilibria, each agent sets the same threshold when determining which developments to accept, and matches with the same probability.

For part b), first consider the waitlist with deferral. Because the equilibrium is stationary, once an agent is offered one development, that agent will be offered every future development. Therefore, agents in the waitlist with deferral must wait (for approximately τ periods) before playing guaranteed choice. Now consider the ticket-saving lottery, and recall that regardless of when it is used, each ticket wins with some fixed probability p. Consider a variant of the ticket-saving lottery in which each ticket, when given to a participant, is visibly labeled as a "winning ticket" (with probability p) or a "losing ticket" (with probability $1-p$). It is clear that in this variant, agents must wait a geometric number of periods before receiving a winning ticket, and from that point onward, will set an acceptance threshold as in guaranteed choice (and use all tickets when entering). As in part a), the distribution of idle time does not matter to an agent, but only the probability q that she becomes eligible before her life event. Because both the waitlist with deferral and the ticket-saving variant match the same number of agents, they must have the same value of q, and therefore lead to equivalent outcomes.

Of course, in the actual ticket-saving lottery, the labels of "winning ticket" and "losing ticket" are revealed only after the tickets are used. But this knowledge does not an change agent's optimal strategy, because when she does not hold a winning ticket, her actions do not matter.[15] Therefore, she should always behave as though she holds a winning ticket: set an acceptance threshold as in guaranteed choice, and use all tickets upon seeing such a development.

For part c), note that when participation cost $c = 0$, delays are costless, so the delayed guaranteed mechanisms in part b) are equivalent to selecting a random subset of agents to face guaranteed choice, which is the definition of the common lottery. Similarly, the single-entry lottery effectively selects some agents (those with winning tickets) to play guaranteed choice, while eliminating others. Although agents in a single-entry lottery do not know whether they have been selected until after entering a lottery, the reasoning from part b) implies that they will set acceptance thresholds as though they held a winning ticket.

Note that the arguments for part c) no longer hold when $c > 0$. In particular, the waitlist with deferral is no longer equivalent to the common lottery, as agents in the former must incur a significant cost in order to be given the option to match to a development. Moreover, agents in the single-entry lottery have an incentive to use their ticket early, so that they can exit and stop incurring participation costs.

[15] One crucial but subtle point is that agents in a ticket-saving lottery are never worse off than they are upon entry, and therefore any agent who chooses to participate will not quit even if told that she does not hold a winning ticket.

4.2 Achieving Effective Matching

In our model, effective matching requires that agents are matched to items that are a good fit, at minimal participation cost. The common lottery accomplishes both of these goals. Agents with good lottery numbers have high continuation values, and therefore an incentive to be selective; agents with poor lottery numbers learn immediately that they will not match, and therefore do not incur participation costs while clinging to a false hope. In fact, Theorem 2 shows that the common lottery not only match dominates all the other mechanisms we study, but also converges to the best possible outcome in terms of matching when agents are eligible for many periods ($\delta \to 1$). We think that this is a reasonable limit to study for the application of affordable housing in New York: there are lotteries for over 70 new developments each year, so if agents expect to remain eligible and interested for at least 18 months, then δ exceeds 0.99.

For the asymptotic limit to be defined, we require that values are bounded ($\overline{v} < \infty$).

Definition 4. *When $\overline{v} < \infty$, define* **perfect matching** *(PM) to be the outcome in which agents of all types match with probability μ, and conditioned on matching, have the highest possible value $\overline{v}$ for their assignment, with negligible participation cost: $P^{PM}(\alpha, v) = \mu \mathbb{1}(v \geq \overline{v} > \alpha)$, with $t^{PM}(\alpha) = 0$.*

Definition 5. *A sequence of equilibrium outcomes E^n converges to outcome E if the outcome functions converge point-wise: that is, $P^{E_n}(\alpha, v) \to P^E(\alpha, v)$ for all α, v.*[16]

Theorem 2 (Match Dominance of the Common Lottery)**.**

a) The unique equilibrium outcome of the common lottery match dominates any equilibrium outcome of independent lotteries, single entry lottery, ticket-saving lottery, waitlist with deferral, and waitlist without deferral.

b) When $\overline{v} < \infty$, as $\delta \to 1$, the equilibrium outcome of the common lottery converges to perfect matching, which match dominates any equilibrium outcome.

The proof of Theorem 2 is in Appendix F.3. Part a) is based on the structural results derived in the proof of Theorem 1, which allow us to derive explicit expressions for the value per match in each mechanism. Part b) is based on showing that when δ is high, agents will only accept developments close to the maximum value of $\overline{v}$. Moreover, any utility loss they incur while waiting for such a development is negligible compared to the many periods they get to enjoy their apartment after matching. Finally, almost every agent who wins the common lottery eventually matches, so the probability of matching is nearly the same (and equal to μ) for all agents.

[16] Definition 5 does not mention waiting times, because for any equilibrium outcome E, the waiting time function t^E is determined by the outcome function P^E by Proposition 1. Convergence of outcome function as $\delta \to 1$ does not imply convergence of the waiting time function, but does imply convergence of the scaled participation cost: that is, $(1 - \delta_n)c|t^{E_n}(\alpha) - t^E(\alpha)| \to 0$. This is sufficient to imply that utilities, the value per match, and the match distribution all converge point-wise: $u^{E_n}(\alpha) \to u^E(\alpha)$, $\nu^{E_n}(\alpha) \to \nu^E(\alpha)$ and $F^{E_n}(\alpha) \to F^E(\alpha)$ for all α.

When participation cost is negligible, the equilibrium outcomes of the single-entry lottery, waitlist with deferral, and ticket-saving lottery also converge to perfect matching by Theorem 1. When $c > 0$, Appendix F.3 shows that the single-entry lottery converges to perfect matching, but the waitlist with deferral and ticket-saving lottery do not. The reason is that under these mechanisms, high values of δ result in long expected wait times, causing some agents not to participate and the remainder to experience significant participation costs. By contrast, in the single-entry lottery, agents can leave as soon as they use their ticket, so the participation cost they incur is minimal.

4.3 Tradeoff Between Matching and Targeting

An anonymous mechanism cannot benefit agents with the greatest need without allocating also agents with less need, because low-need agents can always copy the behavior of high-need agents. This is formalized in Proposition 1, which states that in any equilibrium outcome, the utility of an agent is equal to the integral of the match rate for all agents with better outside options. Hence, in order to increase the utility of agents with poor outside options, it is necessary to increase the match probability of those with better outside options.

Proposition 1. *For any partial equilibrium outcome E, the allocation function $\pi^E(\alpha)$ is weakly decreasing, and the expected utility function is given by*

$$u^E(\alpha) = \int_\alpha^\infty \pi^E(x)\, dx. \tag{9}$$

Proposition 1 can be used to show that when there are many low-need agents, there is a tradeoff between providing high-quality matches and targeting need effectively.

Definition 6. *There are* **many low-need individuals** *if $\overline{\alpha} \geq \overline{v}$ and the density of outside options f is increasing on $(\underline{\alpha}, \overline{\alpha})$.*

Theorem 3 (Matching vs Targeting)**.** *Let E and E' be equilibrium outcomes. If E match dominates E', and if there are many low-need individuals (see Definition 6), then E' targeting dominates E.*

When matching and targeting are in conflict with one another, it is natural to wonder which objective is more important. Theorem 4 shows that the answer to this question depends on both the shape of F and the support of F and G.

Definition 7.

F has a **light left tail** *if $F(x)/f(x)$ is weakly increasing in the domain $(\underline{\alpha}, \overline{\alpha})$.*

F has a **heavy left tail** *if $F(x)/f(x)$ is weakly decreasing in the domain $(\underline{\alpha}, \overline{\alpha})$.*

Examples of light tailed distributions include the uniform, the normal, and the Gumbel distribution, as well as truncated versions of these distributions. The exponential distribution has a constant hazard rate (and thus is the dividing line between light and heavy-tailed distributions). The Pareto and the log-normal distributions are examples of heavy-tailed distributions.[17]

Theorem 4 (Welfare Comparisons)**.** *Let E and E' be equilibrium outcomes. If $\overline{\alpha} \geq \overline{v}$ and E' targeting dominates E, then the following hold:*

a) If F has a light left tail, then $W^E \geq W^{E'}$.

b) If F has a heavy left tail, then $W^E \leq W^{E'}$.

Theorem 3 and 4 b) together imply that when there are many low-need individuals and the outside option distribution has a heavy left tail, optimizing for match quality is detrimental to aggregate welfare. In this case, a common lottery might lower utilitarian welfare compared to independent lotteries.

For affordable housing allocation in New York, we believe that this is not the case. The reason is that those who qualify for housing already fall within a narrow income range, so it seems reasonable that many agents have similar outside options. Moreover, the developments being allocated in New York are newly-constructed and designed to be attractive to market-rate renters, so we expect that most eligible applicants consider many of these units preferable to their current living situation. Hence, the setting in New York may be better approximated by the conditions of Theorem 5, which states that if outside options are light-tailed or sufficiently poor[18], then utilitarian welfare is maximized by prioritizing good matching. Under these conditions, it follows from Theorem 2 b) that when δ is high, a common lottery achieves near-optimal utilitarian welfare.

Theorem 5 (Optimality of Perfect Matching)**.** *Perfect matching achieves weakly higher utilitarian welfare than any equilibrium outcome if either*

a) F has a light left tail, or

b) $\overline{v} - \overline{\alpha} \geq \overline{\alpha} - \underline{\alpha}$.

We interpret Theorems 4 and 5 to mean that matching is more important than targeting whenever outside options follow a light-tailed distribution or are sufficiently low. Figure 1 reinforces this point[19]: it displays the welfare difference between common and independent lotteries, as the shape and support of the outside option distribution vary. The common lottery is superior unless the outside option

[17] Typically, the tail of a distribution refers to the right tail. In Definition 7, we refer to the left tail. This is because the agents with the highest value for being matched are those with the worst outside options.

[18] In particular, condition b) of Theorem 5 is that the value of matching any agent well is greater than the difference in need between any two agents.

[19] The findings in Figure 1 are in alignment with our interpretation, despite the fact that the example violates the conditions of Theorems 3, 4 and 5: G follows a normal distribution, for which $\overline{v} = \infty$; and F follows a (negated) Weibull distribution, for which the density is not increasing whenever $\log(k) > 0$.

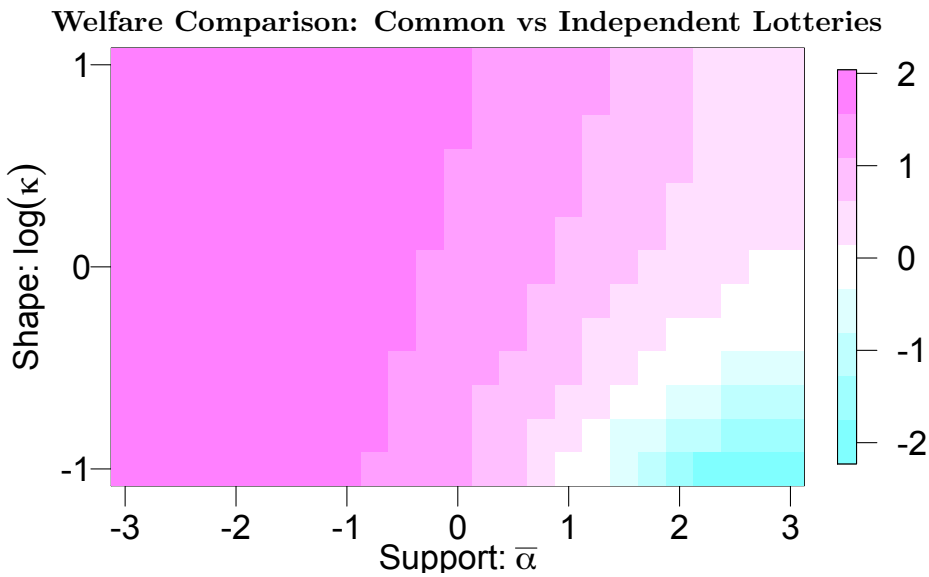

Figure 1: Heatmap of the difference in welfare between the common and independent lotteries, varying the shape and support of the outside option distribution F while holding other parameters fixed.[21] Positive values (pink) correspond to higher welfare under the common lottery. Moving from bottom to top, the tail of F becomes lighter, with $\log(\kappa) = 0$ corresponding to the exponential distribution. Moving from right to left, the outside option distribution shifts downward. A common lottery attains higher welfare whenever outside options are light tailed (top region) or sufficiently poor (left region). A difference of 2 means that the improvement in per-match welfare is equal to two standard deviations of the outside option distribution G.

distribution is heavy-tailed and outside options are good (the lower right region). Furthermore, the differences are significant: a welfare difference of 1.5 implies that the difference between the two mechanisms is equal to the difference between matching agents to random developments and matching them to something that they prefer to 93% of developments.

The proof of Theorem 3 and a more general version of Theorem 4 are in Appendix F.5. The proof of Theorem 5 is in Appendix F.6. All make use of Proposition 1, which we prove in Appendix D.

4.4 Achieving Effective Targeting

Although a common lottery may not be effective at targeting need, the same is true of independent lotteries and the waitlist without deferral. In fact, Proposition 5 in Appendix E shows that in some cases, these approaches result in no targeting at all! Even when there are many low need individuals – in which case Theorems 2a) and 3 jointly imply that these approaches targeting dominate a common lottery – there are generally more effective ways to target need.

A simple approach to achieve good targeting regardless of distributional assumptions is as follows:

[21] We take G to be $Normal(0,1)$, $\mu = 0.1$, $\delta = 0.99$, and $c = 0$. The outside option distribution F is a (negated) Weibull distribution, given by $F(\alpha) = \exp\left(-\left(\Gamma(1+\frac{1}{\kappa})(\overline{\alpha}-\alpha)\right)^{\kappa}\right)$. This distribution has expected value $\overline{\alpha}-1$. It is light-tailed for $\kappa > 1$ and heavy-tailed for $\kappa < 1$, with $\kappa = 1$ corresponding to an exponential distribution.

artificially increase participation cost until it is possible to match every agent who is willing to participate. In practice, this may mean requiring agents to undergo a costly ordeal to remain eligible, such as to complete endless paperwork or to physically line up at a central office every week.[22] While we do not believe that this is a good solution for allocating affordable housing, such practices may be reasonable in settings with loose eligibility criteria, such as in the allocation of discounted Broadway tickets.

Precisely speaking, participation cost c is said to be **market clearing** if under this participation cost, the average match rate under the guaranteed choice matching rule is equal to the supply-demand ratio μ. We show in Appendix F.7 that a market clearing cost always exists; although market clearing costs may not be unique, there always exists a *highest* market clearing cost $\overline{c} < \infty$. Define the **costly guaranteed choice** outcome be the guaranteed choice outcome under the highest market clearing cost $\overline{c}$. When participation cost is increased to $\overline{c}$, all of the mechanisms studied in this paper implement the costly guaranteed choice outcome.

Theorem 6 shows that costly guaranteed choice always targeting dominates the common lottery. Furthermore, it converges to the best possible outcome in terms of targeting when agents are long-lived ($\delta \to 1$) and values are bounded ($\overline{v} < \infty$).

Definition 8. *Define* **perfect targeting** *(PT) to be the outcome in which agents with outside option* $\alpha \leq F^{-1}(\mu)$ *are matched with certainty, and no other agents are matched:* $P^{PT}(\alpha, v) = \mathbb{1}(\alpha \leq F^{-1}(\mu))\mathbb{P}_{v' \sim G}(v \geq v' | v' \geq F^{-1}(\mu))$*, and* $t^{PT}(\alpha) = \frac{1}{c(1-\delta)}\mathbb{1}(\alpha \leq F^{-1}(\mu))\mathbb{E}_{v' \sim G}[v' - F^{-1}(\mu) | v' \geq F^{-1}(\mu)]$.[23]

Theorem 6 (Targeting Dominance of Costly Guaranteed Choice)**.**

a) The costly guaranteed choice outcome target dominates any equilibrium outcome of the common lottery, single-entry lottery, ticket-saving lottery, and waitlist with deferral.

b) When $\overline{v} < \infty$*, as* $\delta \to 1$*, the sequence of costly guaranteed choice outcomes converges to perfect targeting, which targeting dominates any equilibrium outcome.*

The proof is given in Appendix F.7. Part a) is based on structural results derived in the proof of Theorem 1. The Appendix also shows that costly guaranteed choice targeting dominates independent lotteries and the waitlist without deferral when the value distribution G is light tailed. For part b),

[22] Our model assumes that the participation cost c is identical for all agents, and therefore that the agents who participate will be those with the greatest need (worst outside options). In practice, an important caveat when adding participation costs is that the designer should ensure that the type of cost added does not disproportionately impact high-need agents. For example, if wealthier applicants are systematically more adept at filling out forms, or more able to take an afternoon off of work, then adding paperwork to the application process or mandating physical presence may have the opposite of the desired effect. Alatas et al. (2016) and Deshpande and Li (2017) explore this concern using (quasi-)random experiments to empirically estimate the effect of certain types of friction on various sub-populations. We discuss this point further in Section 5.1.

[23] This waiting time ensures that agents with outside option $\alpha = F^{-1}(\mu)$ are indifferent about whether to participate.

we show that if agents remain in the system for many periods, then almost everyone who chooses to participate will eventually find a development that they are willing to accept. Moreover, the agents who choose to participate will be those with the greatest need. Since the most needy are matched with near-certainty and everyone else is not matched at all, this is the best possible outcome in terms of targeting.

Hence, it is rarely a good idea to use independent lotteries with low participation cost: if matching is more important, the designer should adopt a common lottery. If targeting is more important, the designer should increase participation cost so that low need agents do not participate at all.

5 Discussion

In this paper, we argue that two common systems for allocating affordable housing – independent lotteries and a waitlist in which applicants lose priority after declining an offer – incentivize prospective tenants to accept buildings that are only marginally better than their outside options. The resulting allocation is inefficient, in that many or all agents could be simultaneously made better off. We discuss several reforms that could improve the quality of the assignment, including limiting lottery entry, allowing applicants to keep their position in a waitlist after rejecting an offer, allowing applicants to save and combine lottery tickets, and (perhaps most simply) using a common lottery to determine priority for all buildings.

We believe that using a common lottery is a practical solution that could easily be adapted to accommodate institutional details not captured in our model. For example, in New York, eligibility is building-specific,[24] and certain groups – such as city employees or neighborhood residents – get priority for a certain number of units in each building. These practices could be maintained under a common lottery: treat units that give priority to specific groups as separate buildings, and for each building, offer units to those eligible for them, in the order determined by the (universal) ranking of applicants.[25]

There are of course many other ways in which our model oversimplifies reality. We conclude by discussing the robustness of our findings when our modeling assumptions are relaxed.

5.1 Multi-dimensional Agent Types

Consider a richer model in which agents differ not only in their outside option, but also in their value distribution, participation cost, and expected eligibility time. The type of an agent is represented by a tuple $\theta = (\alpha, G, c, \delta)$, and is distributed according to distribution Θ. This is a straightforward

[24] Buildings have different income ranges, and different apartment types (studio, one bedroom, two bedroom, etc).

[25] In the event that there are not enough applicants from a group to fill all units that give this group priority, the remaining units could be offered to other applicants.

generalization of our current framework, and the Bellman equation for the optimal strategy of each agent remains the same as in (1). The only difference is that we must define the outcome E as a function of the tuple (θ, v), instead of only (α, v).

In this model, our result on the matching efficacy of the common lottery (Theorem 2) continues to hold, as the proof is based on analyzing each agent type separately. For a similar reason, the equivalence results continue to hold if agents are homogeneous in δ. However, if the expected eligibility time $\frac{1}{1-\delta}$ varies across agents, then the equivalences break down: agents who are eligible for more periods are more likely to match in waitlist-based mechanisms, whereas short-lived agents prefer lottery-based mechanisms.

Analysis of targeting becomes nuanced under such a model. First, it is unclear whom to target: does someone with very high value for one development but not another have greater or less need than someone with a moderate value for all developments? Second, even if the market designer identifies which types to target, the answer to the question of how to target effectively will depend on the distribution of types. For example, if it happens that agents with the worst outside options also have higher participation costs, then it is possible that a common lottery simultaneously match dominates and targeting-dominates independent lotteries, even if there are many agents with low need. The reason is that independent lotteries may require participation for many periods before matching (and therefore deter entry by those with high participation costs), whereas the winners of a common lottery are matched very quickly.

5.2 Vertical Differentiation of Developments

Our model assumes purely horizontal differentiation between developments, so that all are equally popular in the aggregate. One might consider a model in which development j has quality q_j, and values are distributed as $v_{ij} = q_j + \epsilon_{ij}$, where $\epsilon_{ij} \sim G$ is the horizontal component of preferences. It is much harder to analyze an equilibrium under this model because it is no longer stationary: the type distribution of agents remaining in the system depends on the history of development qualities, and agents must reason about how this type distribution will evolve when making decisions.

Nevertheless, under such a model, we would still expect independent lotteries and waitlist without deferral to yield low match quality when supply is scarce, because agents' acceptance thresholds on the added value of a development $(v_j - \alpha)$ will still equal their continuation value, which is approximately zero if μ is small. Meanwhile, we expect the common lottery to result in better match quality, as agents offered a building j for which their idiosyncratic term ϵ_{ij} is small could wait for a building of similar quality that was better-suited to their needs. Furthermore, the matching/targeting tradeoff described in Theorem 3, 4, and 5 would continue to hold, as the proofs rely only on anonymity of the mechanism,

and not on any assumptions about the nature of the dynamic game being played by agents.[26]

5.3 Partially Observable Outside Options

In practice, observable information is often used to prioritize certain agents. This can be captured by an extension of our model in which agents are classified into groups based on characteristics such as income, family status, current residence, etc. Within each group k, the arrival rate of agents is λ_k per period, and the primitives F, G, c and δ may also be indexed by k. A natural extension of the common lottery to this setting is as follows: assign a priority to each group, and a lottery number to each agent; agent-level priorities are induced by the group priorities, and the lottery numbers are used to break ties. Under this mechanism, it will be the case that high-priority groups can choose whatever they want, and low priority groups are never matched; agents in borderline priority groups are selected based on their lottery numbers.

Our results imply that this extension of the common lottery would work well if priority is given to groups with the lowest average outside option. In particular, if agents remain eligible for many periods, $\overline{v}$ is the same for each group and finite, and outside options are light-tailed within each group, one can show that this version of the common lottery assigns every matched agent to a development where her value is close to the upper bound $\overline{v}$, thus generalizing Theorems 2 b). Moreover, this mechanism achieves near-optimal utilitarian welfare among all stationary mechanisms that are anonymous within each group, thus generalizing Theorem 5 a).

References

Agarwal N, Ashlagi I, Rees M, Somaini P, Waldinger D (2017) An empirical framework for sequential assignments: The allocation of deceased donor kidneys. Technical report, MIT.

Akbarpour M, Li S, Gharan SO (2014) Dynamic matching market design. ArXiv preprint arXiv:1402.3643.

Alatas V, Purnamasari R, Wai-Poi M, Banerjee A, Olken BA, Hanna R (2016) Self-targeting: Evidence from a field experiment in indonesia. *Journal of Political Economy* 124(2):371–427.

Ashlagi I, Burq M, Jaillet P, Manshadi V (2018) On matching and thickness in heterogeneous dynamic markets. Available at SSRN: `https://ssrn.com/abstract=3067596`.

Baccara M, Lee S, Yariv L (2018) Optimal dynamic matching. Available at SSRN: `https://ssrn.com/abstract=3198136`.

[26]In particular, Theorems 3, 4 and 5 also apply in settings where agents have information about future developments, where the designer delays allocation of some units, or where the designer observes agent values and uses this information to determine the allocation.

Bertsekas DP (2012) *Dynamic Programming and Optimal Control, Vol. II* (Athena Scientific), 4th edition, ISBN 1886529442.

Blackorby C, Donaldson D (1988) Cash versus kind, self-selection, and efficient transfers. *The American Economic Review* 691–700.

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

Chakravarty S, Kaplan TR (2013) Optimal allocation without transfer payments. *Games and Economic Behavior* 77(1):1–20.

Clementi F, Gallegati M (2005) Paretos law of income distribution: Evidence for germany, the united kingdom, and the united states. *Econophysics of wealth distributions*, 3–14 (Springer).

Condorelli D (2012) What money can't buy: Efficient mechanism design with costly signals. *Games and Economic Behavior* 75(2):613–624.

Damiano E, Lam R (2005) Stability in dynamic matching markets. *Games and Economic Behavior* 52(1):34–53.

Deshpande M, Li Y (2017) Who is screened out? application costs and the targeting of disability programs. Technical report, National Bureau of Economic Research.

Dickerson JP, Procaccia AD, Sandholm T (2012) Dynamic matching via weighted myopia with application to kidney exchange. *AAAI*.

Doval L (2018) A theory of stability in dynamic matching markets. Technical report, Working paper.

Geyer J, Sieg H (2013) Estimating a model of excess demand for public housing. *Quantitative Economics* 4(3):483–513.

Glaeser EL, Luttmer EFP (2003) The misallocation of housing under rent control. *American Economic Review* 93(4):1027–1046.

Gurvich I, Ward A (2014) On the dynamic control of matching queues. *Stochastic Systems* 4(2):479–523.

Hartline JD, Roughgarden T (2008) Optimal mechanism design and money burning. *Proceedings of the 40th annual ACM symposium on Theory of computing*, 75–84, STOC '08 (New York, NY, USA: ACM).

Hoppe HC, Moldovanu B, Sela A (2009) The theory of assortative matching based on costly signals. *The Review of Economic Studies* 76(1):253–281.

Kaplan EH (1987a) Analyzing tenant assignment policies. *Management Science* 33(3):395–408.

Kaplan EH (1987b) Tenant assignment policies with time-dependent priorities. *Socio-Economic Planning Sciences* 21(5):305–310.

Kaplan EH (1988) A public housing queue with reneging and task-specific servers. *Decision Sciences* 19(2):383–391.

Kennes J, Monte D, Tumennasan N (2014) The day care assignment: A dynamic matching problem. *American Economic Journal: Microeconomics* 6(4):362–406.

Kurino M (2009) Credibility, efficiency, and stability: A theory of dynamic matching markets. *Jena economic research papers, JENA* 7:41.

Leshno JD (2015) Dynamic matching in overloaded waiting lists. Technical report, Columbia University.

Nichols AL, Zeckhauser RJ (1982) Targeting transfers through restrictions on recipients. *The American Economic Review* 72(2):372–377.

Pereyra JS (2013) A dynamic school choice model. *Games and economic behavior* 80:100–114.

Prendergast C (2017) The allocation of food to food banks. Technical report, University of Chicago.

Roth AE, Sönmez T, Ünver MU (2004) Kidney exchange. *The Quarterly Journal of Economics* 119(2):457–488.

Roth AE, Sotomayor M (1992) Two-sided matching. *Handbook of game theory with economic applications* 1:485–541.

Schummer J (2016) Influencing waiting lists. Technical report, Working paper, Kellogg School of Management, 2015. 41.

Sieg H, Yoon C (2017) Waiting for affordable housing. Technical report, University of Pennsylvania.

Sönmez T, Ünver MU (2011) Matching, allocation, and exchange of discrete resources. *Handbook of social Economics* 1(781-852):2.

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

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

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

Thakral N (2016) The public-housing allocation problem. Technical report, Harvard University.

Van Ommeren JN, Van der Vlist AJ (2016) Households' willingness to pay for public housing. *Journal of Urban Economics* 92:91–105.

Waldinger D (2018) Targeting in-kind transfers through market design: A revealed preference analysis of public housing allocation. Technical report, MIT.

A Alternate Payout Models

In this section, we give two alternative ways to formulate our model in Section 3 that are mathematically equivalent. These alternative formulations enrich the interpretation of our results.

A.1 One-Time Payoffs

In the first formulation, payoffs are incurred upon matching or exit, instead of in each period. This model is more natural for allocating Broadway tickets or other experience goods such as hiking, camping, and hunting permits. The modified timeline is as follows:

1. **Arrival and Participation Choice**: As in Section 3.
2. **Life Event**: Every agent exits exogenously with probability $1-\delta$ and receive their outside option α_i.
3. **New Development and Matching:** As in Section 3.
4. **Payoff**: Every matched agent i exits with a one-time payoff of v_{ij}. The unmatched agents continue to the next period.

Agents seek to maximize their expected payout before matching or exiting, minus any participation costs. The updated Bellman Equation is as follows.

$$V(s) = \max\left(0, \delta\mathbb{E}_{v\sim G}\left[\max_{a\in A_s}\left\{T_a(s,m)(v-\alpha) + \sum_{s'\in S} T_a(s,s')V(s')\right\}\right] - c\right). \tag{10}$$

The only change from the original Bellman Equation (1) is that there is no longer a multiplicative factor of $\frac{1}{1-\delta}$ before the $(v-\alpha)$ term, which does not change the mathematical structure. Correspondingly, we remove the $(1-\delta)$ scaling term in an agent's expected utility, so $u^E(\alpha) = \int_{\underline{v}}^{\overline{v}}(v-\alpha)\,dP^E(\alpha,v) - ct^E(\alpha)$. The definitions for outcome, match distribution, value per match, and utilitarian welfare are unchanged. All the results are unchanged, except that the asymptotic results in which $\delta \to 1$ also require scaling c so that $\frac{c}{1-\delta}$ is bounded.

A.2 Reward for Voluntary Exit

Instead of incurring a participation cost c for each period before exiting, agents get a one-time bonus[27] of $r := \frac{c}{1-\delta}$ for voluntarily exiting, and the outside options are all shifted downward by c.

To see that this formulation is equivalent, note that a function $V(s)$ satisfies the Bellman (1) if and

[27] Instead of a one-time bonus, agents who voluntarily exit can equivalently receive a subsidy of c per period until their life event.

only if the function $\tilde{V}(s) := V(s) + r$ satisfies

$$\tilde{V}(s) = \max\left(r, \delta\mathbb{E}_{v\sim G}\left[\max_{a\in A_s}\left\{T_a(s,m)\left(\frac{v-(\alpha-c)}{1-\delta}\right) + \sum_{s'\in S} T_a(s,s')\tilde{V}(s')\right\}\right]\right), \tag{11}$$

which is the Bellman equation for the formulation with exit reward r, no participation cost, and outside option α shifted down by c.

B Formal Definition of Matching MDP

Given a matching rule $R = (S, D, A, T)$ and an outside option α, value distribution G, persistence δ, and participation cost c, a **matching MDP** $\Psi(R) = (S', A', T', \Gamma)$ is a Markov Decision Process with the following parameters:

- State space $S' = S \cup (S \times \mathbb{R}) \cup \{m, e\}$. The states $\{m, e\}$ are terminal states, corresponding respectively to matching and to exiting without a match.
- Action set $A' = \bigcup_{s\in S'} A'_s$. $A'_s = \{l, r\}$ for every state $s \in S$, where l corresponds to voluntarily leaving and r to remaining; $A'_{(s,v)} = A_s$ for every state $(s, v) \in S \times \mathbb{R}$.
- Transition probability function $T' : S' \times A' \times S' \to \mathbb{R}$ and reward function $\Gamma : S' \times A' \to \mathbb{R}$ are as follows.
  - If the current state is $s \in S$, the action l results in transition to state e with no reward, and the action r results in transition with probability $1 - \delta$ to e with reward $-c$, and transition with probability δ to (s, v) with reward $-c$, where v is a new draw from G.
  - If the current state is (s, v), the action $a \in A_s$ results in transition with probability $T_a(a, m)$ to state m with reward $\frac{v-\alpha}{1-\delta}$, and transitions with probability $T_a(s, s')$ to state $s' \in S$ and no reward.

A **strategy** σ to the above MDP is represented by functions $a : S \times \mathbb{R} \to A$ and $b : S \to \{0, 1\}$, where $a(s, v) \in A_s$ is what action to take in state (s, v) and $b(s)$ is whether to take action r in state s.

Every strategy $\sigma = (a, b)$ defines a Markov chain with state space $S'' = S \cup \{m, e_1, e_2\}$, where m, e_1 and e_2 are absorbing. e_1 corresponds to a voluntary exit and e_2 to a forced exit due to the life event. The transition probabilities are $p_{se_1} = 1 - b(s)$, $p_{se_2} = (1 - \delta)b(s)$, and $p_{ss'} = \delta b(s)E_{v\sim G}[T_{a(s,v)}(s, s')]$ for all $s \in S$ and $s' \in S \cup \{m\}$. This is an absorbing Markov chain with a countable state space, in which the chance of transitioning to an absorbing state from any state is at least $1 - \delta > 0$. Hence, given the initial state distribution D, the expected number of times a transition occurs between each pair of states $(s, s') \in S'' \times S''$ before the chain reaches an absorbing state is well-defined. Denote this quantity by $\chi(s, s')$. The **outcome function** and **waiting time function** corresponding to this

strategy for this agent type α are given by:

$$P(\alpha, v) := \sum_{s:p_{sm}>0} \frac{\chi(s,m)}{p_{sm}} E_{v'\sim G}[T_{a(s,v')}(s,m)\mathbb{1}(v' \leq v)]. \tag{12}$$

$$t(\alpha) := \sum_{s\in S} \sum_{s'\in S\cup\{m,e_2\}} \chi(s,s') \tag{13}$$

A strategy (a, b) is optimal if and only if

$$a(s,v) \in \arg\max_{a\in A_s}\{Q(s,v,a)\},$$
$$\text{and } b(s) \in \arg\max\left(0, \delta\mathbb{E}_{v\sim G}[Q(s,v,a(s,v))] - c\right),$$

where $Q(s,v,a) = T_a(s,m)\left(\frac{v-\alpha}{1-\delta}\right) + \sum_{s'\in S} T_a(s,s')V(s')$ is the term inside the max of the Bellman equation (1), and $V(s)$ is the unique solution to (1).

The claim the Bellman Equation (1) has an unique solution $V(s)$ whenever $\delta < 1$ follows from Proposition 1.6.1 of Bertsekas (2012). The claim that $V(s)$ represents the maximum attainable continuation value in state s from any strategy follows from Proposition 1.6.2 of Bertsekas (2012). The conditions needed to apply these two propositions are verified in Lemma 1 below.

Definition 9. *Let $R(S)$ denote the set of functions $S \to \mathbb{R}$ and $B(S) \subseteq R(S)$ denote the set of functions with a finite sup-norm. A mapping $\Pi : R(S) \to R(S)$ is* monotonic *if for any two functions $J, J' \in R(S)$, $J \leq J'$ implies that $\Pi J \leq \Pi J'$, where comparisons of functions are defined point-wise. It is a* contraction with modulus ρ *if $J \in B(S)$ implies $\Pi J \in B(S)$, and for any $J, J' \in B(S)$, $\|\Pi J - \Pi J'\| \leq \rho\|J - J'\|$, where $\|\cdot\|$ is the sup-norm, with $\|J\| := \sup_{s\in S} J(s)$.*

Lemma 1. *Building on the notation of Definition 9, define the mapping $\Pi : R(S) \to R(S)$ where $(\Pi J)(s)$ is given by the right hand side of the Bellman Equation (1) with V replaced by J. For a given strategy $\sigma = (a, b)$, define the mapping $\Pi_\sigma : R(S) \to R(S)$ where*

$$(\Pi_\sigma J)(s) = b(s)\left(\delta\mathbb{E}_{v\sim G}\left[T_{a(s,v)}(s,m)\left(\frac{v-\alpha}{1-\delta}\right) + \sum_{s'\in S} T_{a(s,v)}(s,s')J(s')\right] - c\right). \tag{14}$$

Both Π and Π_σ are monotonic contraction mappings with modulus $\delta < 1$.

Proof of Lemma 1. The monotonicity of Π_σ follows from its linearity, and that of Π follows from $\Pi(s) = \sup_\sigma \Pi_\sigma(s)$. Define $H(x) = \mathbb{E}_{v\sim G}[\max(v - x, 0)]$. For any $J \in B(S)$,

$$\|\Pi_\sigma J\| \leq \|\Pi J\| \leq \delta\mathbb{E}_{v\sim G}\left[\max\left(\|J|, \frac{v-\alpha}{1-\delta}\right)\right] \leq \delta\left(\|J\| + H(\alpha + (1-\delta)\|J\|)\right) < \infty.$$

Finally, we have $\|\Pi_\sigma J - \Pi_\sigma J'\| \leq \delta\|J - J'\|$ since $\sum_s T_{a(s,v)}(s,s') \leq 1$, and $\|\Pi J - \Pi J'\| \leq \sup_\sigma \|\Pi_\sigma J -$

$\Pi_\sigma J'\| \leq \delta \|J - J'\|$. $\square$

C Characterizing Outcomes under Independent Lotteries

For simplicity, we use $E(\alpha, v)$ instead of $P^E(\alpha, v)$ to denote the outcome function of equilibrium outcome E. Define

$$H(x) = \mathbb{E}_{v\sim G}[\max(0, v - x)] = \int_x^\infty \overline{G}(y)\, dy. \tag{15}$$

This is continuous and strictly decreasing for $x < \overline{x}$ and identically zero for $x \geq \overline{x}$. Hence, its inverse $H^{-1}(x)$ is well-defined for $x > 0$. For convenience, define $H^{-1}(0) = \overline{v}$.

Proposition 2 (Optimal Strategy under Independent Lotteries)**.** *Under independent lotteries with success probability p, an optimal strategy for an agent with outside option α is as follows: define $k := \frac{p\delta}{1-\delta}$; if $\alpha > \alpha_0 := H^{-1}(\frac{c}{k})$, then the agent exits immediately upon entry; if $\alpha \leq \alpha_0$, then the agent never voluntarily exits, and participates in lottery j if and only if her value $v_{ij} > \phi(\alpha, p)$, where the threshold $\phi(\alpha, p)$ is the unique solution ϕ to the equation*

$$\phi - \alpha = \max(0, kH(\phi) - c). \tag{16}$$

In all optimal strategies, the outcome is the same,[28] *and is given by*

$$IL_p(\alpha, v) := \begin{cases} \frac{k(\overline{G}(\phi(\alpha,p)) - \overline{G}(v))}{1 + k\overline{G}(\phi(\alpha,p))} & \text{if } \alpha \leq \alpha_0 \text{ and } v > \phi(\alpha, p), \\ 0 & \text{otherwise.} \end{cases} \tag{17}$$

The utility of an agent with outside option α is given by

$$u^{IL_p}(\alpha) = \phi(\alpha, p) - \alpha. \tag{18}$$

Proof of Proposition 2. For any matching rule R and a given type α, if $V(s)$ is the solution to the Bellman Equation (1), and E is any equilibrium outcome of R, we have the identity

$$u^E(\alpha) = (1 - \delta)\mathbb{E}_{s_0\sim D}[V(s_0)], \tag{19}$$

where D is the initial state distribution specified in R. This identity follows from the definition of u^E in section 3.4.

[28] This result assumes that agents who are exactly indifferent between staying and leaving will always stay.

For the independent lotteries matching rule, the Bellman Equation can be simplified as follows. (Since there is only one state, we use V to denote the optimal continuation utility.)

$$V = \max\left(0, \delta \mathbb{E}_{v\sim G}[\max(V, p\frac{(v-\alpha)}{1-\delta} + (1-p)V)] - c\right) = \max\left(0, \delta V + kH(\alpha + (1-\delta)V) - c\right). \quad (20)$$

If $kH(\alpha) - c < 0$, then (20) implies that $V = 0$. If $kH(\alpha) - c \geq 0$, then V satisfies

$$(1-\delta)V = kH(\alpha + (1-\delta)V) - c.$$

Either way, we have the equation

$$(1-\delta)V = \max(0, kH(\alpha + (1-\delta)V) - c), \quad (21)$$

which is identical to (16) with the change of variables $\alpha + (1-\delta)V = \phi(\alpha, p)$. From examining 20, we see that an optimal solution is to exit if $0 > kH(\phi(\alpha, p)) - c$, which is equivalent to $\alpha > \alpha_0$.[29] If $\alpha \leq \alpha_0$, then the agent should join every lottery in which the value $v \geq \alpha + (1-\delta)V = \phi(\alpha, p)$. In fact, every optimal solution must voluntarily exit if $\alpha > \alpha_0$, never exit if $\alpha < \alpha_0$, and while in the system, join a lottery if the value $v > \phi(\alpha, p)$, and not join if $v < \phi(\alpha, p)$. The only flexibility is what to do in the knife edge cases in which $\alpha = \alpha_0$ or if $v = \phi(\alpha, p)$. However, we assumed that agents with $\alpha = \alpha_0$ will participate, and $v = \phi(\alpha, p)$ with probability zero since G is continuous. Hence, the outcome function is unique.

We now derive the expression (17). Suppose that $\alpha \leq \alpha_0$, then the agent never chooses to exit. In each period j, the agent matches with probability $\delta p \overline{G}(\phi(\alpha, p))$, and exits due to the life event with probability $1 - \delta$. Hence, the probability of eventually matching before exiting is

$$\pi^{IL_p}(\alpha) = \frac{p\delta\overline{G}(\phi(\alpha, p))}{1 - \delta + p\delta\overline{G}(\phi(\alpha, p))} = \frac{k\overline{G}(\phi(\alpha, p))}{1 + k\overline{G}(\phi(\alpha, p))}. \quad (22)$$

Now, since the threshold on the value is fixed, the value that the agent is matched with is drawn according to the distribution $v \sim G$ conditional on $v > \phi(\alpha, p)$. This yields (17).

Finally, the utility $u^{IL_p}(\alpha) = (1-\delta)V = \phi(\alpha, p) - \alpha$.

□

Proposition 3 (Comparative Statics of the Threshold Function ϕ)**.** *Let $\phi(\alpha, p) : \mathbb{R} \times (0, 1] \to \mathbb{R}$ be the optimal threshold for an agent with outside-value α under independent lotteries with success probability p, given by* (16)*. The function $\phi(\alpha, p)$ is continuous in both arguments. Furthermore, let*

[29] To see this, note that $kH(\phi(\alpha, p)) < c$ implies $\phi(\alpha, p) = \alpha$, so this implies $kH(\alpha) < c$, which is equivalent to $\alpha > \alpha_0$. Conversely, since H is weakly decreasing and $\phi(\alpha, p) \geq \alpha$, $kH(\alpha) < c$ also implies $kH(\phi(\alpha, p)) < c$.

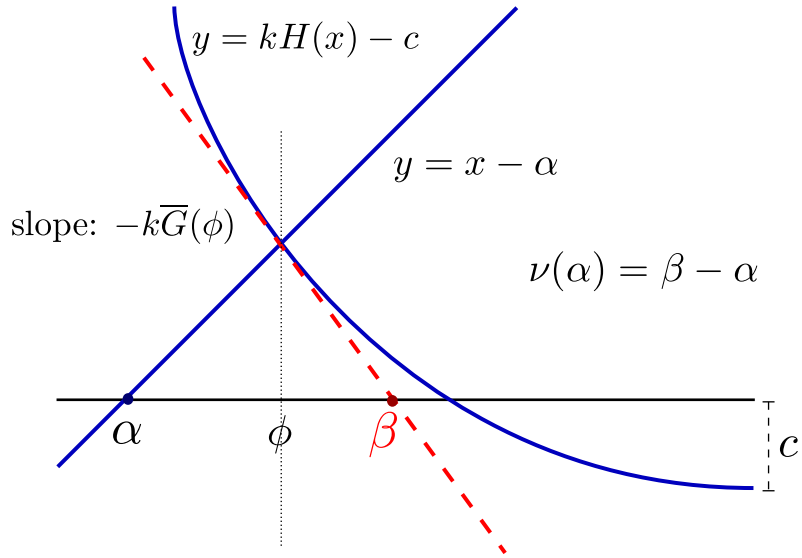

Figure 2: Geometry of the fixed point equation (16) that determines the optimal threshold ϕ under independent lotteries. The curves $y = x - \alpha$ and $y = kH(x) - c$ intersect at $x = \phi$. Let β denote the x-intercept of the tangent to the curve $y = kH(x) - c$ at $x = \phi$. Proposition 2 implies that the distance $\phi - \alpha$ is equal to the utility $u(\alpha)$ of an agent of outside value α; the distance $\beta - \alpha$ is equal to her value per match $\nu(\alpha)$; the ratio $\frac{\phi-\alpha}{\beta-\alpha}$ is equal to her probability of matching $\pi(\alpha)$.

$\overline{v} = \sup\{v : G(v) < 1\}$, $k = \frac{p\delta}{1-\delta}$, *and* $\alpha_0 = H^{-1}(\frac{c}{k})$. *We have*

a) If $\alpha < \alpha_0$, then $\phi(\alpha, p)$ is strictly increasing in both p and α. Furthermore, $\phi(\alpha, p) - \alpha$ is strictly positive and strictly decreasing in α, and converges to zero as $\alpha \to \alpha_0$.

b) When $\alpha \geq \alpha_0$, then $\phi(\alpha, p) = \alpha$. (In particular, the above monotonicity conditions are still weakly true.)

c) For any $p > 0$, $\lim_{\alpha \to -\infty} \phi(\alpha, p) = -\infty$.

d) $\lim_{p \to 0} \phi(\alpha, p) = \alpha$.

Proof of Proposition 3. As illustrated in Figure 2, $\phi(\alpha, p)$ is defined to be the point at which the curves $LS(x) = x - \alpha$ and $RS(x) = \max(0, kH(x) - c)$ intersect (the names LS and RS are given because these represent the left and right sides of (16), respectively). Note that LS is strictly increasing, RS is weakly decreasing, and both curves are continuous in α and p, implying that $\phi(\alpha, p)$ is also continuous in both arguments.

When $\alpha < \alpha_0$, then $\overline{G}(\alpha) > 0$. Increasing p strictly increases RS, and does not affect LS. Increasing α strictly decreases LS, and does not affect the RS. In either case, the intersection strictly shifts to the right. This proves the strict monotonicity of $\phi(\alpha, p)$. Furthermore, since $\phi(\alpha, p) - \alpha$ is the vertical component of the intersection, this strictly decreases when α increases. Moreover, the fact that $\phi(\alpha, p) - \alpha$ is strictly positive follows from the fact that it is equal to $RS(\phi(\alpha, p))$, which is strictly positive for $\alpha < \alpha_0$. The convergence of $\phi(\alpha, p) - \alpha$ to zero as α approaches α_0 from the left

follows from the fact that $RS(\phi(\alpha, p)) \leq \max(0, kH(\alpha) - c)$, which goes to zero as $\alpha \to \alpha_0$.

Note also that $RS(x)$ is identically zero for $\alpha \geq \alpha_0$, so for such α, $\phi(\alpha, p) = \alpha$.

For the third claim, note that when $\alpha \to -\infty$, LS shifts to the left. Since RS is weakly decreasing with slope bounded between $[-1, 0]$, it must be that $\phi(\alpha, p)$ (which is the horizontal component of the point of intersection of the two curves) also tends to $-\infty$.

For the final claim, note that since $\phi(\alpha, p) \geq \alpha$ and H is weakly decreasing, we have $0 \leq \phi(\alpha) - \alpha \leq \max(0, kH(\alpha) - c)$, which tends to zero as p does since $k = \frac{p\delta}{1-\delta}$ and $c \geq 0$. □

D Envelope Theorem

Proposition 1 follows immediately from Lemma 2. The basic mathematical argument parallels that of the revenue equivalence theorem in auction theory, which shows that incentive compatibility conditions pin down a precise relationship between the allocation function and each agent's utility. We show that this mathematical recipe can be applied in our model, which features dynamics and no monetary transfers.

Lemma 2. *For any partial equilibrium outcome E, we have:*

a) The allocation function $\pi^E(\alpha)$ is weakly decreasing, which implies that it is continuous almost everywhere.

b) The utility function $u^E(\alpha)$ is continuous everywhere, and differentiable wherever $\pi^E(\alpha)$ is continuous. For such α,

$$\frac{du^E(\alpha)}{d\alpha} = -\pi^E(\alpha). \tag{23}$$

c) The utility of the highest type is zero: $u^E(\overline{v}) = 0$. (For simplicity, define $u^E(\infty)$ to be a short hand for $\lim_{x\to\infty} u^E(x)$.)

Proof of Lemma 2. Fix a partial equilibrium outcome E. The main idea is that agents should not envy the allocation outcome of other agent types, otherwise their strategy profile would not be optimal. By inspecting the equation for utility $u^E(\alpha)$ in (4), we see that when an agent i of outside option α copies the behavior of an agent i' with outside option α', she obtains the utility

$$u^E(\alpha') + \int_{\underline{v}}^{\overline{v}} (\alpha' - \alpha)\, dP^E(\alpha', v) = u^E(\alpha') + \pi^E(\alpha')(\alpha' - \alpha).$$

Therefore, we have the following incentive-compatibility condition:

$$u^E(\alpha) \geq u^E(\alpha') + \pi^E(\alpha')(\alpha' - \alpha). \tag{24}$$

For the remainder of the proof, the superscript E will be omitted to keep the notation uncluttered. Writing the same inequality above with the roles of α and α' exchanged and adding, we get

$$(\pi(\alpha) - \pi(\alpha'))(\alpha - \alpha') \leq 0 \qquad \forall \alpha, \alpha' \in \mathbb{R} \tag{25}$$

This implies that $\pi(\alpha)$ is weakly decreasing (point 1 of Lemma 2).

Rearranging (24) and the analogous equation with the roles of α and α' exchanged, we get that for all $\alpha > \alpha'$

$$-\pi(\alpha) \geq \frac{u(\alpha) - u(\alpha')}{\alpha - \alpha'} \geq -\pi(\alpha') \tag{26}$$

This implies that $u(\alpha)$ is continuous everywhere. Moreover, whenever $\pi(\alpha)$ is continuous, $u(\alpha)$ is differentiable with derivative equal to $-\pi(\alpha)$.

We now show that $u(\overline{v}) = 0$. Suppose first that $\overline{v} < \infty$. We have

$$0 \leq u(\overline{v}) \leq \phi(\overline{v}, 1) - \overline{v},$$

where the first inequality follows because agents have the option of exiting immediately, and the second follows from Proposition 4 below. Because $\overline{v} = H^{-1}(0) \geq H^{-1}(\frac{1-\delta}{\delta}c)$, Proposition 3 b) states that $\phi(\overline{v}, 1) = \overline{v}$, so $u(\overline{v}) = 0$.

Suppose that $\overline{v} = \infty$, then the same argument as above holds, except that we replace $u(\overline{v})$ by $\lim_{\alpha\to\infty} u(\alpha)$ and $\phi(\overline{v}, 1) - \overline{v}$ by $\lim_{\alpha\to\infty} \phi(\alpha, 1) - \alpha$. This last limit is equal to zero by the Proposition 3 a). □

Proposition 4 (Upper Bound on Utility of an Agent)**.** *In any equilibrium outcome E, the utility of an agent is no more than her utility under guaranteed-choice:*

$$u^E(\alpha) \leq u^{GC}(\alpha) = \phi(\alpha, 1) - \alpha, \tag{27}$$

where $\phi(\alpha, 1)$ is the optimal cutoff under guaranteed choice, as defined in (16) *with $k = \frac{\delta}{1-\delta}$.*

Proof of Proposition 4. Let $V(s)$ be the optimal solution to the Bellman Equation (1). By Lemma 1, we can apply Proposition 1.6.1 of Bertsekas (2012), which implies that $\sup_{s\in S} V(s) < \infty$. Let $\overline{V} = \sup_{s\in S} V(s)$.

$$\overline{V} \leq \max\left(0, \delta\mathbb{E}_{v\sim G}\left[\max\left(\frac{v-\alpha}{1-\delta}, \overline{V}\right)\right] - c\right) = \max\left(0, \delta\overline{V} + \frac{\delta}{1-\delta}H(\alpha + (1-\delta)\overline{V}) - c\right). \tag{28}$$

Let $y = \frac{\delta}{1-\delta}H(\alpha + (1-\delta)\overline{V})$. If $y \leq c$, then the above is equivalent to $\overline{V} \leq 0$. If $y > c$, then since $\overline{V} \geq 0$ (as agents can always exit upon entry), we have $\delta\overline{V} + y - c > 0$, so the above is equivalent to

$(1-\delta)\overline{V} \leq \frac{\delta}{1-\delta}H(\alpha+\overline{V}) - c$. Either way, the above is equivalent to

$$(1-\delta)\overline{V} \leq \max\left(0, \frac{\delta}{1-\delta}H(\alpha+(1-\delta)\overline{V}) - c\right). \tag{29}$$

Let the function ϕ be as in (16). Setting $(1-\delta)\overline{V} = \phi(\alpha,1) - \alpha$ would make the above left side (LS) and right side (RS) equal. Since the LS is an increasing function of $\overline{V}$ for $\delta < 1$ and since the RS is weakly decreasing in $\overline{V}$, it must be that any $\overline{V}$ satisfying the above inequality also satisfies $(1-\delta)\overline{V} \leq \phi(\alpha,1) - \alpha$. By (19), $u^E(\alpha) \leq (1-\delta)\overline{V} \leq \phi(\alpha,1) - \alpha$. □

E Pessimality of Independent Lotteries

In this section, we show that when participation cost is negligible, values are high so that everyone prefers every development to their outside option, and supply is scarce, then independent lotteries are simply matching random people to random developments, which is a bad outcome in terms of both matching and targeting.

Definition 10. *Define the* **random matching** *outcome function as*

$$RM(\alpha, v) = \mu G(v). \tag{30}$$

This matches random agents to random developments, irrespective of values or outside options. It is straightforward to show that this outcome is targeting dominated by all equilibrium outcomes with arbitrary $c \geq 0$, and match dominated by any equilibrium outcome with $c = 0$.

Proposition 5. *If $c = 0$ and $\overline{\alpha} < \underline{v}$, then there exists a threshold $\mu_0 > 0$ on the supply-demand ratio such that whenever $\mu < \mu_0$, every equilibrium outcome under independent lotteries is equal to random matching.*

The proof of Proposition 5 is in Appendix F.9. The intuition is that without participation costs, agents under independent lotteries apply to any development for which the added value $v - \alpha$ is greater than their continuation value, which is near zero if supply is scarce. Since everyone prefers every development to their outside option, they will apply everywhere, and the resultant matching is completely random.

Online Appendix

F Additional Proofs

F.1 Proofs of Outcome Equivalence (Theorem 1)

Propositions 6 and 7 not only imply Theorem 1, but also characterize the structure of equilibrium outcomes. Proposition 1 implies that for any equilibrium outcome, the waiting function can be expressed in terms of the outcome function, so to prove outcome equivalence, it suffices to show that the set of outcome functions are equal. As explained in Section 3.4, we use the simplified notation of $E(\alpha, v)$ to denote the outcome function of equilibrium outcome E.

Define $u^{GC}(\alpha) = \phi(\alpha, 1) - \alpha$ to be the utility for an agent of outside option α under guaranteed choice, and let $\alpha^{GC}(x)$ be its inverse. By Proposition 3, this is well defined in the domain $x \in (0, \infty)$ and is strictly decreasing. $\alpha^{GC}(x)$ is the agent type that obtains utility of exactly x under guaranteed choice. Define $\alpha^{GC}(0) = H^{-1}\left(\frac{1-\delta}{\delta}c\right)$.

Proposition 6 (Equivalence to Waitlist with Deferral)**.** *For any participation cost $c \geq 0$, the unique equilibrium outcome of the waitlist with deferral and the ticket-saving lottery is given by*

$$WD(\alpha, v) = \begin{cases} qGC(\alpha, v) & \text{if } \alpha \leq \alpha_1 := \alpha^{GC}(\frac{1-q}{q}c), \\ 0 & \text{otherwise,} \end{cases} \tag{31}$$

where $GC = IL_1$ is the outcome under guaranteed choice as given in (17)*, and $q \in (0, 1]$ is the unique value that satisfies the fixed point equation*

$$\mu = q \int_{-\infty}^{\alpha^{GC}\left(\frac{1-q}{q}c\right)} \pi^{GC}(x)\, dF(x). \tag{32}$$

The corresponding utilities are

$$u^{WD}(\alpha) = \max\{0, q(\phi(\alpha, 1) - \alpha) - (1 - q)c\}, \tag{33}$$

where $\phi(\alpha, p)$ is given by (16)*.*

When $c = 0$, WD is also the unique equilibrium outcome of the single-entry lottery and of the common lottery.

The above expressions for the outcome function and utility have the following intuitive explanation. Both the waitlist with deferral and the ticket-saving lottery are equivalent to having agents wait $\tau \sim \Lambda$ periods before playing the guaranteed choice game. The equilibrium parameter q represents $\mathbb{E}_{\tau\sim\Lambda}[\delta^\tau]$,

which is the probability that an agent has not yet exited exogenously by period τ, and the expected cost one has to pay before reaching period τ, multiplied by the scaling term of $(1-\delta)$ from (4), is

$$(1-\delta)\mathbb{E}_{\tau\sim\Lambda}\left[\sum_{k=1}^{\tau}\delta^{k-1}\right]c = (1-q)c. \tag{34}$$

so only agents whose utility under guaranteed choice is more than $\frac{1-q}{q}c$ will choose to join, and those who do obtain the guaranteed choice outcome function scaled down by $q \leq 1$.

Proposition 7 (Equivalence to Independent Lotteries)**.** *For any participation cost $c \geq 0$, the set of equilibrium outcomes under the waitlist without deferral and independent lotteries are both non-empty and equal to*

$$\{IL_p : p \in (0,1], \overline{\pi}^{IL_p} = \mu\}, \tag{35}$$

where IL_p is the outcome under independent lotteries with success probability p as given in (17)*.*

Proof of Proposition 6. The idea is to show that the waitlist with deferral, the ticket-saving lottery, and the single-entry lottery are special cases of a general class of mechanism which we call *delayed guaranteed choice* (Definition 11). This equivalence is described in Definition 12 and Lemma 3. The proof concludes by showing that under certain assumptions, any delayed guaranteed choice mechanism yields outcome functions of a particular form (Lemma 4).

Definition 11. *Given an arbitrary distribution Λ on the non-negative integers plus infinity, define the* **delayed guaranteed choice** *matching rule with idle-time distribution Λ as follows. Each agent plays guaranteed choice after being idle for $\tau_i \sim \Lambda$ periods. The agent knows the distribution Λ but not her realization τ_i.*

To precisely encode this as a matching rule, let the set of actions be $\{Accept, Reject\}$. Let the state space be $\{(s_1, s_2) : s_1, s_2 \in \mathbb{N}, s_1 \geq s_2\}$. The first component represents the number of periods the agent has waited since entry; the second encodes the last period the agent chose to accept, counting from time of entry. The initial state is $(0,0)$. The transitions are as follows.

- *Accept:*
  - *with probability $p_\Lambda(s_1, s_2) := \mathbb{P}_{\tau\sim\Lambda}(\tau \leq s_1 | \tau \geq s_2)$, the idle time has been reached and the agent is matched to the current development;*
  - *otherwise, the state transitions to (s_1+1, s_1+1). (The idle time has not been reached and the agent knows it.)*
- *Reject: the state transitions to (s_1+1, s_2). (The agent waits another period but does not learn anything new about her idle time.)*

The following definition gives a strong notion of equivalence of matching rules. In colloquial terms, two rules are strategically equivalent if one can transform both to a common rule by relabeling states and actions, and by collapsing multiple states into equivalence classes. An implication is that strategically equivalent matching rules have the same set of equilibrium outcomes.

Definition 12. *A matching rule* (S', D', A', T') *is said to be a* simpler representation *of matching rule* (S, D, A, T) *if there exists a surjective mapping* $\omega : S \cup A \to S' \cup A'$ *taking state to state and action to action,*[30] *such that for each state* $s \in S$*, the initial state distribution is preserved:*

$$\sum_{s \in \omega^{-1}(s')} D(s) = D'(s') \qquad \text{for all } s' \in S'. \tag{36}$$

The action set is preserved:

$$A'_{\omega(s)} = \{\omega(a) : a \in A_s\}, \tag{37}$$

and for each action $a \in A_s$*, and each state* $\tilde{s} \in S'$*, we have:*

$$T'_{\omega(a)}(\omega(s), \tilde{s}) = \sum_{s' : \omega(s') = \tilde{s}} T_a(s, s') \tag{38}$$

$$T'_{\omega(a)}(\omega(s), m) = T_a(s, m) \tag{39}$$

Two matching rules are said to be **strategically equivalent** *if there is a common matching rule that is a simpler representation of both.*

Define the ticket-saving lottery **with all-or-nothing constraint** as the modification of the ticket-saving lottery in which the set of action each period is $j \in \{0, s\}$. In other words, the agent is constrained each period to either use all tickets available, or use none at all.

Lemma 3. *The following matching rules are strategically equivalent to a delayed guaranteed choice matching rule with a certain idle-time distribution* Λ*.*

a) Waitlist with deferral with expected idle-time τ*:* Λ *has support* $\{\lfloor \tau \rfloor, \lceil \tau \rceil\}$ *and mean* τ*.*

b) Ticket-saving lottery with success probability p *and all-or-nothing constraint:* $\Lambda = Geom(p) - 1$*, where* $Geom(p)$ *is the geometric distribution with parameter* p*.*

c) Single-entry lottery with success probability p*:* Λ *has support* $\{0, \infty\}$*, with probability of being zero equal to* p*.*

Let the probability that starting in state (s_1, s_2), the agent will reach her idle time before exoge-

[30] The mapping encodes an equivalence relationship; each element of $S' \cup A'$ is a representative of a equivalence class of $S \cup A$.

nously exiting be denoted

$$q_\Lambda(s_1, s_2) = \mathbb{E}_{\tau\sim\Lambda}[\delta^{\max(\tau - s_1, 0)} | \tau \geq s_2], \tag{40}$$

the following assumption is satisfied by the delayed guaranteed choice description of the waitlist with deferral and the ticket-saving lottery.

Assumption 2. *The idle-time distribution Λ is such that the function $q_\Lambda(s_1, s_2)$ is minimized at $(s_1, s_2) = (0, 0)$, with $q_\Lambda(0, 0) = \mathbb{E}_{\tau\sim\Lambda}[\delta^\tau] > 0$.*

The above assumption implies that agents who choose initially to participate will never choose to exit, which also holds when $c = 0$.

Lemma 4. *When $c = 0$ or when Λ satisfies Assumption 2, an optimal strategy for an agent with outside option α under the delayed guaranteed choice matching rule with idle time distribution Λ is as follows: if $\alpha > \alpha_1 := \alpha^{GC}(\frac{1-q_\Lambda(0,0)}{q_\Lambda(0,0)}c)$, then exit immediately; otherwise, do not voluntarily exit, and accept every development with value exceeding $\phi(\alpha, 1)$, which is the same threshold as used in guaranteed choice.*

Moreover, in all optimal strategy profiles, the outcome function and utilities are equal to $WD(\alpha, v)$ and $u^{WD}(\alpha)$ as given in (31) and (33) with parameter $q = q_\Lambda(0, 0)$. [31]

Note that when $c = 0$, the description in Lemma 4 is equivalent to selecting a random proportion of agents of measure q to play guaranteed choice, which is exactly the common lottery.

Moreover, in any equilibrium, the total match rate must equal to μ, so the value of q must satisfy the fixed point equation (32), which always has an unique solution since the left side is a constant, $\alpha^{GC}\left(\frac{1-q}{q}c\right)$ is strictly increasing in q, $\pi^{GC}(x) > 0$ for any $x < \alpha^{GC}(0)$, and $\overline{\pi}^{GC} \geq \mu$.

To finish off the proof of Proposition 6, we show that assuming the all-or-nothing constraint for the ticket-saving lottery is without loss of generality.

Lemma 5. *The ticket-saving lottery with all-or-nothing constraint is outcome equivalent to the ticket-saving lottery without this constraint.*

The proofs to Lemmas 3, 4, and 5 are in Section F.2. □

Proof of Proposition 7. The proof is analogous to that of Proposition 6. We show that both classes of matching rules are strategically equivalent to special cases of a more general class of matching rules, which we call *periodic-offer*: offer each agent a new development every certain number of periods, with the inter-offer time possibly random but following a fixed distribution (Definition 13 and Lemma 6). We then derive the set of equilibrium outcomes for such matching rules via a common analysis (Lemma 7).

[31] The uniqueness of the outcome function assumes that agents who are indifferent between participating and leaving will participate; otherwise the outcome is not unique for the agents who are indifferent.

Definition 13. *Given an arbitrary distribution Λ on the non-negative integers plus infinity, define the* **periodic-offer** *matching rule with idle-time distribution Λ as follows. The actions are "Accept" and "Reject," and the state space is $\mathbb{N}$. The initial state is 0, and in this state, the an idle time $\tau_i \sim \Lambda$ for the agent is drawn. The agent knows the distribution Λ but not the realization τ_i. The transitions are as follows:*

- *With probability $p_\Lambda(s) := \mathbb{P}_{\tau\sim\Lambda}(\tau = s|\tau \geq s)$,*
  - *"Accept" results in being matched to the current development,*
  - *"Reject" results in the agent going back to state 0 and receiving a new draw of τ_i.*
- *With probability $1 - p_\Lambda(s)$, either action results in the state incrementing to $s + 1$.*

Lemma 6. *The following matching rules are strategic equivalent (see Definition 12) to the periodic offer matching rule with a certain idle-time distribution Λ.*

a) Waitlist without deferral with expected idle-time τ: Λ has support $\{\lfloor\tau\rfloor, \lceil\tau\rceil\}$ and mean τ.

b) Independent lotteries with success probability p: $\Lambda = Geom(p)-1$, where $Geom(p)$ is the geometric distribution with parameter p.

Define $q_\Lambda(s_1, s_2)$ as in (40), then $\delta q_\Lambda(s, s)$ is the probability that starting from state s, the agent will be able to receive an offer before exiting exogenously. Note that both of the idle-time distributions above satisfy Assumption 2: the probability of receiving another offer is minimized in the initial state $(0, 0)$.

Lemma 7. *Under Assumption 2, an optimal strategy under periodic-offer with idle-time distribution Λ is to follow an optimal strategy to independent lotteries with success probability $p_0 := \frac{(1-\delta)q_\Lambda(0,0)}{1-\delta q_\Lambda(0,0)}$. Moreover, every optimal strategy yields the same outcome function IL_{p_0}.*

Finally, the set of equilibrium outcomes is non-empty because by continuity, $\{\overline{\pi}^{IL_p} : p \in (0, 1]\} \supseteq (0, \overline{\pi}^{GC}] \ni \mu$. This completes the proof of Proposition 7. The proofs of Lemmas 6 and 7 are given in Appendix F.2. □

F.2 Proofs of Lemmas used in the Equivalence Proof (Appendix F.1)

Proof of Lemma 3. For the waitlist with deferral, consider first the strategically equivalent matching rule in which the set of actions is $\{Accept, Reject\}$ in every state, but the action is ignored and the state is incremented if the idle time has not been reached.[32] Consider the following mapping from the

[32]Precisely speaking, the action is ignored with probability 1 for every state $s < \lfloor\tau\rfloor$, with probability $\tau - \lfloor\tau\rfloor$ in state $s = \lfloor\tau\rfloor$, and with probability 0 in state $s \geq \lfloor\tau\rfloor + 1$. This matching rule is a simpler representation of waitlist with deferral as presented in Section 3.2 because the action "Wait" can map to "Accept," which is ignored anyway.

delayed guaranteed choice matching rule to this: the mapping preserves each action, and maps state $(s_1, s_2) \to s_1$. When Λ has support $\{\lfloor\tau\rfloor, \lceil\tau\rceil\}$ and mean τ, this mapping preserves transitions because,

$$p_\Lambda(s_1, s_2) = \mathbb{P}_{\tau\sim\Lambda}(\tau \leq s_1 | \tau \geq s_2) = \begin{cases} 0 & \text{if } s_1 < \lfloor\tau\rfloor, \\ 1 + \lfloor\tau\rfloor - \tau & \text{if } s_1 = \lfloor\tau\rfloor, \\ 1 & \text{if } s_1 \geq \lfloor\tau\rfloor + 1., \end{cases} \tag{41}$$

which is equal to the probability of matching in a waitlist with deferral in state s_1 when the agent accepts.

For the ticket-saving lottery with all-or-nothing constraint, consider the mapping from delayed guaranteed choice which takes state (s_1, s_2) to state $s_1 - s_2$ in the ticket saving lottery, action "Accept" in delayed guaranteed choice to the action in the ticket saving lottery of using all tickets, and action "Reject" in delayed guaranteed choice to using no tickets. When $\Lambda = Geom(p) - 1$, it preserves the transitions because

$$p_\Lambda(s_1, s_2) = \mathbb{P}_{\tau\sim\Lambda}(\tau \leq s_1 | \tau \geq s_2) = 1 - (1-p)^{s_1 - s_2}, \tag{42}$$

which is equal to the probability of matching in the ticket-saving lottery when the agent spends $s_1 - s_2$ tickets.

For the single-entry lottery, consider the mapping which takes "Accept" in delayed guaranteed choice to "Enter" in the single-entry lottery, and "Reject" to "Abstain." It takes any state $(s_1, 0)$ in delayed guaranteed choice to the initial state in the single-entry lottery (corresponding to not having entered a lottery yet), and any state (s_1, s_2) with $s_2 \geq 1$ to the null state in the single-entry lottery (corresponding to already having entered a lottery). The transitions are preserved because

$$p_\Lambda(s_1, s_2) = \mathbb{P}_{\tau\sim\Lambda}(\tau \leq s_1 | \tau \geq s_2) = \begin{cases} p & \text{if } s_2 = 0, \\ 0 & \text{otherwise}, \end{cases} \tag{43}$$

which is equal to the probability of matching in the single-entry lottery when the agent chooses "*Enter*" in states 1 and 0 respectively. □

Proof of Lemma 4. First note the following fact about the threshold α_1 as defined in the lemma: if $c > 0$, then Assumption 2 implies that $\alpha_1 < \infty$. Moreover, $q_\Lambda(0,0)u^{GC}(\alpha) - (1 - q_\Lambda(0,0))c \geq 0$ if and only if $\alpha \leq \alpha_1$, because the definition of α_1 implies this expression is identically zero at $\alpha = \alpha_1$, and Proposition 3 implies that $u^{GC}(\alpha)$ is weakly decreasing in α everywhere and strictly decreasing at $\alpha = \alpha_1$.

Consider for now only agents with outside option $\alpha \leq \alpha_1$. Consider also the restriction on the strategy space in which agents cannot voluntarily exit. Under this constraint, it is optimal to play

delayed guaranteed choice the same way as guaranteed choice, since whatever actions have no effect before the idle time has been reached. Therefore, an optimal strategy is to accept development j if and only if $v_{ij} > \phi(\alpha, 1)$, which is the optimal threshold under guaranteed choice (see Proposition 2). This strategy results in continuation values,

$$V(s_1, s_2) = q_\Lambda(s_1, s_2)\frac{u^{GC}(\alpha)}{1-\delta} - (1 - q_\Lambda(s_1, s_2))\frac{c}{1-\delta}, \tag{44}$$

because $q_\Lambda(s_1, s_2)$ is the probability that the agent will not exogenously exit before her idle time is reached, and

$$(1 - q_\Lambda(s_1, s_2))\frac{c}{1-\delta} = \mathbb{E}_{\tau\sim\Lambda}[\sum_{k=1}^{\tau-s_1} \delta^{k-1}c|\tau \geq s_2],$$

is the expected participation cost incurred before her idle time is reached. This implies that the utility of the agent from state 0, scaled as in (4), is

$$u(\alpha) = (1-\delta)V(0,0) = q_\Lambda(0,0)(\phi(\alpha,1) - \alpha) + (1 - q_\Lambda(0,0))c.$$

Now, we show that the above restriction on the strategy space is without loss of generality. It is clearly true if participation cost $c = 0$. If $c > 0$, then $\alpha < \alpha_1$ implies that the expression in (44) satisfies

$$(1-\delta)V(s_1, s_2) + c = q_\Lambda(s_1, s_2)(u^{GC}(\alpha) + c) > q_\Lambda(0,0)(u^{GC}(\alpha_1) + c) = c.$$

So $V(s_1, s_2) > 0$ for all states (s_1, s_2), so the agent should never choose to exit when following the above strategy. This proves that the above strategy is optimal even allowing voluntary exits, and that (44) describes the optimal continuation values. This also proves that (33) with q set to $q_\Lambda(0,0)$ describes the optimal utility for these agents.

For agents with outside $\alpha > \alpha_1$, we show that an optimal strategy for them is to exit immediately. This is because the arguments above showed that when $\alpha_1 < \infty$, then an agent with outside value exactly α_1 receives zero utility in state $(0,0)$. Since utilities are weakly decreasing in α, all agents with outside option $\alpha > \alpha_1$ must also receive zero utility in state $(0,0)$, so it is optimal to exit immediately.

It remains to prove that set of outcome functions under optimal strategies are as given in Lemma 4. For $\alpha > \alpha_1$, Proposition 1 implies that in any optimal strategy, the agent matches with probability zero.[33] So the outcome is zero.

For $\alpha < \alpha_1$, when $c = 0$ or Λ satisfies Assumption 2, it is never optimal to voluntarily exit, since 44 implies that $V(s_1, s_2) \geq V(0,0) > 0$. Moreover, in order to obtain utilities as in (33), the agent must

[33] To apply Proposition 1, let E be a partial equilibrium outcome in which agents follow a certain optimal strategy profile to delayed guaranteed choice. If $\alpha_1 < \infty$, then we have by the formula for utilities that we have already proved that $u^E(\alpha_1) = 0$, which implies that $\pi^E(\alpha) = 0$ for any $\alpha > \alpha_1$.

play an optimal strategy to guaranteed choice when her idle time has been reached, and Proposition 2 implies that conditional on reaching her idle time, the outcome must be as in (17) with $p = 1$. Since without exiting, she cannot affect the probability that she remains in the system until her idle time is reached, the outcome must be as in (31) with $q = q_\Lambda(0, 0)$.

For $\alpha = \alpha_1$, the agent is indifferent between exiting and staying in the initial state, and in each state (s_1, s_2) with $V(s_1, s_2) = 0$. Since we assume that she always stays, the outcome function is as in the case with $\alpha < \alpha_1$. □

Proof of Lemma 5. Let $p_k = 1 - (1 - p)^k$. For a fixed outside option α, the Bellman equation for the ticket-saving lottery is

$$V(s) = \max(0, \delta \mathbb{E}_{v \sim G}[\max_{0 \le k \le s} \{p_k \frac{v - \alpha}{1 - \delta} + (1 - p_k) V(s - k + 1)\}] - c). \tag{45}$$

Let $V_0(s)$ be the DP value for the ticket-saving lottery with all or nothing constraint. Then V_0 satisfies the modification of (45) with the domain of the inner maximum replaced with $k \in \{0, s\}$. By Lemma 3 and 4, the optimal threshold in the inner maximum for V_0 is $\phi(\alpha, 1)$, which is the optimal threshold under guaranteed choice. Let $u^{GC}(\alpha) = \phi(\alpha, 1) - \alpha$. This implies that

$$V_0(s + 1) = p_s \frac{u^{GC}(\alpha)}{1 - \delta} + (1 - p_s) V_0(1). \tag{46}$$

Now, let $\lambda = \frac{u^{GC}(\alpha)}{1-\delta}$, using the above and the identity $(1 - p_a)(1 - p_b) = 1 - p_{a+b}$, we have that for any $1 \le k \le s$,

$$\begin{aligned} p_k \lambda + (1 - p_k) V_0(s - k + 1) &= p_k \lambda + (1 - p_k)[p_{s-k} \lambda + (1 - p_{s-k}) V_0(1)] \\ &= [p_k + (1 - p_k) p_{s-k}] \lambda + (1 - p_s) V_0(1) \\ &= p_s \lambda + (1 - p_s) V_0(1). \end{aligned}$$

This implies that $V_0(s)$ also satisfies the Bellman Equation (45). This is because after we plug in V_0, the inner maximum of (45) is the maximum of $k+1$ linear function of v, all of of which have a common intersection point at $v = \phi(\alpha, 1)$. Since p_s is increasing in s, the maximum is achieved with action $k = 0$ if $v \le \phi(\alpha, 1)$ and with $k = s$ if $v \ge \phi(\alpha, 1)$, so we can remove the actions $k \in \{1, 2, \cdots, s - 1\}$ without loss of generality and get back the Bellman equation for the ticket-saving lottery with all-or-nothing constraint.

Moreover, since G is continuous, the action $k \in \{1, 2, \cdots, s - 1\}$ is optimal with probability zero in state s, so imposing the all-or-nothing constraint does not affect the set of equilibrium outcomes. □

Proof of Lemma 6. For the waitlist without deferral, consider the identity mapping that takes state

s in periodic-offer to s in the waitlist, while preserving the actions *Accept* and *Reject*. This mapping preserves the transitions by the definitions of the two matching rules.

For independent lotteries, consider the mapping which takes every state in periodic-offer to the only state under independent lotteries, the action *Accept* in periodic-offer to *Enter* under independent lotteries, and the action *Reject* to *Abstain*. This preserves the transitions because with $\Lambda = Geom(p) - 1$, $p_\Lambda(s) = \mathbb{P}_{\tau\sim\Lambda}(\tau = s|\tau \geq s) = p$ always. □

Proof of Lemma 7. Note that $\delta' := \delta q_\Lambda(0,0)$ is the probability that starting from entry, or from receiving a previous offer, the agent will get to decide on another offer before exiting. By inspecting (16) and (17) in Proposition 2, we see that the outcome under independent lotteries with success probability p_0 is the same as the outcome under guaranteed choice when the continuation probability is changed to δ'.[34] By Proposition 2, an optimal strategy is to exit immediately if $\alpha > \alpha_0 := H^{-1}(c/k)$, where $k = \frac{p_0\delta}{1-\delta}$, and to not exit otherwise and accept every lottery with value exceeding $\phi(\alpha, p_0)$.

We show that under Assumption 2, the above strategy is optimal under periodic-choice with idle-time distribution Λ. Consider first a restriction on the strategy space that disallows voluntary exit in states $s > 0$. In this case, the agent can effectively do nothing before her idle-time is reached. In state 0, the agent knows that her chance of still being around for the next offer is $\delta' = \delta q_\Lambda(0,0)$ and she can exit now and get a reward of r. In every state $s > 0$, the agent can behave as if the current state is one with an offer, because in other states the action does not matter anyway. Therefore, the agent's decision problem is the same as in under guaranteed choice with continuation probability δ', so the above strategy is optimal. Furthermore, by the argument above that allows us to ignore states $s > 0$ before the idle time has been reached, and by Proposition 2, the outcome under any optimal strategy (with the above restriction) is the same and is equal to IL_{p_0}.

Now, it suffices to show that allowing voluntary exits does not affect the optimality of the above strategy and does not change the set of equilibrium outcomes. For ease of computation, adopt the change of variable $\tilde{V}(s) = V(s) + r$, where $r = \frac{c}{1-\delta}$. As explained in Appendix A.2, this can be interpreted as offering a one-time reward of r for voluntary exits and not charging any participation cost. For agent types $\alpha \leq \alpha_0$, the transformed continuation value upon entry is

$$\tilde{V}(0) = \frac{\phi(\alpha, p_0) - \alpha}{1-\delta} + r,$$

and the transformed continuation value in state s under the above strategy is:

$$\tilde{V}(s) = \frac{q_\Lambda(s,s)}{q_\Lambda(0,0)}\tilde{V}(0) \tag{47}$$

[34] The value of p_0 should be such that $p_0\frac{\delta}{1-\delta} = \frac{\delta'}{1-\delta'}$, which implies that $p_0 = \frac{(1-\delta)q_\Lambda(0,0)}{1-\delta q_\Lambda(0,0)}$.

This is because her transformed continuation value at the time of receiving an offer is $\tilde{V}_0/(\delta q_\Lambda(0,0))$, and $\delta q_\Lambda(s,s)$ is the chance that the agent will not be forced to exit before receiving the next offer. When $r = 0$ or when Λ satisfies Assumption 2, it is impossible for $\tilde{V}(0) \geq r$ but $\tilde{V}(s) < r$ for $s > 0$, so agents never have strict incentives to exit in states $s > 0$, and by our assumption that only agents with strict incentives will exit, they will not do so. This proves the desired result for $\alpha \leq \alpha_0$.

For agents with outside option $\alpha > \alpha_0$, we get by an analogous argument as found in the proof of Proposition 4 that in any optimal strategy of periodic-offer, these agents must match with probability zero[35], which yields the desired result. □

F.3 Efficacy of the Common Lottery for Matching (Proof of Theorem 2)

We first show two structural results of independent lotteries and single entry lottery (Propositions 8 and 9).

Proposition 8. *The value per match under independent lotteries matching rule is weakly increasing in the success probability p and weakly decreasing in the participation cost c.*

Geometric Proof of Proposition 8. We first give a geometric proof based on analyzing Figure 2 in Section C, which shows that the value per match $\nu(\alpha) = \beta - \alpha$, where β is the x-intercept of the tangent of the function $y = kH(x) - c$ at $x = \phi$. From the figure, we see that increasing c shifts the function $y = kH(x) - c$ downward, which moves the x-intercept β to the left. The decrease in β is strict until $c \geq kH(\alpha)$, at which point the value per match becomes identically zero as it is optimal for the agent to exit immediately upon entry.

Increasing p corresponds to increasing $k = \frac{p\delta}{1-\delta}$. Figure 3 illustrates how the x-intercept β weakly increases as k increases. Let $k_1 < k_2$ be the two values of k, with corresponding acceptance thresholds ϕ_1 and ϕ_2 from the fixed point equation (16). Let l_1 be the tangent to the curve $y = k_1H(x) - c$ at $x = \phi_1$ and l_2 be the tangent to the higher curve $y = k_2H(x) - c$ at $x = \phi_2$. Draw also the tangent l' to the higher curve $y = k_2H(x) - c$ at the lower threshold $x = \phi_1$. Let the x-intercept of lines l_1, l', and l_2 be β_1, β' and β_2 respectively. Since the higher curve $y = k_2H(x) - c$ is the result of vertically scaling $y = k_1H(x) - c$ by the factor k_2/k_1 away from the horizontal line $y = -c$, the tangents l_1 and l' must intersect on $y = -c$. Since $c \geq 0$, this implies that the x-intercept $\beta' \geq \beta_1$, with the inequality being strict if $c > 0$. Now, by Proposition 3, $\phi_2 \geq \phi_1$, so by the convexity of the curve $y = k_2H(x) - c$, the x-intercept $\beta_2 \geq \beta'$. Combining, we get that $\beta_2 \geq \beta_1$ with the inequality strict if $c > 0$. This implies that the value per match is weakly higher under the higher success probability. □

Algebraic Proof of Proposition 8. By Proposition 2, under independent lotteries with success

[35] The argument is as follows. We have already shown that under periodic-choice with participation cost c, the utility of an agent with outside value α_0 is zero, so by Proposition 1, the probability of matching must be zero for $\alpha > \alpha_0$.

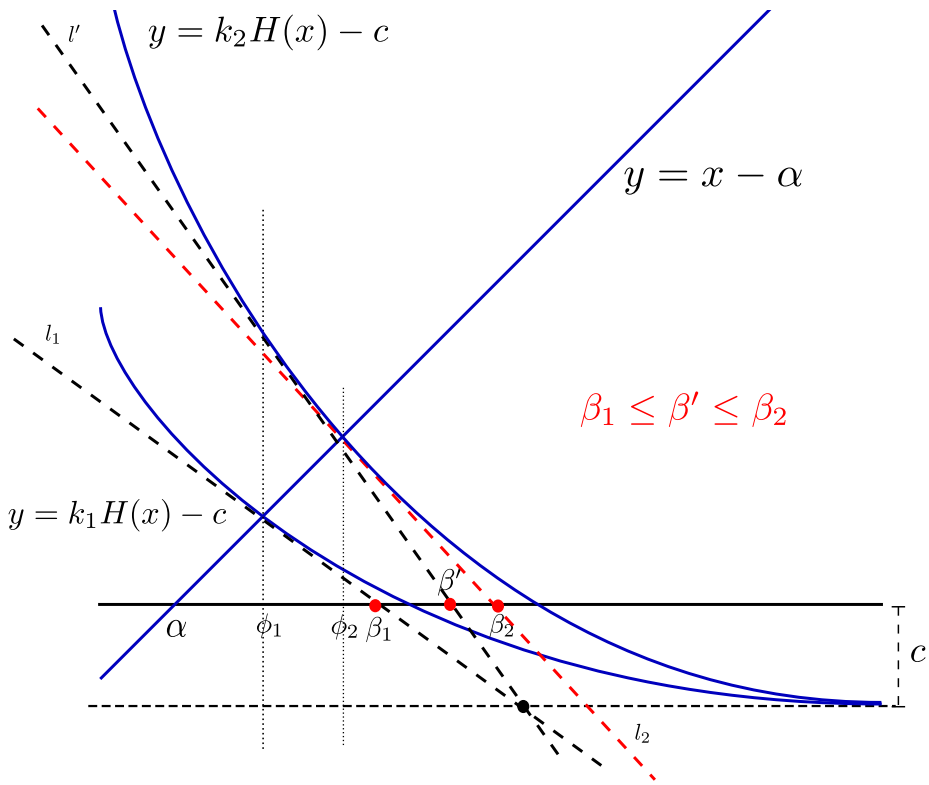

Figure 3: Illustration of how the x-intercept β shifts to the right as k increases. As can be seen, the curve $y = k_2H(x) - c$ is the result of vertical scaling the curve $y = k_1H(x) - c$ away from the horizontal line $y = -c$. As a result, their tangents at $x = \phi_1$ (the lines l' and l_1) intersect on the horizontal line $y = -c$, which implies that the x-intercept $\beta' \geq \beta_1$. Now, l_2 is the tangent to $y = k_2H(x) - c$ at $x = \phi_2 \geq \phi_1$. By convexity, the x-intercept β_2 of the tangent line l_2 shifts to the right from β', which implies that $\beta_2 \geq \beta' \geq \beta_1$.

probability p, the value per match of an agent type α is

$$\nu(\alpha) = \phi - \alpha + \frac{\phi - \alpha}{k\overline{G}(\phi)}, \tag{48}$$

where $\phi = \phi(\alpha, p)$ as given in Proposition 2 and $k = \frac{p\delta}{1-\delta}$.

When the participation cost c increases, we see by inspection of the fixed point equation (16) that $\phi - \alpha$ weakly decreases. Moreover, ϕ weakly decreases, which means that $\overline{G}(\phi)$ weakly increases. Hence, the value per match is weakly decreasing in c.

When the success probability p increases, $k = \frac{p\delta}{1-\delta}$ increases. Suppose it increases from k_1 to k_2. From Proposition 2, the threshold on α for exiting is weakly higher under k_2 than k_1, so it suffices to consider the case in which the agent matches with positive probability under both success probabilities. In this case, $\overline{G}(\phi_1) \geq \overline{G}(\phi_2) > 0$, and the above expression for value per match becomes

$$\nu(\alpha) = \phi - \alpha + \frac{H(\phi) - c/k}{\overline{G}(\phi)}. \tag{49}$$

Now, $D(x) := H(x) - c/k_2$ is convex, non-negative, and strictly decreasing in the interval $(-\infty, \phi_2]$, with subgradient $-\overline{G}(x)$. Hence, we can apply the following lemma.

Lemma 8. *If $D(x)$ is a strictly decreasing, non-negative convex function on the domain $(-\infty, \overline{x}]$, and $(-d(x))$ is any subgradient to $D(x)$ at x, then*

$$h(x) := x + \frac{D(x)}{d(x)}$$

is a weakly increasing function of x on the domain $(-\infty, \overline{x}]$.

By Lemma 8,

$$\phi_1 - \alpha + \frac{H(\phi_1) - c/k_1}{\overline{G}(\phi_1)} \leq \phi_1 - \alpha + \frac{H(\phi_1) - c/k_2}{\overline{G}(\phi_1)} \leq \phi_2 - \alpha + \frac{H(\phi_2) - c/k_2}{\overline{G}(\phi_2)}, \tag{50}$$

so the match value is weakly higher under k_2 than under k_1. □

Proof of Lemma 8. For any $x < x' \leq \overline{x}$, we have by convexity,

$$D(x') \geq D(x) + (x' - x)(-d(x)).$$

Since $d(x) > 0$, and $d(x) \geq d(x')$, the above implies that

$$\begin{aligned} x' - x &\geq \frac{D(x)}{d(x)} - \frac{D(x')}{d(x)} \\ &\geq \frac{D(x)}{d(x)} - \frac{D(x')}{d(x')}, \end{aligned}$$

which implies that $h(x') \geq h(x)$. □

Proposition 9. *For any $c \geq 0$ and $p \leq 1$, an optimal strategy for an agent with outside option α under the single entry lottery is to follow that of the agent under the guaranteed choice matching rule with participation cost c/p. Furthermore, in all optimal strategies, the outcome is the same and given by $SE_p(\alpha, v) = p \cdot GC_{\frac{c}{p}}(\alpha, v)$, and utility $u^{SE_p}(\alpha) = p \cdot u^{GC_{\frac{c}{p}}}(\alpha, v)$, where $GC_{\frac{c}{p}}$ denotes the guaranteed choice outcome under participation cost c/p.*

Proof of Proposition 9. The proof is analogous to that of Proposition 2, except the Bellman equation for the continuation value of agent type α is now

$$V = \max\left(0, \delta \mathbb{E}_{v \in G}\left[\max\left(p\frac{v-\alpha}{1-\delta}, V\right)\right] - c\right) = \max\left(0, \delta V + \frac{p\delta}{1-\delta} H\left(\alpha + \frac{1-\delta}{p} V\right) - c\right).$$

Let $\psi = \alpha + \frac{1-\delta}{p} V$ be the optimal acceptance threshold, we follow the same argument as in Proposition 2 and get that ψ satisfies the fixed point equation

$$\psi - \alpha = \max\left(0, \frac{\delta}{1-\delta} H(\psi) - \frac{c}{p}\right), \tag{51}$$

which is the same as (16) for the guaranteed choice problem with participation cost $\frac{c}{p}$. The utility is $(1-\delta)V = p(\psi - \alpha) = p \cdot GC_{\frac{c}{p}}(\alpha)$. The outcome function is p multiplied by the outcome conditional on winning the lottery, which is the same as $GC_{\frac{c}{p}}(\alpha, v)$. As in the proof of Proposition 2, uniqueness of the outcome follows from our assumption that agents indifferent about staying will stay, and from G being continuous. □

Proof of Theorem 2 a). The common lottery has an unique equilibrium outcome because the probability that an agent is selected to be eligible must be equal to $\mu/(\overline{\pi}^{GC})$. Furthermore, it has the same value per match as in guaranteed choice, which match dominates any equilibrium outcome of independent lotteries by Proposition 8.

By Proposition 6, the value per match in the unique equilibrium outcome of the waitlist with deferral and the ticket-saving lottery is

$$\nu^{WD}(\alpha) = \max\left(\nu^{GC}(\alpha) - \frac{1-q}{q\pi^{GC}(\alpha)} c, 0\right),$$

where q is as defined in Proposition 6 and GC is the guaranteed choice outcome. Since $c \geq 0$ and $q \leq 1$, this is weakly smaller than $\nu^{GC}(\alpha)$, which is the same as the value per match under the common lottery.

Finally, the common lottery match dominates any equilibrium outcome of the single-entry lottery, because Proposition 9 shows that the latter has the same value per match as guaranteed choice with a higher participation cost. However, Proposition 8 implies that the value per match under guaranteed choice is weakly decreasing in the participation cost. □

Proof of Theorem 2 b). Perfect matching match dominates any equilibrium outcome because the value per match in any equilibrium outcome is equal to the expected value conditional on matching, minus the outside option, minus the expected total participation cost. The first term is upper-bounded by $\overline{v}$, and the third term is non-negative, so an overall upper-bound is $\overline{v} - \alpha$.

We now show that the common lottery converges to perfect matching when $\delta \to 1$. Define $GC_c^\delta(\alpha, v)$ to be the guaranteed choice outcome under continuation probability δ and participation cost c. Define $\phi_c^\delta(\alpha)$ to be the optimal acceptance threshold for agent type α as in Proposition 2, and $\pi_c^\delta(\alpha)$ the probability of matching. By Proposition 2,

$$GC_c^\delta(\alpha, v) = \mathbb{P}_{v' \sim G}(v \geq v' | v' \geq \phi_c^\delta(\alpha))\pi_c^\delta(\alpha). \tag{52}$$

For any $c \geq 0$ and $\epsilon < 1$, define $\overline{\delta}(\epsilon, c)$ to be the unique value of $\overline{\delta} < 1$ such that

$$\frac{\overline{\delta}}{1 - \overline{\delta}} = \frac{1 - \epsilon^2 + c}{H(\overline{v} - \epsilon^2)}. \tag{53}$$

Lemma 9. *For any $c \geq 0$, any $\epsilon > 0$ (such that $\epsilon < 1$), and any $\delta \geq \overline{\delta}(\epsilon, c)$ (such that $\delta < 1$), we have*

a) $\overline{v} - \epsilon \leq \phi_c^\delta(\alpha) \leq \overline{v}$ for all $\overline{v} - \frac{1}{\epsilon} \leq \alpha < \overline{v}$.

b) $(1 - \epsilon)\mathbb{1}(\alpha \leq \overline{v} - \epsilon) \leq \pi_c^\delta(\alpha) \leq \mathbb{1}(\alpha < \overline{v})$ for all $\alpha \in \mathbb{R}$.

Proof of Lemma 9. For part a), it suffices by Proposition 3 to show that $\phi_c^\delta(\alpha_3) \geq \overline{v} - \epsilon$ for $\alpha_3 = \overline{v} - \frac{1}{\epsilon}$. This is true because $H(\overline{v} - \epsilon) \geq \frac{1}{\epsilon} H(\overline{v} - \epsilon^2)$ by convexity of H and $H(\overline{v}) = 0$. So for any $x < \overline{v} - \epsilon$,

$$x - \alpha_3 < (\overline{v} - \epsilon) - (\overline{v} - \frac{1}{\epsilon}) \leq \frac{1}{\epsilon}\left(\frac{\delta}{1-\delta} H(\overline{v} - \epsilon^2) - c\right) \leq \frac{\delta}{1-\delta} H(\overline{v} - \epsilon) - c \leq \max\left(\frac{\delta}{1-\delta} H(x) - c, 0\right),$$

which means that x cannot satisfy the fixed point equation (16) determining $\phi_c^\delta(\alpha_3)$.

For part b), it suffices to show that anyone with outside option $\alpha \leq \overline{v} - \epsilon$ matches with probability at least $1 - \epsilon$. To do this, it suffices to show that $\pi_c^\delta(\alpha_4) \geq 1 - \epsilon$ for $\alpha_4 = \overline{v} - \epsilon$, since Lemma 2 implies that $\pi_c^\delta(\alpha)$ is weakly decreasing. Let $\phi_4 = \phi_c^\delta(\alpha_4)$. We first show that $\phi_4 \geq \overline{v} - \epsilon^2$. This is because for

$x \leq \overline{v} - \epsilon^2$, we have

$$x - \alpha_4 < (\overline{v} - \epsilon^2) - (\overline{v} - \epsilon) \leq \frac{\delta}{1-\delta} H(\overline{v} - \epsilon^2) - c \leq \max\left(\frac{\delta}{1-\delta} H(x) - c, 0\right),$$

so x cannot satisfy the fixed point equation (16) determining ϕ_4. Now,

$$\frac{\delta}{1-\delta}\overline{G}(\phi_4) \geq \frac{\frac{\delta}{1-\delta}\int_{\phi_4}^{\overline{v}} \overline{G}(x)\,dx}{\overline{v} - \phi_4} \geq \frac{\phi_4 - \alpha_4}{\overline{v} - \phi_4} \geq \frac{1-\epsilon}{\epsilon},$$

which implies that

$$\pi_c^{\delta}(\alpha_4) = \frac{\frac{\delta}{1-\delta}\overline{G}(\phi_4)}{1 + \frac{\delta}{1-\delta}\overline{G}(\phi_4)} \geq 1 - \epsilon.$$

□

Returning to the proof of Theorem 2 b), Lemma 9 implies that for any sequence $\delta_1, \delta_2, \cdots$ of such that $\lim_{n\to\infty} \delta_n = 1$ and $\delta_n < 1$, and any sequence $c_1, c_2, \cdots$, such that $c_n < \overline{c} < \infty$ for some constant $\overline{c}$, we have

$$\lim_{n\to\infty} GC_{c_n}^{\delta_n}(\alpha, v) = \mathbb{1}(\alpha < \overline{v})\mathbb{1}(v \geq \overline{v}).$$

Now, define $E_n(\alpha, v) = \mu GC_c^{\delta}(\alpha, v)/\overline{\pi}^{GC_c^{\delta}}$, then since the average match rate $\overline{\pi}^{GC_c^{\delta}} \to 1$, $E_n \to PM$.

This immediately implies the convergence of any sequence of the common lottery outcome to perfect matching, as long as the participation costs in this sequence is upper-bounded by a constant. Now, Proposition 9 implies that any sequence of equilibrium outcomes of the single-entry lottery with success probabilities (p_n) is outcome equivalent to a sequence of guaranteed choice outcomes with participation costs $c_n = \frac{c}{p_n}$. Moreover, each $p_n \geq \mu$, so $c_n < \frac{c}{\mu} < \infty$. Hence, any sequence of single-entry lottery outcomes also converges to perfect matching as $\delta \to 1$. □

F.4 Machinery Used in the Proofs of Theorems 3 and 4

F.4.1 Theory of Hazard Rate Dominance

The proofs of Theorems 3 and 4 make use of the following tools from theory of stochastic ordering. For ease of exposition, we develop the tools in general before applying to our problem.

Definition 14 (Hazard Rate Dominance)). *Let $a(x)$ and $b(x)$ be non-negative real-valued functions, with $A(x) := \int_x^{\infty} a(y)\,dy$ and $B(x) := \int_x^{\infty} b(y)\,dy$ both finite for all $x \in \mathbb{R}$. We say that function $a(x)$* hazard-rate dominates *$b(x)$ if $a(x)/A(x) \leq b(x)/B(x)$ wherever the denominators are both positive.*

Given a function $a(x)$ as in Definition 14, and given $y \in \mathbb{R}$ such that $A(y) > 0$, define random

variable X_y^a taking on values (y, ∞) with CDF

$$\mathbb{P}(X_y^a \leq x) = \frac{\int_y^x a(z)\, dz}{A(y)}. \tag{54}$$

Define random variable X_y^b similarly.

Lemma 10 (Equivalence Definitions of Hazard Rate Dominance)**.** *Let $a(x)$, $b(x)$, $A(x)$, and $B(x)$ be as in Definition 14. The following statements are equivalent.*

1. *$a(x)$ hazard rate dominates $b(x)$.*
2. *$A(x)/B(x)$ is weakly increasing in x wherever the denominator is positive.*
3. *X_y^a first order stochastic dominates X_y^b for every y such that both $A(y)$ and $B(y)$ are positive.*
4. *$a(x)h(x)$ hazard rate dominates $b(x)h(x)$ for every non-negative, weakly increasing, and bounded function $h(x)$.*

Our proof of Lemma 10 makes use of a well-known fact and a technical lemma given below.

Fact 1. *If X and Y are random variables taking values on set D, and if X first order stochastic dominates Y, then for any weakly increasing function $\gamma : D \to \mathbb{R}$, we have*

$$\mathbb{E}[\gamma(X)] \geq \mathbb{E}[\gamma(Y)].$$

Moreover, if γ is instead weakly decreasing, then the above inequality is reversed.

Lemma 11. *If $A(x)$ and $B(x)$ are continuous, weakly-decreasing and non-negative functions in x, and the ratio $A(x)/B(x)$ is weakly increasing in x whenever the denominator is positive, then $A(x) = 0$ implies $B(x) = 0$.*

Proof of Lemma 11. Define $\overline{x} = \sup\{x : A(x) > 0 \text{ and } B(x) > 0\}$. Suppose on the contrary that $A(\overline{x}) = 0$ while $B(\overline{x}) > 0$. Then for any $x < \overline{x}$, $\frac{A(x)}{B(x)} > 0 = \frac{A(\overline{x})}{B(\overline{x})}$, which contradicts the weakly decreasing property of $A(x)/B(x)$. □

Proof of Lemma 10. Define $\overline{x} = \sup\{x : A(x) > 0 \text{ and } B(x) > 0\}$. By Lemma 11, statement 2) is equivalent to $A(x)/B(x)$ being weakly increasing for $x \in (-\infty, \overline{x})$. We prove the equivalence in steps.

1) $\iff$ 2): The equivalence follows from the fact that for any $x < \overline{x}$, we have $A(x) > 0$ and $B(x) > 0$, and $A(x)/B(x)$ is differentiable, with

$$\frac{d}{dx} \log(\frac{A(x)}{B(x)}) = -\frac{a(x)}{A(x)} + \frac{b(x)}{B(x)}.$$

2) $\iff$ 3): The equivalence follows from the fact that for any $y < x < \overline{x}$, we have $A(y) \geq A(x) > 0$, and $B(y) \geq B(x) > 0$, with

$$\mathbb{P}(X_y^a \leq x) - \mathbb{P}(X_y^b \leq x) = \frac{B(x)}{B(y)} - \frac{A(x)}{A(y)}.$$

3) $\implies$ 4): If $h(x) = 0$ for every $x < \overline{x}$, then there is nothing to prove. Otherwise, we have that for any $y < \overline{x}$, both $\mathbb{E}[h(X_y^a)]$ and $\mathbb{E}[h(X_y^b)]$ are strictly positive and finite (strictly positive because $h(x) > 0$ in some open interval containing $\overline{x}$, and finite because h is bounded). Using condition 3) and Fact 1 we have

$$\frac{A(y)}{\int_y^\infty a(x)h(x)\,dx} = \frac{1}{\mathbb{E}[h(X_y^a)]} \leq \frac{1}{\mathbb{E}[h(X_y^b)]} = \frac{B(y)}{\int_y^\infty b(x)h(x)\,dx}.$$

Multiplying this by $h(y)a(y)/A(y) \leq h(y)b(y)/B(y)$ yields the desired condition.

4) $\implies$ 1): This is clear since $h(x) = 1$ is a non-negative, weakly increasing and bounded function.

□

A sufficient condition for hazard rate dominance is as follows.

Lemma 12 (Monotone Likelihood Ratio Implies Hazard Rate Dominance)**.** *Let $a(x)$ and $b(x)$ be as in Definition 14, then if $a(x)/b(x)$ is weakly increasing in x on the domain where the denominator is positive and if $b(x)$ is weakly decreasing in x everywhere, then $a(x)$ hazard rate dominates $b(x)$.*

Proof of Lemma 12. Under the conditions of the Lemma, for any x such that $A(x) > 0$, $B(x) > 0$, then it must be that $b(x) > 0$, and we have

$$\frac{A(x)}{B(x)} = \frac{\int_x^\infty a(y)\,dy}{B(x)} \geq \frac{\int_x^\infty \frac{a(x)}{b(x)} b(y)\,dy}{B(x)} = \frac{a(x)}{b(x)}.$$

So $a(x)/A(x) \leq b(x)/B(x)$, as desired. □

F.4.2 Strong Targeting Dominance

We prove that a notion of targeting dominance also implies match dominance, thus providing a partial converse to Theorem 3.

Definition 15. *We say that partial equilibrium outcome E'* **strongly targeting dominates** *E if the ratio $\pi^E(\alpha)/\pi^{E'}(\alpha)$ is weakly increasing in α wherever the denominator is positive.*[36]

It is possible to show that if no agent desires all possible units ($\underline{\alpha} \geq \underline{v}$), and either

- G follows a uniform or an exponential distribution,

[36] This condition states that the ratio of the density of matched agents under E to the density of matched agents under E' is increasing – that is, the densities associated with F^E and $F^{E'}$ have a monotone likelihood ratio. Because monotone likelihood ratio implies first order stochastic stochastic dominance, strong targeting dominance implies targeting dominance.

- or G is heavy-tailed and μ is sufficiently small,

then any equilibrium outcome of independent lotteries strongly targeting dominates the unique equilibrium outcome of the common lottery.

Proposition 10. *Any partial equilibrium outcome E strongly targeting domintes perfect matching (PM from Definition 4).*

Proof of Proposition 10. This follows immediately from the weakly-decreasing property of the allocation function π^E from Lemma 2, and the fact that PM is constant in the domain $(-\infty, \overline{v})$. □

Proposition 11. *If partial equilibrium outcome E' strongly targeting dominates E', then E match dominates E'.*

Proof of Proposition 11. By Lemma 2, $a(x) = \pi^E(x)$ and $b(x) = \pi^{E'}(x)$ satisfy the condition of Lemma 12. So π^E hazard rate dominates $\pi^{E'}$, which implies match dominance by Lemma 13. □

F.5 Matching and Targeting Tradeoff (Proofs of Theorems 3 and 4)

The following lemma allows us to analyze match dominance using the machinery of hazard rate dominance developed in Section F.4.1.

Lemma 13. *An equilibrium outcome E match dominates outcome E' if and only if the allocation function π^E hazard rate dominates $\pi^{E'}$.*

Proof of Lemma 13. For any equilibrium outcome E, by the definition of value per match and by Proposition 1, we have

$$\nu^E(\alpha) = \frac{u^E(\alpha)}{\pi^E(\alpha)} = \frac{\int_\alpha^\infty \pi^E(x)\,dx}{\pi^E(\alpha)}. \tag{55}$$

Therefore, match dominance implies that

$$\frac{\pi^E(\alpha)}{\int_\alpha^\infty \pi^E(x)\,dx} = \frac{\pi^E(\alpha)}{u^E(\alpha)} \leq \frac{\pi^{E'}(\alpha)}{u^{E'}(\alpha)} = \frac{\pi^{E'}(\alpha)}{\int_\alpha^\infty \pi^{E'}(x)\,dx}, \tag{56}$$

wherever the $u^E(\alpha)$ and $u^{E'}(\alpha)$ are both positive. So π^E hazard rate dominates $\pi^{E'}$.

Because (55) is an identity, the above argument can be used to prove the converse implication when all four of the following are strictly positive: $\pi^E(\alpha)$, $\pi^{E'}(\alpha)$,$u^E(\alpha)$, and $u^{E'}(\alpha)$. To complete the proof, it suffices to check that if π^E hazard rate dominates $\pi^{E'}$, then $u^E(\alpha) = 0$ implies $u^{E'}(\alpha) = 0$, which follows by Lemma 11. □

Proof of Theorem 3. Define

$$h(x) = \begin{cases} f(x) & \text{if } x \leq \overline{\alpha} \\ f(\overline{\alpha}) & \text{otherwise.} \end{cases} \tag{57}$$

This is the smallest function larger than f that is weakly increasing throughout $\mathbb{R}$. If $\pi^E(\overline{\alpha}) = 0$, then $u^E(\overline{\alpha}) = 0$, which implies by match dominance that $u^{E'}(\overline{\alpha}) = 0$. This in turn implies that $\pi^E(x) = \pi^{E'}(x) = 0$ for any $x > \overline{\alpha}$. Therefore,

$$F^E(\alpha) = \frac{\int_{-\infty}^{\alpha} \pi^E(x)\, dF(x)}{\int_{-\infty}^{\infty} \pi^E(x)\, dF(x)} = \frac{\int_{-\infty}^{\alpha} \pi^E(x) h(x)\, dx}{\int_{-\infty}^{\infty} \pi^E(x) h(x)\, dx}. \tag{58}$$

An analogous equation holds for $F^{E'}(\alpha)$. By Lemma 10, condition 3), it suffices to prove that $\pi^E(x)h(x)$ hazard rate dominates $\pi^{E'}(x)h(x)$, which follows by Lemma 10, condition 4), since $h(x)$ is non-negative, weakly increasing and bounded, while $\pi^E(x)$ hazard rate dominates $\pi^{E'}(x)$ by Lemma 13. □

Proof of Theorem 4. We prove the following more general of Theorem 4, in which the condition $\overline{\alpha} \geq \overline{v}$ is replaced by the weaker conditions of $\pi^{E'}(\overline{\alpha}) = 0$ and $\pi^E(\overline{\alpha}) = 0$. Moreover, we also include alternative condition based on match dominance and strong targeting dominance (Definition 15).

Theorem 7 (Generalization of Theorem 4)**.** *Let E and E' be equilibrium outcomes. If any of the following conditions are satisfied:*

- *E match dominates E' and f is weakly increasing in $(\underline{\alpha}, \overline{\alpha})$, or*
- *E' targeting dominates E and $\pi^{E'}(\overline{\alpha}) = 0$, or*
- *E' strongly targeting dominates E,*

then the following hold:

a) If F has a light left tail, then $W^E \geq W^{E'}$.

b) If F has a heavy left tail and $\pi^E(\overline{\alpha}) = 0$, then $W^E \leq W^{E'}$.

For clarity, we split up the proof into the following two parts.

Lemma 14. *If equilibrium outcome E' targeting dominates E, then*

a) $W^E \geq W^{E'}$ if F has a light left tail and $\pi^{E'}(\overline{\alpha}) = 0$.

b) $W^E \leq W^{E'}$ if F has a heavy left tail and $\pi^E(\overline{\alpha}) = 0$.

Lemma 15. *If F has a light left tail, then $W^E \geq W^{E'}$ if either*

a) E' strongly targeting dominates E, or

b) E match dominates E' and f is weakly increasing in $(\underline{\alpha}, \overline{\alpha})$.

These two lemmas together imply Theorem 7, because when F has a heavy left tail, then f is guaranteed to be weakly increasing on its domain. If in addition we have $\pi^E(\overline{\alpha}) = 0$, then all three conditions of Theorem 4 imply targeting dominance[37], so Theorem 7 reduces to part b) of Lemma 14.

[37] Appendix F.5 shows that if E match dominates E', f is increasing on $(\underline{\alpha}, \overline{\alpha})$ and $\pi^E(\overline{\alpha}) = 0$, then E' targeting dominates E. Furthermore, strong targeting dominance always implies targeting dominance.

When F has a light left tail, Theorem 7 reduces to part a) of Lemma 14 and to both parts of Lemma 15.

□

The proof of both Lemma 14 and 15 are based on the following identity.

Lemma 16. *The welfare of any equilibrium outcome E can be expressed as*

$$W^E = \frac{\int_{-\infty}^{\infty} \pi^E(x) F(x)\, dx}{\int_{-\infty}^{\infty} \pi^E(x)\, dF(x)} \tag{59}$$

$$= \mathbb{E}_{\alpha \sim F^E}[\gamma(\alpha)] + u^E(\overline{\alpha})/\mu, \tag{60}$$

where $\gamma(\alpha) := F(\alpha)/f(\alpha)$ for $\alpha \in (\underline{\alpha}, \overline{\alpha})$, and $u^E(\infty) := 0$ for convenience.

Proof of Lemma 14. Consider the first term of (60). Suppose that F is light-tailed, so that γ is increasing. In that case, we have $\mathbb{E}_{\alpha \sim F^E}[\gamma(\alpha)] \geq \mathbb{E}_{\alpha \sim F^{E'}}[\gamma(\alpha)]$. If $\pi^{E'}(\overline{\alpha}) = 0$, then by Proposition 1 we have $u^{E'}(\overline{\alpha}) = 0 \leq u^E(\overline{\alpha})$, implying that $W^E \geq W^{E'}$. Analogous logic implies that if F is heavy-tailed (γ is decreasing) and $\pi^E(\overline{\alpha}) = 0$, then $W^{E'} \geq W^E$. □

Proof of Lemma 15. We first prove part a). Because E' targeting dominates E and γ is weakly increasing, we have that

$$\mathbb{E}_{\alpha \sim F^E}[\gamma(\alpha)] \geq \mathbb{E}_{\alpha \sim F^{E'}}[\gamma(\alpha)].$$

By (60), it suffices to prove that $u^E(\overline{\alpha}) \geq u^{E'}(\overline{\alpha})$. Suppose that $u^{E'}(\overline{\alpha}) = 0$, then we have nothing to prove. Suppose that $u^{E'}(\overline{\alpha}) > 0$, then $\pi^{E'}(\overline{\alpha}) > 0$, and by strong targeting dominance

$$\frac{u^E(\overline{\alpha})}{u^{E'}(\overline{\alpha})} = \frac{\int_{\overline{\alpha}}^{\infty} \pi^E(x)\, dx}{\int_{\overline{\alpha}}^{\infty} \pi^{E'}(x)\, dx} \geq \frac{\pi^E(\overline{\alpha})}{\pi^{E'}(\overline{\alpha})} \geq \frac{\int_{-\infty}^{\overline{\alpha}} \pi^E(x)\, dF(x)}{\int_{-\infty}^{\overline{\alpha}} \pi^{E'}(x)\, dF(x)} = 1.$$

The first inequality above follows from $\pi^E(x) \geq \frac{\pi^E(\overline{\alpha})}{\pi^{E'}(\overline{\alpha})} \pi^{E'}(x)$ for all $x \geq \overline{\alpha}$, and the second from the reverse inequality when $x \leq \overline{\alpha}$. The last inequality follows because both equilibria induce the same average match rate μ by definition.

We now prove part b). Define $h(x)$ as in (57). By Lemma 10, $\pi^E(x)h(x)$ hazard rate dominates $\pi^{E'}(x)h(x)$, and

$$\frac{\int_{-\infty}^{\overline{\alpha}} \pi^E(x)h(x)\, dx}{\int_{-\infty}^{\infty} \pi^E(x)h(x)\, dx} \leq \frac{\int_{-\infty}^{\overline{\alpha}} \pi^{E'}(x)h(x)\, dx}{\int_{-\infty}^{\infty} \pi^{E'}(x)h(x)\, dx} \tag{61}$$

Define random variables X and X' with densities proportional to $\pi^E(x)h(x)$ and $\pi^{E'}(x)h(x)$ respectively. This is well defined because the integral $\int_{-\infty}^{\infty} \pi^E(x)h(x)\, dx = f(\overline{\alpha})u^E(\overline{\alpha}) + \mu < \infty$. We

have that X first order stochastic dominates X'. Extend $\gamma(x)$ to domain $\mathbb{R}$ as follows,

$$\gamma(x) = \begin{cases} 0 & \text{if } x \leq \underline{\alpha} \\ F(x)/f(x) & \text{if } x \in (\underline{\alpha}, \overline{\alpha}) \\ 1/f(\overline{\alpha}) & \text{if } x \geq \overline{\alpha}. \end{cases} \tag{62}$$

Since F has a light tail, $\gamma(x)$ is weakly increasing on $\mathbb{R}$. Moreover, $\gamma(x)h(x) = F(x)$ everywhere. Therefore,

$$\frac{\int_{-\infty}^{\infty} \pi^E(x)F(x)\,dx}{\int_{-\infty}^{\infty} \pi^E(x)h(x)\,dx} = \mathbb{E}[\gamma(X)] \geq \mathbb{E}[\gamma(X')] = \frac{\int_{-\infty}^{\infty} \pi^{E'}(x)F(x)\,dx}{\int_{-\infty}^{\infty} \pi^{E'}(x)h(x)\,dx}. \tag{63}$$

Dividing inequality (63) by (61), and noting that $h(x) = f(x)$ on the domain $(-\infty, \overline{\alpha})$, we get that $W^E \geq W^{E'}$, as desired. □

Proof of Lemma 16. Equation (59) follows from Proposition 1 and switching the order of integrals.

$$\begin{aligned} W^E \cdot \overline{\pi}^E &= \int_{-\infty}^{\infty} u^E(\alpha) dF(\alpha) \\ &= \int_{-\infty}^{\infty} \int_{\alpha}^{\infty} \pi^E(x)\,dx\,dF(\alpha) \\ &= \int_{-\infty}^{\infty} \pi^E(x) \int_{-\infty}^{x} dF(\alpha)\,dx \\ &= \int_{-\infty}^{\infty} \pi^E(x)F(x)\,dx, \end{aligned}$$

Equation (60) follows from the fact that for equilibrium outcomes, $\overline{\pi}^E = \mu$, and that the last integral above can be split into two components,

$$\begin{aligned} \int_{-\infty}^{\infty} \pi^E(x)F(x)\,dx &= \int_{-\infty}^{\overline{\alpha}} \pi^E(x)\gamma(x)\,dF(x) + \int_{\overline{\alpha}}^{\infty} \pi^E(x)\,dx \\ &= \overline{\pi}^E \mathbb{E}_{\alpha \sim F^E}[\gamma(\alpha)] + u^E(\overline{\alpha}), \end{aligned}$$

where the last line follows from $\overline{\pi}^E = \int_{-\infty}^{\overline{\alpha}} \pi^E(x) dF(x)$, and from $\int_{\overline{\alpha}}^{\infty} \pi^E(x)\,dx = u^E(\overline{\alpha})$ (Proposition 1). □

F.6 Welfare Optimality of Perfect Matching (Proof of Theorem 5)

Part a) of Theorem 5 follows from Theorem 7 in Appendix F.5, as every equilibrium outcome strongly targeting dominates perfect matching (see Proposition 10 in section F.4.2).

Proof of Theorem 5 b). Let E be any equilibrium outcome with allocation function $\pi(\alpha)$. Note that there must exist an α with $\pi(\alpha) \leq \mu$, as otherwise $\overline{\pi}^E > \mu$. Define $\beta := \inf\{x : \pi(x) \leq \mu\}$.

If $\beta \leq \underline{\alpha}$, then $\pi(\alpha) \leq \mu$ for all $\alpha \geq \underline{\alpha}$, so $W^E \leq W^{PM}$ by Lemma 16. If $\beta \geq \overline{\alpha}$, then it must be $\pi(\alpha) = \mu$ for all $\alpha \in (\underline{\alpha}, \overline{\alpha})$, so we again have $\pi(\alpha) \leq \mu$ for all $\alpha \geq \underline{\alpha}$, since $\pi(x)$ is weakly decreasing by Lemma 2.

It remains to consider the case when $\beta \in (\underline{\alpha}, \overline{\alpha})$. Define

$$\theta := \int_{\underline{\alpha}}^{\beta} \pi(x)\, dF(x) \geq \mu F(\beta).$$

We have by the identity $\overline{\pi}^E = \mu$ that

$$\theta - \mu F(\beta) = (1 - F(\beta))\mu - \int_{\beta}^{\overline{\alpha}} \pi(x)\, dF(x) \leq (1 - F(\beta))(\mu - \pi(\overline{\alpha})) \leq \mu - \pi(\overline{\alpha}). \tag{64}$$

Moreover, we have

$$\begin{aligned}
\int_{\underline{\alpha}}^{\beta} (\pi(x) - \mu) F(x)\, dx &= \int_{\underline{\alpha}}^{\beta} \int_{x}^{\beta} \pi(y) - \mu\, dy\, dF(x) \\
&\leq \int_{\underline{\alpha}}^{\beta} (\pi(x) - \mu)(\beta - x)\, dF(x) \\
&\leq (\overline{\alpha} - \underline{\alpha}) \int_{\underline{\alpha}}^{\beta} (\pi(x) - \mu)\, dF(x) \\
&\leq (\overline{v} - \overline{\alpha})(\theta - \mu F(\beta)).
\end{aligned}$$

In addition,

$$\begin{aligned}
\int_{\beta}^{\overline{v}} (\pi(x) - \mu) F(x)\, dx &= \int_{\beta}^{\overline{\alpha}} (\pi(x) - \mu) F(x)\, dx + \int_{\overline{\alpha}}^{\overline{v}} (\pi(x) - \mu)\, dx \\
&\leq 0 - (\overline{v} - \overline{\alpha})(\mu - \pi(\overline{\alpha})).
\end{aligned}$$

Adding the two above inequalities and applying Lemma 16, we have

$$\frac{\mu(W^E - W^{PM})}{\overline{v} - \overline{\alpha}} = \frac{\int_{\underline{\alpha}}^{\overline{v}} (\pi(x) - \mu) F(x)\, dx}{\overline{v} - \overline{\alpha}} \leq \theta - \mu F(\beta) - (\mu - \pi(\overline{\alpha})) \leq 0,$$

where the last inequality is by (64). So $W^E \leq W^{PM}$, as desired. □

F.7 Efficacy of Costly Guaranteed Choice for Targeting (Proof of Theorem 6)

Proposition 12.

a) Every market clearing cost $c' \leq \frac{\delta}{1-\delta} H(F^{-1}(\mu))$.

b) If the average match rate $\overline{\pi}^{GC_c} \geq \mu$, then there exists a market clearing cost $c' \geq c$.

c) Under Assumption 1, a highest market clearing cost $\overline{c}$ always exists.

Proof of Proposition 12. Let GC_c be the guaranteed choice outcome under participation cost c. For part a), note that for any $c' > \frac{\delta}{1-\delta}H(F^{-1}(\mu))$, we have $\overline{\pi}^{GC_{c'}} \leq F(H^{-1}(c')) < \mu$. Hence, c' cannot be a market clearing cost.

For part b), note that $\overline{\pi}^{GC_c} = \int_{-\infty}^{\infty} \pi^{GC_c}(\alpha) f(\alpha)\, d\alpha$ is continuous in c, because by Proposition 2, $\pi^{GC_c}(\alpha)$ is continuous in c except at $\alpha = H^{-1}(\frac{1-\delta}{\delta}c)$. Now, since $\pi^{GC_c} \geq \mu$, and $\pi^{GC_{c''}} < \mu$ for any $c'' > \frac{\delta}{1-\delta}H(F^{-1}(\mu))$, continuity implies that there must exists a market clearing cost $c' \geq c$.

For part c), note that $\overline{c} = \max\{c : \overline{\pi}^{GC_c} \geq \mu\}$. The maximum exists because the set is non-empty by Assumption 1, closed by the continuity of $\overline{\pi}^{GC_c}$ as a function of c, and has a finite upper-bound as established in a). □

Proof of Theorem 6 a). Let $\overline{c}$ be the largest market clearing cost. Let E be an equilibrium outcome of the common lottery, waitlist with deferral, or the ticket-saving lottery with participation cost c. We show that

$$\pi^E(\alpha) \leq \pi^{GC_c}(\alpha) \leq \pi^{GC_{\overline{c}}}(\alpha) \qquad \text{for any } \alpha \leq H^{-1}\left(\frac{1-\delta}{\delta}\overline{c}\right). \tag{65}$$

The first inequality holds

- for the common lottery because in this case $\pi^E(\alpha)$ is equal to $\pi^{GC_c}(\alpha)$ times the probability of being selected.
- for the waitlist with deferral (and equivalently the ticket-saving lottery), because in this case $\pi^E(\alpha)$ is either zero or equal to $\pi^{GC_c}(\alpha)$ times the quantity q as defined in Proposition 6.

The second inequality of (65) holds by the fact that $\overline{c} \geq c$ (Proposition 12 b)), the fact that $\pi^{GC_c}(\alpha)$ is weakly decreasing in α, and the following identity, which follows directly from Proposition 2.

Lemma 17. *If $c' \geq c$, with $\Delta = c' - c \geq 0$,*

$$GC_{c'}(\alpha, v) = \begin{cases} GC_c(\alpha - \Delta, v) & \text{if } \alpha \leq \alpha_0 = H^{-1}(\frac{(1-\delta)}{\delta}c'), \\ 0 & \text{otherwise.} \end{cases}$$

Note that by Proposition 2,

$$\pi^{GC_{\overline{c}}}(\alpha) = 0 \leq \pi^E(\alpha) \qquad \text{for any } \alpha > H^{-1}\left(\frac{1-\delta}{\delta}\overline{c}\right). \tag{66}$$

By the definition of the match distribution F^E, and by Lemma 18 (given below), inequalities (65) and (66) together imply that $GC_{\overline{c}}$ targeting dominates E.

To complete the proof, note that by Proposition 9, any equilibrium outcome of the single-entry lottery is equivalent to a equilibrium outcome of the common lottery with a higher participation cost, so it is targeting dominated by the costly guaranteed choice outcome as well. □

Lemma 18. *Given two continuous distributions with densities f_1 and f_2. If there is a threshold x^* such that*

$$f_1(x) \leq f_2(x) \quad \textit{for all } x \leq x^*,$$
$$f_1(x) \geq f_2(x) \quad \textit{for all } x > x^*,$$

then the first distribution first order stochastic dominates the second: $F_1(x) \leq F_2(x)$ for all $x \in \mathbb{R}$.

Proof of Lemma 18. Let the CDFs of the two distribution be F_1 and F_2. It suffices to prove that for all $x \in \mathbb{R}$, $F_1(x) \leq F_2(x)$. For any $x \leq x^*$, this is true because

$$F_1(x) = \int_{-\infty}^{x} f_1(y)\, dy \leq \int_{-\infty}^{x} f_2(y)\, dy = F_2(x).$$

For any $x > x^*$, this is true because

$$1 - F_1(x) = \int_{x}^{\infty} f_1(y)\, dy \geq \int_{x}^{\infty} f_2(y)\, dy = 1 - F_2(x).$$

□

Proof of Theorem 6 b). We first show that PT targeting dominates any equilibrium outcome E. This holds because for $\alpha \leq F^{-1}(\mu)$,

$$F^E(\alpha) = \frac{1}{\mu} \int_{-\infty}^{\alpha} \pi^E(x) dF(x) \leq \frac{1}{\mu} \int_{-\infty}^{\alpha} dF(x) = F(\alpha)/\mu = F^{PT}(\alpha),$$

and for $\alpha > F^{-1}(\mu)$, we have $F^{PT}(\alpha) = 1 \geq F^E(\alpha)$.

We now show that costly guaranteed choice converges to perfect targeting when $\delta \to 1$. Define $GC_c^\delta(\alpha, v)$, $\pi_c^\delta(\alpha)$, $\pi_c^\delta(\alpha)$, and $\overline{\delta}(\epsilon, c)$ as in the proof of Theorem 2 b) in Appendix F.3. Without loss of generality, let $\mu < 1$, otherwise guaranteed choice with $c = 0$ would be feasible. So $F^{-1}(\mu) < \overline{\alpha} \leq \overline{v}$. Define $\epsilon > 0$ to be any small constant such that $F^{-1}(\frac{\mu}{1-\epsilon_0}) < \overline{v} - \epsilon_0$.

By the expression of GC_c^δ in terms of π_c^δ and ϕ_c^δ in Equation (52) (from proof of Theorem 2 b) in Appendix F.3), Theorem 6 b) follows from the following lemma, which implies as $\delta \to 1$, all agent types $\alpha > F^{-1}(\mu)$ will not participate under costly guaranteed choice, and any type $\alpha \leq F^{-1}(\mu)$ will participate and set an acceptance threshold of $F^{-1}(\mu)$.

□

Lemma 19. *For any $\epsilon > 0$ (such that $\epsilon \leq \epsilon_0$), and any $\delta \geq \overline{\delta}(\epsilon, 0)$ (such that $\delta < 1$), we have*

a) $(1-\epsilon)\mathbb{1}(\alpha \leq F^{-1}(\mu)) \leq \pi_c^\delta(\alpha) \leq \mathbb{1}(\alpha \leq F^{-1}(\frac{\mu}{1-\epsilon}))$ *for any* $\alpha \in \mathbb{R}$.

b) $F^{-1}(\mu) - \epsilon \leq \phi_c^\delta(\alpha) \leq F^{-1}(\frac{\mu}{1-\epsilon})$ *for any* $\alpha \in [F^{-1}(\mu) - \frac{1}{\epsilon}, F^{-1}(\frac{\mu}{1-\epsilon})]$.

Proof of Lemma 19. Let $\alpha_0 = H^{-1}(\frac{c(1-\delta)}{\delta})$ be the type that is indifferent between participating and leaving as in Proposition 2. By Lemma 9 in Appendix F.3 and by Lemma 17 above, we have that for any $\delta \geq \overline{\delta}(\epsilon, 0)$, and any $\alpha \leq \min(\overline{v} - \epsilon, \alpha_0)$,

$$\pi_c^\delta(\alpha) \geq \pi_0^\delta(\alpha) \geq 1 - \epsilon.$$

Thus, since $F^{-1}(\frac{\mu}{1-\epsilon}) \leq \overline{v} - \epsilon$ for any $\epsilon \leq \epsilon_0$, we have

$$\pi_c^\delta(\alpha) \in \{0\} \cup [1-\epsilon, 1] \qquad \text{for all } \alpha \leq F^{-1}(\frac{\mu}{1-\epsilon}),$$

where $\pi_c^\delta(\alpha) = 0$ if $\alpha > \alpha_0$. The above implies that $\pi_c^\delta(\alpha) = 0$ for any $\alpha > F^{-1}(\frac{\mu}{1-\epsilon})$, because otherwise the total match rate would exceed μ by the monotonicity of $\pi_c^\delta(\alpha)$ (see Lemma 2). Furthermore, $\pi_c^\delta(\alpha) \geq 1 - \epsilon$ for any $\alpha \leq F^{-1}(\mu)$, as otherwise the overall match rate would be strictly less than μ. This proves part a).

For part b), note that part a) implies that $F^{-1}(\mu) \leq \alpha_0 \leq F^{-1}(\frac{\mu}{1-\epsilon})$. Proposition 3 implies that for any $\alpha < \alpha_0$, $\phi_c^\delta(\alpha) \leq \alpha_0 \leq F^{-1}(\frac{\mu}{1-\epsilon})$. Now, let $\alpha_3 = F^{-1}(\mu) - \frac{1}{\epsilon}$ and $\phi_3 = \phi_c^\delta(\alpha)$. It suffices to show that $\phi_3 \geq F^{-1}(\mu) - \epsilon$. This is because for any $x < F^{-1}(\mu) - \epsilon$, we have

$$x - \alpha_3 < (F^{-1}(\mu) - \epsilon) - (F^{-1}(\mu) - \frac{1}{\epsilon}) \leq \frac{\epsilon\delta}{1-\delta}\overline{G}(F^{-1}(\mu)) \leq \max\left(\frac{\delta}{1-\delta}H(x) - c, 0\right),$$

so x cannot satisfy the fixed point equation (16) determining ϕ_3. In the above, the second inequality holds because for any $\epsilon \leq \epsilon_0$, $\delta \geq \overline{\delta}(\epsilon, 0)$, the convexity of H implies $H(\overline{v} - \epsilon^2) \leq \epsilon^2\overline{G}(\overline{v} - \epsilon^2) \leq \epsilon^2\overline{G}(F^{-1}(\mu))$, so $\frac{\delta\epsilon}{1-\delta}\overline{G}(F^{-1}(\mu) \geq \frac{1}{\epsilon} - \epsilon$ by the definition of $\delta(\epsilon, 0)$ in Equation (53). The third inequality holds because

$$\max\left(\frac{\delta}{1-\delta}H(x) - c, 0\right) \geq \frac{\delta}{1-\delta}\int_x^{F^{-1}(\mu)} \overline{G}(y)\, dy \geq \frac{\epsilon\delta}{1-\delta}\overline{G}(F^{-1}(\mu)).$$

□

F.8 Alternative Conditions for Costly Guaranteed Choice Targeting Dominating Independent Lotteries

Theorem 6 b) implies that cost guaranteed choice targeting dominates independent lotteries when δ is sufficiently high. Proposition 13 establishes the same comparison under alternative conditions.

Definition 16. *The value distribution G is said to be light tailed if the conditional expectation $\mathbb{E}_{v\sim G}[v-x|v \geq x]$ is weakly decreasing in $x \in (-\infty, \overline{v})$.*[38]

Proposition 13. *Any costly guaranteed choice outcome targeting dominates any equilibrium outcome of independent lotteries and the waitlist without deferral if either of the following conditions hold:*

a) G is light tailed (Definition 16), or

b) $c = 0$ and μ is less than a threshold $\mu_0 > 0$ that depends on F, G and δ,

Proof of Proposition 13. Let $\alpha_3 = H^{-1}(\frac{1-\delta}{\delta}\overline{c})$. Note that $\pi^{GC_{\overline{c}}}(\alpha) = 0 \leq \pi^{IL_p}(\alpha)$ for any $\alpha > \alpha_3$. By Lemma 18, it suffices to prove the following.

Lemma 20. *Fix F, G, δ. If either of the following hold:*

a) G is light-tailed, or

b) $c = 0$ and $\mu < \mu_0$ (a constant depending on F, G, δ),

then if $\overline{c}$ is the largest possible market clearing cost and p is an equilibrium success probability of independent lotteries, we have $\pi^{GC_{\overline{c}}}(\alpha) \geq \pi^{IL_p}(\alpha)$ for all $\alpha \leq \alpha_3 := H^{-1}(\frac{1-\delta}{\delta}\overline{c})$.

Part a) of Lemma 20 follows from the following two comparative statics.

i) When G is light-tailed, increasing the success probability p of independent lotteries causes all agents to match at weakly higher rates. (Follows from Lemma 21, given below.)

ii) When $p = 1$, increasing c to $\overline{c}$ causes all agents with $\alpha \leq \alpha_3$ to match at higher rates. (Follows from Lemma 17.)

For part b) of Lemma 20, we construct the threshold $\mu_0 > 0$ as follows. Choose any $c_0 \in (0, \frac{\delta}{1-\delta}H(\underline{\alpha}))$. Define

$$\mu_0 = \min\left(\min_{c' \in [0, c_0]} \overline{\pi}^{GC_{c'}}, (1-\delta)\overline{\pi}^{GC_0}\pi^{GC_0}(H^{-1}(\frac{1-\delta}{\delta}c_0))\right). \tag{67}$$

Note that both of the two minimands are strictly positive. The first is positive because by Proposition 2, for any $c' \in [0, c_0]$, agents with $\alpha < H^{-1}(\frac{1-\delta}{\delta}c_0)$ participate and match with positive probability under guaranteed choice, and there is a positive measure of such agents since $c_0 < \frac{\delta}{1-\delta}H(\underline{\alpha})$. Furthermore

[38] A sufficient condition is that $(1 - G(v))/g(v)$ is weakly decreasing in the domain $v \in (\underline{v}, \overline{v})$.

$\overline{\pi}^{GC_c}$ is continuous in c so the minimum over the compact set $[0, c_0]$ is attained within the set. The second minimand is positive because $\pi^{GC_0}(\alpha) > 0$ for any $\alpha < \overline{v}$ and $H^{-1}(\frac{1-\delta}{\delta}c_0) < \overline{v}$ since $c_0 > 0$.

For any supply-demand ratio μ less than the threshold μ_0, we have that $\overline{c} \geq c_0$, because otherwise $\overline{\pi}^{GC_{\overline{c}}} \geq \mu_0 > \mu$, which contradicts $\overline{c}$ being market clearing. This along with Lemma 22 (see below) implies

$$\mu_0 > \mu = \mathbb{E}_{\alpha\sim F}[\pi^{IL_p}(\alpha)] \geq \mathbb{E}_{\alpha\sim F}[p\pi^{GC_0}(\alpha)] = p\overline{\pi}^{GC_0}. \tag{68}$$

Therefore, for any $\alpha \leq H^{-1}(\frac{1-\delta}{\delta}\overline{c})$,

$$\pi^{IL_p}(\alpha) \leq \frac{p}{1-\delta} < \frac{\mu_0}{(1-\delta)\overline{\pi}^{GC_0}} \leq \pi^{GC_0}(H^{-1}(\frac{1-\delta}{\delta}c_0)) \leq \pi^{GC_0}(H^{-1}(\frac{1-\delta}{\delta}\overline{c})) \leq \pi^{GC_0}(\alpha) \leq \pi^{GC_{\overline{c}}}(\alpha).$$

The first inequality follows from Lemma 22 (see below); the second from (68); the third from the definition of μ_0; the fourth because both π^{GC_0} and H^{-1} are weakly decreasing and $\overline{c} \geq c_0$; the fifth because π^{GC_0} is weakly decreasing; and the last from Lemma 17. □

Lemma 21. *Suppose that G is light tailed, then the match probability for every agent under independent lotteries is weakly increasing in the success probability p.*

Proof of Lemma 21. We first prove that the lemma in the special case when $c = 0$. For $\alpha \geq \overline{v}$, the match probability is always zero, so is trivially weakly increasing in p. For $\alpha < \overline{v}$, we have by rearrangement of (16) with $c = 0$,

$$\frac{p\delta}{1-\delta}\overline{G}(\phi(\alpha, p)) = \frac{\phi(\alpha, p) - \alpha}{\mathbb{E}[v - \phi(\alpha, p) | v \geq \phi(\alpha, p)]}.$$

By Proposition 3, $\phi(\alpha, p)$ is weakly increasing in p, so the numerator is weakly increasing in p. For $\alpha < \overline{v}$, the denominator is positive, and is weakly decreasing in p by G being light-tailed. Thus, the entire expression above is weakly increasing in p. The desired result follows from the fact that the match probability $\pi(\alpha, p)$ is equal to the above under the monotone transformation $y = \frac{x}{1+x}$.

We now generalize the above result to the case with $c > 0$. Define IL_p^c to be independent lotteries outcome with success probability p and participation cost c. We have by Proposition 2 the identity

$$IL_p^c = \begin{cases} IL_p^0(\alpha - c) & \text{for all } \alpha \leq H^{-1}\left(\frac{(1-\delta)c}{p\delta}\right), \\ 0 & \text{otherwise.} \end{cases} \tag{69}$$

Let $0 < p < p'$ be two success probabilities. We have already established that $IL_p^0(\alpha) \leq IL_{p'}^0(\alpha)$ for all α. For any $c > 0$, note that $H^{-1}(\frac{(1-\delta)c}{p\delta}) < H^{-1}(\frac{(1-\delta)c}{p'\delta})$ since H^{-1} is strictly decreasing in the domain $(0, \infty)$. Hence, the above identity implies that $IL_p^c(\alpha) \leq IL_{p'}^c(\alpha)$, as desired. □

Lemma 22. *Fix F, G, δ. For any $p \in (0,1)$, let π^{IL_p} be the allocation function under independent lotteries with success probability p and no participation cost, as defined in* (22)*. Then*

$$p\,\pi^{GC_0}(\alpha) \leq \pi^{IL_p}(\alpha) \leq \frac{p\delta}{1-\delta}.$$

Proof of Lemma 22. By Proposition 3, $\phi(\alpha, p) \leq \phi(\alpha, 1)$. Since $\overline{G}(x)$ is increasing and $\overline{G}(\phi(\alpha, p)) \leq 1$, Proposition 2 implies the following:

$$\pi^{IL_p}(\alpha) = \frac{\delta p \overline{G}(\phi(\alpha,p))}{1-\delta+\delta p\overline{G}(\phi(\alpha,p))} \geq \frac{\delta p \overline{G}(\phi(\alpha,1))}{1-\delta+\delta p\overline{G}(\phi(\alpha,1))} \geq p\frac{\delta \overline{G}(\phi(\alpha,1))}{1-\delta+\delta \overline{G}(\phi(\alpha,1))} = p\pi^{GC_0}(\alpha).$$
$$\pi^{IL_p}(\alpha) = \frac{\delta p \overline{G}(\phi(\alpha,p))}{1-\delta+\delta p\overline{G}(\phi(\alpha,p))} \leq \frac{p\delta}{1-\delta+p\delta} \leq \frac{p\delta}{1-\delta}.$$

□

F.9 Convergence of Independent Lotteries to Random Matching

Proof of Proposition 5. By Lemma 23 (stated after this proof), we can define a threshold $\mu_0 > 0$ such that if $\mu < \mu_0$, then for any equilibrium outcome of independent lotteries, the success probability p satisfies $\frac{p\delta}{1-\delta}H(\underline{v}) < \underline{v} - \overline{\alpha}$.

By Proposition 2, it suffices to show that $\phi(\alpha, p) \leq \underline{v}$ for all $\alpha \leq \overline{\alpha}$, as then every agent would be accepting every development and the outcome must be random matching. By Proposition 3 a), it suffices to show that $\phi(\overline{\alpha}, p) \leq \underline{v}$. Suppose on the contrary that $\underline{v} < \phi(\overline{\alpha}, p)$, then

$$\frac{p\delta}{1-\delta}H(\phi(\overline{\alpha},p)) \leq \frac{p\delta}{1-\delta}H(\underline{v}) < \underline{v} - \overline{\alpha} < \phi(\overline{\alpha},p) - \overline{\alpha},$$

which contradicts the fixed point equation (16). □

Lemma 23. *If $c = 0$, then for every $\epsilon > 0$, there exists a threshold $\mu_0(\epsilon) > 0$ on the supply-demand ratio such that whenever $\mu < \mu_0(\epsilon)$, every equilibrium outcome of independent lotteries has success probability $p < \epsilon$.*

Proof of Lemma 23. Define the function

$$h(p) = \overline{\pi}^{IL_p} = \int_{-\infty}^{\infty} \pi^{IL_p}(\alpha)\, dF(\alpha),$$

where π^{IL_p} is as in (22). By Proposition 3 and the Dominated Convergence Theorem (since $0 \leq \pi^{IL_p}(\alpha) \leq 1$), $h(p)$ is continuous in p. Moreover, by Lemma 22 in Appendix F.8, $h(p) \geq ph(1) > 0$

for every $p > 0$. For every $\epsilon > 0$, define

$$\mu_0(\epsilon) = \min_{p \in [\epsilon, 1]} h(p).$$

The minimum exists because $h(p)$ is continuous and $[\epsilon, 1]$ is compact. Moreover, $\mu_0(\epsilon) > 0$ since $h(p) > 0$ throughout this interval. This $\mu_0(\epsilon)$ satisfies the condition of Lemma 23 by construction.

$\square$