Alaskan Hunting License Lotteries Are Flexible and Approximately Efficient

Nick Arnosti
Columbia Business School

Tim Randolph
Columbia University

Abstract

We analyze the k-ticket lottery, which is used to allocate hunting permits in the state of Alaska. Each participant is given k tickets to distribute among lotteries for different types of items. Participants who win multiple items receive their favorite, and new winners are drawn from the lotteries with unclaimed items.

When supply is scarce, equilibrium outcomes of the k-ticket lottery approximate a competitive equilibrium from equal incomes (CEEI), which is Pareto efficient. When supply is moderate, k-ticket lotteries exhibit two sources of inefficiency. First, some agents may benefit from trading probability shares. Second, outcomes may be "wasteful": agents may receive nothing even if acceptable items remain unallocated. We bound both sources of inefficiency, and show that each is eliminated by a suitable choice of k: trades are never beneficial when $k = 1$, and waste is eliminated as $k \to \infty$.

The wastefulness of the k-ticket lottery has some benefits: agents with strong preferences may prefer k-ticket lottery outcomes to those of any nonwasteful envy-free mechanism. These agents prefer small values of k, while agents with weak preferences prefer large values of k. Together, these results suggest that the k-ticket lottery performs well under most circumstances, and may be suitable for other settings where items are rationed.

1 Introduction

We address the challenge of allocating heterogeneous resources to unit-demand agents without using money. Real-world examples include the allocation of seats at public schools, affordable housing units, spots in popular running races, and permits for popular hiking and camping destinations.

We take inspiration from the allocation of hunting permits. In the United States, this allocation is conducted by state agencies. Prospective hunters must apply for a permit, which are limited in supply. Typically, permitting fees are nominal, and demand exceeds supply. Most states conduct annual lotteries to award permits. Though procedures vary, states generally conduct separate allocations for each species, so that a hunter might apply for and win permits for multiple species. However, for each species, hunters are constrained to receive at most one permit.

In Alaska, the allocation for each individual species proceeds as follows. The state offers different types of permit, which specify when and where hunting may occur, as well as the number and gender of animals that may be killed. It sets a quota for each permit type, and publishes this information in an annual "Draw Supplement." [1] Each applicant can submit up to six applications, and may apply for a given type of permit multiple times to increase their odds of winning (for example, one might apply five times for a popular permit, and once for a second, less popular permit). Each applicant also submits a preference ranking over permits for which they have applied. After the application deadline, applications are drawn randomly for each permit. Applicants who win multiple permits keep only their favorite among these, and new names are drawn for the remainder. This process continues until each permit type has either been fully allocated or offered to all applicants.[2]

This inspires several natural questions. How do outcomes of this lottery compare to those that would result from using other procedures? Are these outcomes efficient, or at least approximately so? And what would be the effect of increasing or decreasing the maximum number of applications?

We address these questions using a model with a continuum of agents and n types of items. In Section 3 we define the *k-ticket lottery*, in which agents may allocate $k \in \mathbb{N}$ tickets across lotteries. Given agents' actions, each item has a resulting level of competition, summarized by the probability that a ticket entered into the corresponding lottery will be drawn. Agents take these win probabilities into account when deciding how to use their tickets.

Our first result establishes that when there are many more agents than items, equilibrium outcomes of the k-ticket lottery are close to outcomes of a competitive equilibrium with equal incomes (CEEI): each agent is nearly indifferent between these mechanisms. Thus, when demand is high, the choice of k is not very important, and k-ticket lotteries can be thought of as a virtual currency system. In particular, their equilibrium outcomes are approximately Pareto efficient.

[1]The most recent supplement is available at https://www.adfg.alaska.gov/index.cfm?adfg=huntlicense.drawsupplements, and additional information on the draw can be found at https://www.adfg.alaska.gov/index.cfm?adfg=huntlicense.lottery.

[2]An alternative description of the process is as follows. A hunter who submits t_i applications to permit type i is given a random priority for i that is the first-order statistic of t_i iid uniform draws from $[0, 1]$. Using these priorities and the preference rankings submitted by hunters, the state runs the permit-proposing Deferred Acceptance algorithm.

In general, k-ticket lotteries exhibit two forms of inefficiency. First, some agents may benefit from trading probability shares. Second, outcomes may be "wasteful": agents may receive nothing even if acceptable items remain unallocated. Our second contribution is to bound both sources of inefficiency. In particular, Theorem 2 states that for any equilibrium allocation, no ex ante exchange of probability shares can increase the welfare of every agent to more than $(2 - \frac{1}{k})$ times their welfare under the original allocation. Theorem 3 states that no reallocation of wasted items can increase the welfare of any agent to more than $(1 + \frac{1}{k})$ times their welfare under the original allocation. In particular, these two theorems imply that when $k = 1$, no beneficial trades are possible, and as $k \to \infty$, waste is eliminated.

Our third contribution is to demonstrate that the wastefulness of k-ticket lotteries may actually improve sorting: agents with strong preferences for popular items may prefer the k-ticket lottery to any nonwasteful mechanism. We illustrate this by considering the scenario in which there are two types of items and one item is universally preferred. In this scenario, CEEI is equivalent to random matching. (This occurs because the less demanded item is available at no cost, so all agents spend their entire budget on the highly demanded item.) By contrast, in a k-ticket lottery, securing the second item comes at the cost of spending a ticket, which only agents with a weak preference for the first item are willing to pay. Proposition 4 establishes that in this setting, smaller choices of k are better for agents who strongly prefer the more popular item. In addition, we demonstrate that no other mechanism achieves greater efficiency with the same amount of waste: Proposition 5 states that in the two-item setting, allocating items with a k-ticket lottery is equivalent to discarding the items wasted by the k-ticket lottery and then running CEEI.

Taken together, our results suggest that k-ticket lotteries produce good outcomes across a range of market conditions. When demand far outpaces supply, they approximate Competitive Equilibrium from Equal Incomes. When demand and supply are more balanced, they remain approximately efficient and may allow agents to signal their preference intensities more effectively than CEEI.

2 Related Work

Hylland and Zeckhauser (1979) introduce Competitive Equilibrium from Equal Incomes (CEEI), in which agents use virtual currency to buy probability shares of items. Equilibria of CEEI are ex-ante

Pareto efficient and envy-free, and Ashlagi and Shi (2015) show that this is the only mechanism with these properties. Despite its virtues, CEEI is rarely used in practice due to the challenges of soliciting preference intensities from agents.

Several papers have argued that other mechanisms may approximate CEEI outcomes. Abdulkadiroğlu et al. (2011) and Miralles (2009) note advantages of the "Boston" mechanism, in which items are preferentially awarded to agents who rank them highly. Abdulkadiroğlu et al. (2015) introduce Choice-Augmented Deferred Acceptance (CADA), which allows agents to improve their priority at one targeted item. If items have no inherent preferences over agents, the Boston mechanism is simply a sequence of 1-ticket lotteries, and CADA is a 1-ticket lottery followed by random serial dictatorship. Due to the close relationship between these mechanisms and the one-ticket lottery, our claim that the one-ticket lottery is trade efficient (Corollary 1) parallels the observation that schools that fill in the first round of Boston are efficiently allocated (Miralles, 2009), as well as an analogous claim for CADA (Abdulkadiroğlu et al., 2015).

Immorlica et al. (2017) introduce the "raffle", which is the natural extension of the k-ticket lottery with $k = \infty$. They show that for any equilibrium of this mechanism, it is impossible to simultaneously increase all participants' welfare to $\frac{e}{e-1}$ times their original welfare, and argue that this mechanism is a practical alternative to CEEI. However, it may be difficult to allow each agent to allocate an infinite number of tickets in practice. As such, our work can be viewed as a study of practical variants of the raffle with finitely many tickets.

One key difference between a k-ticket lottery and the ∞-ticket raffle is that the former might fail to allocate items, even if there are unmatched agents who consider them acceptable. In fact, this wastefulness also differentiates the k-ticket lottery from CEEI, the Boston mechanism, CADA, and Random Serial Dictatorship. Although wastefulness seems to be a shortcoming, Section 6 shows that it also has benefits: k-ticket lotteries may achieve more effective sorting than these other mechanisms, benefiting agents with strong preferences. The setting in Section 6 is very similar to that considered by Cavallo (2014). He seeks to maximize utilitarian welfare, and shows that in some cases, random allocation is optimal, but in others, wasteful mechanisms offer improvements.

At a technical level, our paper leverages several results from other papers. Lemma 1, which gives an agent's best response, builds upon the work of Chade and Smith (2006). Lemma 2, which states that every strategy profile induces a unique consistent outcome, builds upon the uniqueness of a stable outcome established by Azevedo and Leshno (2016). The proof that 1-ticket lotteries have a unique equilibrium is similar to a result in Gale (1976). The proof of Theorem 2, which states that k-ticket equilibria are $(2 - 1/k)$-approximately trade efficient, parallels the proof of approximate efficiency in Immorlica et al. (2017).

3 Model

There are n types of items with quantities $\mu = (\mu_1, \mu_2, ..., \mu_n) \in \mathbb{R}^n_+$. There is mass M of agents. Each agent is identified by a type $v = (v_1, v_2, ...v_n) \in \mathbb{R}^n$, where v_i denotes the agent's value for item i. Types are distributed according to a probability measure η that is absolutely continuous with respect to the Lebesgue measure. Agents also have an outside option with a value normalized to zero. Without loss of generality, we consider only agent types with positive value for at least one item and items which a positive measure of agents prefer to their outside option.

In Section 3.1, we define the optimization problem facing a single agent. Section 3.2 discusses how agent strategies determine aggregate outcomes and defines equilibrium. Section 3.3 introduces our efficiency metrics.

3.1 Single-Agent Decision Problem

The agent's interaction with the k-ticket lottery is summarized by a strategy $s = (t, \succ)$ consisting of a ticket allocation vector $t = (t_1, t_2, ..., t_n)$ and a preference ranking $\succ$ over the set of items. For $i, j \in [n]$, t_i denotes the number of tickets entered in lottery i and $j \succ i$ indicates that the agent prefers j to i in her preference ranking. Let S_k denote the set of all possible strategies with at most k tickets. A *strategy profile* $\Phi : \mathbb{R}^n \to S_k$ associates a k-ticket strategy with every agent type.

A strategy profile Φ induces a vector $p = (p_1, p_2, ..., p_n)$ of win probabilities, where p_i denotes the probability that an individual ticket entered in lottery i will be drawn during the lottery resolution process. Agents take p as given when considering which strategy to adopt. We address the calculation of p in Section 3.2.

Given strategy s and win probabilities p, the vector $\chi(s, p) = \{\chi_i(s, p)\}_{i=1}^n$ summarizes the probability that an agent who adopts strategy s wins each item. We define

$$\chi_i(s, p) := (1 - (1 - p_i)^{t_i}) \prod_{j \succ i} (1 - p_j)^{t_j}, \tag{1}$$

meaning that an agent wins item i so long as at least one the tickets entered in lottery i is drawn and the agent does not win any item j for which $j \succ i$.

For fixed p, the expected welfare of an agent with type v under strategy s is

$$w(v, s, p) := v \cdot \chi(s, p). \tag{2}$$

At any equilibrium strategy profile, each agent maximizes his or her expected welfare.

Definition 1. *Let $p = (p_1, p_2, ..., p_n)$ be a fixed vector of win probabilities for each lottery. A strategy profile Φ is* ***optimal*** *given p if*

$$\forall v, \Phi(v) \in \arg\max_{s \in S_k} w(v, s, p). \tag{3}$$

Given p, an agent's optimal strategy can be computed by a simple greedy algorithm. We write $s + \{i\}$ to denote the strategy s with an additional ticket placed in the lottery for item i.

Lemma 1. *For any p and any v, the following algorithm computes an optimal strategy:*

1. *Set $s := (\vec{0}, \succ)$, where $\succ$ is ordered according to the true preferences of v.*
2. *For $j = 1, 2, .., k$: select $i_j \in \arg\max_{i \in [n]}(w(v, s + \{i\}, p) - w(v, s, p))$. If the $\arg\max$ is not unique, select i_j to be the agent's most preferred item in this set. Set $s := s + \{i_j\}$.*

Proof Sketch. The problem of computing an optimal k-ticket strategy can be transformed into a downward-recursive portfolio choice problem over stochastic options. Lemma 1 then follows from the optimality of the Marginal Improvement Algorithm introduced by Chade and Smith (2006). The reduction to downward recursive portfolio choice is elaborated in Appendix B.1. □

Multiple choices of i may maximize $w(v, s + \{i\}, p)$ for two reasons. First, different strategies may lead to the same probabilistic allocation if the agent can secure his or her most preferred item with a single ticket. In this case, the algorithm above ensures that agents continue to enter tickets into the lottery for their most preferred item.[3] Second, an agent may be exactly indifferent between two strategies that lead to different allocations. Because η is absolutely continuous, for any p this occurs only for a set of agents with measure zero. Thus, the algorithm above uniquely determines allocation probabilities for almost every agent type.

[3]This choice is motivated by the fact that this strategy is uniquely optimal under a small perturbation of the vector p.

3.2 k-Ticket Equilibria

Given a strategy profile, any win probability vector p must be consistent with the available quantities μ. That is, for $i \in [n]$, either the entire mass μ_i should be allocated or every ticket entered into lottery i should be drawn.

Definition 2. *Fix a strategy profile Φ. A vector p of win probabilities is **consistent** with Φ if for all $i \in [n]$,*

$$M \sum_{s \in S_k} \chi_i(s, p)\, \eta(\{v : \Phi(v) = s\}) \leq \mu_i, \tag{4}$$

with equality if $p_i < 1$.

Lemma 2. *For each strategy profile, there is a unique consistent win probability vector.*

Proof Sketch. For a fixed strategy profile Φ, the k-ticket lottery is equivalent to a deferred acceptance procedure in which agents' priority for each item i increases stochastically with the number of tickets placed in the lottery for item i. It then follows from the results of Azevedo and Leshno (2016) that every strategy profile Φ admits at least one vector p of win probabilities, that the set of win probability vectors forms a lattice, and that every consistent vector p matches the same set of agents. Because the measure of matched agents is strictly increasing in p, the lattice property implies that the consistent p is unique. The full proof is in Appendix B.2. □

Definition 3. *A **k-ticket equilibrium** is a pair (Φ, p) consisting of a strategy profile and a vector of win probabilities such that Φ is optimal given p and p is consistent with Φ.*

Proposition 1 (k-Ticket Equilibria Exist)**.** *Every k-ticket lottery admits an equilibrium.*

Proof Sketch. We describe a continuous function f that maps the space of win probability vectors to itself such that $f(p) = p$ if and only if p corresponds to a k-ticket equilibrium (Φ, p). Applying Brouwer's fixed point theorem to f demonstrates that an equilibrium exists. The full proof is in Appendix B.2. □

In fact, the 1-ticket lottery has a unique vector p of win probabilities at equilibrium. When $k \geq 2$, however, k-ticket lotteries may have multiple equilibria with distinct win probability vectors. We prove the uniqueness of 1-ticket equilibria and provide an example illustrating the non-uniqueness of 2-ticket equilibria in Appendices B.3 and B.4.

3.3 Envy-Free and Pareto-Efficient Allocations

A (probabilistic) allocation is a function $x : \mathbb{R}^n \to [0,1]^n$ that assigns each agent a chance to receive each item. Each agent type v is allocated $x(v) = (x_1(v), x_2(v), ..., x_n(v))$, where $x_i(v)$ denotes the probability of recieving item i. Allocations satisfy the following feasibility constraints:

1. Allocations respect the quantity of each item:

$$\forall i, M \int x_i(v) d\eta \leq \mu_i. \tag{5}$$

2. Each agent receives at most one item:

$$\forall v, \sum_{i=1}^{n} x_i(v) \leq 1. \tag{6}$$

Definition 4. *When the inequality in* (5) *is tight, we say that item i is* **fully allocated** *under the allocation x.*

A k-ticket lottery equilibrium (Φ, p) corresponds to the allocation x such that $x(v) = \chi(\Phi(v), p)$.

Definition 5. *An allocation is* **envy-free** *if every agent weakly prefers her allocation to that of every other agent. That is, for every agent type v, we have*

$$x(v) \cdot v = \max_{u \in \mathbb{R}^n} x(u) \cdot v. \tag{7}$$

The property of envy-freeness is readily interpreted as a fairness constraint as it ensures that no agent envies the allocation of another. However, it also arises naturally if the mechanism designer cannot observe agent types, and must offer the same set of options to every agent. From this perspective, envy-freeness may be thought of as a feasibility constraint that enforces anonymity. In our continuum model, allocations produced by the k-ticket lottery are envy-free because every agent faces the same win probability vector p.

Definition 6. *Allocation x'* ***Pareto dominates*** *allocation x if $x'(v) \cdot v \geq x(v) \cdot v$ for every v, with strict inequality for some set of agents with positive η-measure. An allocation is* ***Pareto-efficient*** *if no feasible allocation Pareto dominates it.*

4 The High Demand Setting

We start by considering the case where the number of agents is much larger than the number of items. In this case, the win probability for each item is small. As a result, an agent's chance of receiving an item is roughly proportional to the number of tickets spent on that item. In other words, tickets serve as virtual currency, and agents spend their tickets on the single item that maximizes "bang for buck." The resulting outcome resembles a competitive equilibrium from equal incomes. We establish this formally by showing that as M grows, each agent becomes indifferent between the k-ticket lottery and CEEI.

Theorem 1. *Fix k, μ, and η. Let X_M denote the set of allocations corresponding to k-ticket equilibria with the given parameters and agent mass M, and let ce_M denote the unique allocation under CEEI.*[4] *For almost every agent v, we have*

$$\lim_{M\to\infty} \sup_{x\in X_M} \frac{x(v)\cdot v}{ce_M(v)\cdot v} = \lim_{M\to\infty} \inf_{x\in X_M} \frac{x(v)\cdot v}{ce_M(v)\cdot v} = 1. \tag{8}$$

Proof sketch. According to the greedy algorithm for optimal strategy choice (Lemma 1), agents place tickets in more than one lottery if the difference in marginal benefit between placing their first and last ticket in the same lottery is sufficient to justify the change. If the odds of winning each lottery are small, the marginal benefit of placing each subsequent ticket in the same lottery is roughly constant. As a result, most agents opt to place all their tickets in a single lottery and the resulting outcome is similar to the unique 1-ticket equilibrium. Finally, when all items are fully allocated, the 1-ticket equilibrium allocation is equivalent to CEEI. The full proof of Theorem 1 is in Appendix C. □

5 Approximate Efficiency

Section 4 establishes that when there are far more agents than items, any k-ticket equilibrium outcome approximates CEEI, which is Pareto efficient. In this section, we make no assumptions about the number of agents and items. Although k-ticket equilibria are not Pareto efficient, we

[4] In general, CEEI may have multiple equilibria. However, if there is an equilibrium in which all items are fully allocated, then this equilibrium is unique. Therefore, for all sufficiently large M, ce_M is well defined.

prove that they are approximately so.

We start by identifying two sources of inefficiency that may arise in a k-ticket equilibrium. First, two or more agents may be able to profitably exchange probability shares. Second, there might be unallocated items which can be assigned to benefit some agents. We formalize these concepts below.

Definition 7. *A* **trade reallocation** *of x is a feasible allocation x' such that $\int x(v)\ d\eta = \int x'(v)\ d\eta$. An allocation x is (ex ante)* **trade efficient** *if no trade reallocation of x Pareto dominates it.*

Definition 8. *A* **waste reallocation** *of x is a feasible allocation x' such that for every v and every i that is fully allocated under x, $x'_i(v) \leq x_i(v)$. An allocation x is* **nonwasteful** *if for any waste reallocation x' and all v, $x(v) \cdot v \geq x'(v) \cdot v$.*

This definition implies that an allocation is wasteful whenever there is some item that is not fully allocated and some agent with a positive probability of an outcome that is worse than receiving this good.

If agents outnumber items and always prefer something to nothing, nonwastefulness and trade-efficiency are equivalent to the familiar property of Pareto efficiency.

Proposition 2. *If $M \geq \sum_i \mu_i$ and all agents have positive values for all items, an allocation x is Pareto efficient if and only if it is nonwasteful and trade efficient.*

Proof. Pareto efficiency trivially implies trade efficiency and nonwastefulness. In this setting, any nonwasteful mechanism allocates all items. When all items are allocated, trade efficiency implies Pareto efficiency. □

The next sections provide definitions of approximate trade-efficiency and nonwastefulness, and show that k-ticket lottery outcomes satisfy these properties. In particular, the k-ticket lottery is always $(2 - \frac{1}{k})$-trade efficient and $\frac{1}{k}$-wasteful. Figure 1 summarizes our results.

5.1 Approximate Trade Efficiency

The k-ticket lottery is not generally trade efficient.[5] In particular, if non-identical agents both place tickets in lotteries i and j and face some risk of receiving nothing, then they can exchange

[5] That is, agents may be able to agree to an ex ante exchange of probability shares that offers a Pareto improvement. The k-ticket lottery *is* ex post trade efficient, meaning that after the allocation occurs, no agents can exchange items to their mutual benefit. For details, see Appendix A.

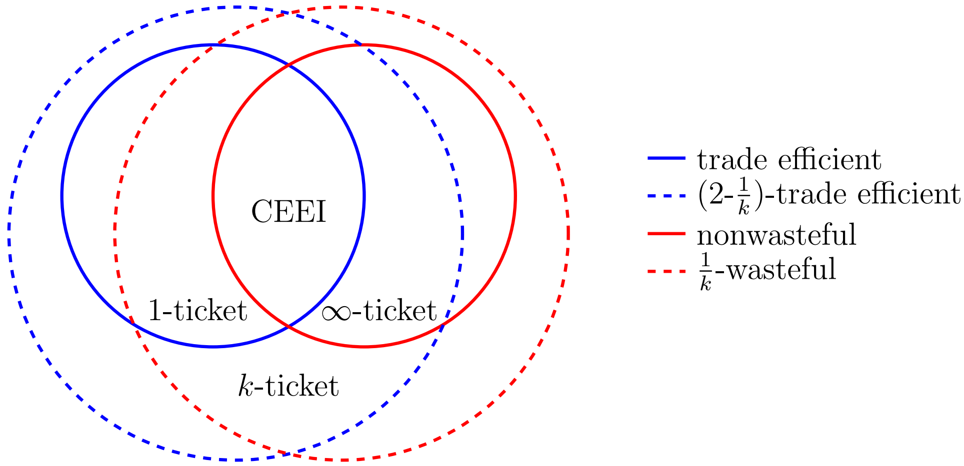

Figure 1: CEEI is both trade efficient and nonwasteful. Although k-ticket lotteries do not generally satisfy either of these conditions, they are approximately trade efficient and approximately nonwasteful. Furthermore, trade efficiency is attained for $k = 1$ and nonwastefulness as $k \to \infty$.

probability shares of i and j to their mutual benefit (Appendix D.1 makes this statement precise). We show, however, that the benefit from such trades is limited.

Definition 9. *Given $\alpha \geq 1$, an allocation x is α-**trade efficient** if there does not exist a trade reallocation x' of x such that for all v,*

$$\frac{x'(v) \cdot v}{x(v) \cdot v} \geq \alpha.$$

Theorem 2. *Any equilibrium allocation of the k-ticket lottery is $(2 - \frac{1}{k})$-trade efficient.*

We prove Theorem 2 in Appendix D.2. For $k = 1$ we obtain the following corollary.

Corollary 1. *The 1-ticket equilibrium allocation is trade-efficient.*

The bound in Theorem 2 is not tight for $k \geq 2$: in particular, the results of Immorlica et al. (2017) imply that the ∞-ticket lottery is $e/(e-1) \approx 1.58$ trade efficient. However, Theorem 2 implies that for any k, it is impossible to double every agent's welfare through an ex-ante exchange of probability shares. By contrast, random serial dictatorship (which always produces an ex post efficient allocation) is not α-trade efficient for any $\alpha < \infty$ (see Immorlica et al. (2017) for details).

5.2 Approximate Nonwastefulness

We quantify waste in terms of the value that can be gained from reallocating wasted items. In particular, we say that an allocation is ϵ-wasteful if reallocating surplus items never improves any agent's utility by more than an ϵ-fraction.

Definition 10. *An allocation x is ϵ-**wasteful** if for all waste reallocations x' and every v,*

$$\frac{x'(v) \cdot v}{x(v) \cdot v} \leq 1 + \epsilon. \tag{9}$$

Theorem 3. *Any equilibrium allocation of the k-ticket lottery is $1/k$-wasteful.*

Proof. Let x be an allocation corresponding to a k-ticket equilibrium and A_x be the set of fully allocated items under x. For a given v, the best possible waste reallocation ensures that v receives $i^* \in \arg\max_{i \in [n] \setminus A_x} v_i$ if she fails to win anything else, and thus

$$x'(v) \cdot v = x(v) \cdot v + (1 - \sum_i x_i(v)) v_{i^*}. \tag{10}$$

Consider the process of optimal strategy selection outlined in Lemma 1. At each stage, the agent chooses to place her j^{th} ticket in the lottery corresponding to the item which maximizes her marginal increase in welfare. At each stage, the agent declines to place a ticket in the lottery for item i^*, so her marginal increase in welfare must be at least $(1 - \sum_i x_i(v)) v_{i^*}$ at each step. It follows that her total welfare

$$x(v) \cdot v \geq k \cdot (1 - \sum_i x_i(v)) v_{i^*}. \tag{11}$$

The theorem follows from substituting (11) into (10). □

Theorem 3 is in fact tight.[6] It implies that the mechanism designer can reduce waste to an arbitrarily small fraction of welfare by choosing a large value of k. The next section explains why this may not lead to desirable outcomes.

6 Wastefulness: Bug, or Feature?

To motivate our analysis, consider the real circumstances faced by moose hunters in Alaska, in which the 6-ticket lottery is used to allocate hunting permits. Moose hunts occur across the state,

[6] To see this, consider a setting with two types of items, and let x be the allocation corresponding to a k-ticket equilibrium with $p = (\epsilon, 1)$. Consider agent type $v = (1/\epsilon, 1 - \epsilon)$, for some small $\epsilon > 0$. By Lemma 1, the optimal strategy is to place all k tickets in the lottery for item 1. As $\epsilon \to 0$, $x(v) \cdot v \to k$. If x' gives the agent item 2 whenever she fails to receive item 1, then $x'(v) \cdot v \to k + 1$ as $\epsilon \to 0$.

and hunt location is correlated with demand. Consider the example of moose hunts near the Nowitna River corridor, in central Alaska. Hunts along the river are generally more popular than hunts in the areas to the east and west of the river (the draw supplement warns hunters that the latter areas “are remote and access is limited”). In 2018, just 10% of permit applications for moose hunts along the river corridor were drawn during the lottery resolution process, while 96% of permit applications for moose hunts in the remote areas were drawn (ADFG (2019a)).

To model this situation, we consider a specific scenario that we refer to as the *two-item setting.* In this setting, there is a unit mass of agents and two items $\mu_1 + \mu_2 \geq 1$. Agents' values for item 1 are normalized to 1, and $v_2 > 0$ for all agents. Given $v \geq 0$, let $F(v)$ denote the fraction of agents with $v_2 \leq v$. We assume that $F(1) > \mu_1$, meaning that not all agents who prefer item 1 can receive it. By contrast, the abundance of supply ensures that agents who wish to claim item 2 can do so. In this sense, the two-item setting captures a situation where demand is low enough that only one type of item is fully allocated, in contrast to the high demand setting considered in Section 4.

In the case of the moose hunters, suppose that item 1 corresponds to accessible hunts, while item 2 corresponds to remote and inaccessible hunts. A hunter for whom the difficulty of traveling to a remote region is prohibitive corresponds to an agent with a small v_2, as a permit for a remote hunt has little value to her. On the other hand, a hunter with more time, money or energy might correspond to an agent with $v_2 \approx v_1$ or $v_2 > v_1$. One might hope to assign most of the accessible hunts to hunters who are unable to access the remote regions. However, Proposition 3 implies that if everyone prefers accessible hunts, then any envy-free, nonwasteful mechanism awards these hunts randomly.

Proposition 3. *In the two-item setting, if all agents prefer item 1 to item 2 ($F(1) = 1$), the only nonwasteful, envy-free allocation randomly allocates item 1 and allocates item 2 to the remaining agents.*

Proof. Because all agents have positive values for both items and $\mu_1 + \mu_2 \geq 1$, nonwastefulness implies $x_1(v) + x_2(v) = 1$ for all v and that item 1 is fully allocated. Envy-freeness thus implies $x_1(v) = x_1(v')$ for every agent $v, v' < 1$: that is, all agents with $v < 1$ receive the same allocation. If all agents prefer item 1, the supply constraints imply that $x(v) = (\mu_1, 1 - \mu_1)$ for all $v < 1$. □

Proposition 3 implies that Competitive Equilibrium from Equal Incomes (Hylland and Zeckhauser, 1979), the Boston mechanism, Choice-Augmented Deferred Acceptance (Abdulkadiroğlu

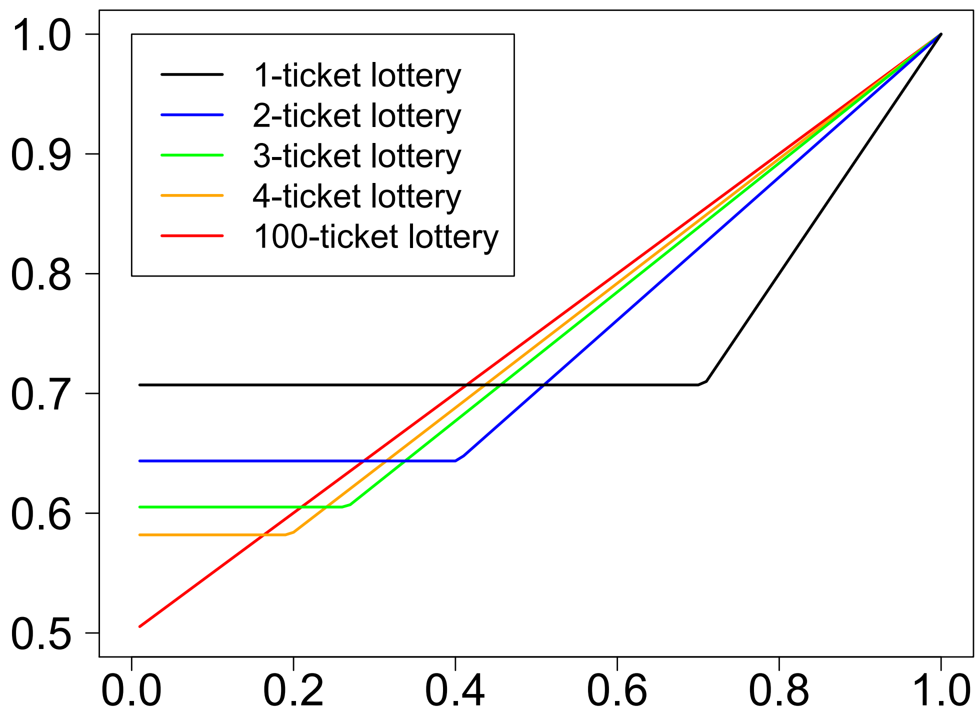

Figure 2: Agent utility (y-axis) as a function of v_2 (x-axis), in a setting with $\mu_1 = \mu_2 = 1/2$, $v_1 = 1$ for all agents, and $v_2 \sim U[0,1]$. In this case, CEEI results in a random allocation, so the utility of an agent of type v_2 is $(1 + v_2)/2$. The 100-ticket lottery results in a nearly random allocation. Decreasing the number of tickets increases welfare for agents with low v_2, at the expense of agents with intermediate or high v_2.

et al., 2015), the ∞-Ticket Lottery (Immorlica et al., 2017), and Random Serial Dictatorship are all equivalent to random allocation in the two-item setting. This holds despite the fact that agents have different relative values for the items.

In a 1-ticket lottery, by contrast, any hunter who competes for an accessible hunt forgoes the possibility of hunting elsewhere. This risk causes hunters with mild preferences for convenient hunts to select hunts in less accessible areas that they are more likely to win. As a result, there are fewer competitors for accessible hunts, and these hunts are allocated to hunters with stronger preferences for accessibility. The same effect is present for larger k, though to a lesser extent: hunters with strong preferences for accessible hunts will apply exclusively to them, while others will hedge their bets by applying to at least one remote hunt.

Figure 2 shows outcomes of the k-ticket lottery when v_2 is uniformly distributed on $[0, 1]$. Agents with strong preferences for item 1 spend all k tickets on it, while those with weaker preferences spend one ticket securing item 2. When k is large, most agents use a ticket on item 2 and the allocation is nearly random. As k shrinks, outcomes improve for agents with low value for item 2. Proposition 4 establishes that this holds more generally: agents with strong preferences for the more highly demanded item prefer smaller values of k.

Proposition 4 (Agents with Strong Preferences Prefer Smaller k)**.** *In the two-item setting, let x^i and x^j denote allocations corresponding to i-ticket and j-ticket equilibria, with $i < j$. If an agent v prefers x^i to x^j, then so do all agents w with $w_2 \leq v_2$.*

The proof of Proposition 4 is in Appendix E.1. Proposition 4 implies that hunters with little value for inaccessible hunts may prefer the k-ticket lottery to a nonwasteful mechanism. However, might there be a better choice of wasteful mechanism? In particular, is it possible to design a mechanism with comparable levels of waste that all hunters prefer? Proposition 5 states that in the two-item setting, the answer is no: the k-ticket lottery is equivalent to running CEEI after discarding some items, and is therefore trade efficient.

Proposition 5. *In the two-item setting, each k-ticket equilibrium allocation x corresponds to a CEEI with quantity vector $\tilde{\mu} := \int x(v)\, dF(v) \leq \mu$.*

The proof of Proposition 5 is in Appendix E.2. Proposition 5 implies that a version of CEEI with preemptive disposal could replicate outcomes from the k-ticket lottery. However, the former approach would face several obstacles to implementation, including the challenges of determining up front how many permits to withhold and asking hunters to report cardinal utilities.

7 Conclusion

Our results have implications for the distribution of hunting permits in Alaska as well as for other settings where items are rationed by lottery. Different results provide lessons for different markets.

Permits for bison hunts in Alaska are in very high demand. In 2018, in each lottery for bison hunt permits, 1% or less of the submitted applications were drawn during the lottery resolution process (ADFG, 2019a). In this case, Theorem 1 predicts that the k-ticket lottery outcome will approximate the outcome of a CEEI, which is envy-free and Pareto efficient. Hunters, meanwhile, have a simple strategy: they should apply exclusively to the hunt which maximizes the ratio of value to success probability. As a consequence, changing the maximum number of applications will have little effect on the outcome. The application process could therefore be simplified by adopting the 1-ticket lottery in cases where hunts are consistently oversubscribed by large margins.

In the submarket for unguided Kodiak bear hunting permits, win probabilities are more variable: depending on the hunt, individual applications were drawn between 1% to 50% of the time in 2018, with a median of 3.5% (ADFG (2019a)). In this case, the choice to set $k = 6$ results in a lottery mechanism that is 1/6-wasteful by Theorem 3, so the value of wasted permits will be low. Theorem 2 implies that the lottery is also 11/6-trade efficient, although we believe that this bound is pessimistic.

Theorem 3 suggests that the Alaska Department of Fish and Game could decrease the number of unallocated permits by increasing k. However, Section 6 outlines the potential downside of this change. Proposition 3 implies that a wasteful allocation is necessary to achieve effective sorting. When $k = 6$, agents who apply to less desirable hunts must forfeit opportunities to win more popular options, thereby increasing the odds for those aiming exclusively for these hunts. Proposition 4 suggests that reducing the number of tickets would benefit agents who are only interested in highly demanded hunts, while potentially increasing waste. In other words, maximizing the number of permits that are allocated is at odds with ensuring that permits are assigned to the hunters who value them the most.

This observation has implications for the handling of unallocated permits. When permits go unallocated, the biologist in charge of the area may reissue these permits later in the year (ADFG (2019b)). Although one might think that reissuing these permits unambiguously helps hunters, in fact the effect is more ambiguous. If leftover permits are reallocated, hunters who consider these options acceptable can use all six applications to compete for popular hunts. By committing not to reallocate leftover permits, the state could incentivize these hunters to apply to less popular hunts, improving outcomes for hunters who are only interested in popular hunts.

Our results find further application in other settings where the k-ticket lottery is employed. For instance, the Singapore Housing and Development Board uses the 1-ticket lottery to distribute leases in new public housing units (Singapore HDB (2019)). Applicants are divided into submarkets based on age and family type. Within a submarket, housing options are differentiated by maturity: leases in mature neighborhoods with established amenities typically experience higher demand than leases in non-mature neighborhoods. As a consequence, the results in Section 6 indicate that the Singapore housing lottery produces good outcomes for agents with strong preferences for mature neighborhoods but might result in some units going unallocated. Additionally, Corollary 1 implies that the outcome is trade-efficient.

Our analysis assumes equilibrium play. Agents may make mistakes if they are poorly informed about the demand for each item. To address this concern, the Singapore Housing and Development Board provides live estimates of success probabilities for each development, and allows applicants to modify their application at any time before the deadline (Singapore HDB (2018)). The annual Alaska hunt supplement lists the number of permits available for each hunt, as well as data from the previous year on the number of applicants and the percentage of applications drawn (see ADFG (2019a)). Though better than nothing, this approach suffers from at least two problems. First, hunters are not provided with information about year-to-year fluctuations in demand caused by the creation of new hunts, changes in the number of permits offered per hunt, and other factors. Second, for $k > 1$, it is impossible to compute the win probability of a single application from the number of successful permits and submitted applications, because some hunters may be offered multiple permits. Creating an on online system to display live win probabilities for each hunt could remedy these issues, helping to ensure that the appealing theoretical performance of the k-ticket lottery is borne out in practice.

References

Abdulkadiroğlu, A., Che, Y.-K., and Yasuda, Y. (2011). Resolving conflicting preferences in school choice: The "boston mechanism" reconsidered. *American Economic Review*, 101(1):399–410.

Abdulkadiroğlu, A., Che, Y.-K., and Yasuda, Y. (2015). Expanding "choice" in school choice. *American Economic Journal: Microeconomics*, 7(1):1–42.

ADFG (2019a). 2019-2020 Alaska drawing permit hunt supplement. Available at `https://www.adfg.alaska.gov/static/applications/web/nocache/license/huntlicense/pdfs/2019-2020_draw_supplement.pdfDDB3E5422A81D382A439F6878368477C/2019-2020_draw_supplement.pdf`.

ADFG (2019b). Undersubscribed hunt overview. Available at `https://www.adfg.alaska.gov/index.cfm?adfg=huntlicense.undersubscribed`.

Ashlagi, I. and Shi, P. (2015). Optimal allocation without money: An engineering approach. *Management Science*, 62(4):1078–1097.

Azevedo, E. M. and Leshno, J. D. (2016). A supply and demand framework for two-sided matching markets. *Journal of Political Economy*, 124(5):1235–1268.

Cavallo, R. (2014). Incentive compatible two-tiered resource allocation without money. In *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems*, pages 1313–1320. International Foundation for Autonomous Agents and Multiagent Systems.

Chade, H. and Smith, L. (2006). Simultaneous search. *Econometrica*, 74(5):1293–1307.

Gale, D. (1976). The linear exchange model. *Journal of Mathematical Economics*, 3(2):205–209.

Hylland, A. and Zeckhauser, R. (1979). The efficient allocation of individuals to positions. *Journal of Political Economy*, 87(2):293–314.

Immorlica, N., Lucier, B., Weyl, G., and Mollner, J. (2017). Approximate efficiency in matching markets. In *International Conference on Web and Internet Economics*, pages 252–265. Springer.

Miralles, A. (2009). School choice: The case for the Boston mechanism. In *International Conference on Auctions, Market Mechanisms and Their Applications*, pages 58–60. Springer.

Singapore HDB (2018). Hdb sales launch. Available at `https://www20.hdb.gov.sg/`.

Singapore HDB (2019). Modes of sale. Available at `https://www.hdb.gov.sg/cs/infoweb/residential/buying-a-flat/new/modes-of-sale`.

A Ex Post Trade Efficiency

We say that an allocation is ex post trade efficient if agents cannot mutually benefit from ex post trades of their received items. Formally, we define the *envy graph* of an allocation x to be a directed graph on $[n]$ containing a directed edge $i \to j$ for every i, j such that $x_i(v) > 0$ and $v_i < v_j$ for every v in a set $V_{i \to j}$ with $\eta(V_{i,j}) > 0$. The edge $i \to j$ in the envy graph indicates that a set of agents with nonzero measure will receive item i and would benefit from an ex post trade for item j.

Definition 11. *An allocation x is* ***ex post trade efficient*** *if its envy graph is acyclic.*

Lemma 3. *Ex ante trade efficiency implies ex post trade efficiency.*

Proof. We prove the contrapositive. Consider an allocation x whose envy graph contains a cycle C. Define

$$capacity(C) := \min_{i \to j \,\in\, C} \int_{V_{i \to j}} x_i(v)\, d\eta, \tag{12}$$

and note that $capacity(C) > 0$ by definition. For each edge $i \to j$ in C, choose a subset $V'_{i \to j} \subseteq V_{i \to j}$ such that $\int_{V'_{i \to j}} x_i(v)\, d\eta = capacity(C)$ and exchange i for j to create an allocation that Pareto dominates x and keeps $\int x(v)\, d\eta$ constant. □

Theorem 4. *Any k-ticket lottery outcome is ex post trade efficient.*

Proof. Consider the envy graph of an allocation x corresponding to a k-ticket equilibrium. If the envy graph contains $i \to j$, then there exists a set $V_{i \to j}$ with $\eta(V_{i \to j}) > 0$ such that for every $v \in V_{i \to j}$, $x_i(v) > 0$ and $v_i < v_j$. This implies $p_i > p_j$, as no agent in $V_{i \to j}$ places a ticket in the lottery for item i otherwise. Thus a cycle in the envy graph would correspond to a cycle of strict inequalities in win probabilities, which is impossible. □

B Proofs from Section 3

B.1 Transformation of k-Ticket Strategy Choice into Downward Recursive Portfolio Choice

Proof of Lemma 1. Fix a vector p of win probabilities, and define a multiset S of items that consists of k copies of each item i. For each item i, we define the cardinal utility of i to be $u_i := v_i$, and the success chance to be $\alpha_i := p_i$. Let $\succ_v$ be a total ordering on S according to the true preferences of an agent type v, with ties between identical items broken in a consistent manner.

The Marginal Improvement Algorithm (MIA) of Chade and Smith (2006) determines a portfolio $\hat{Y} \subseteq S$ that maximizes the gross payoff

$$f(Y) = \sum_{i \in Y} \Big(\alpha_i u_i \prod_{j \in Y \mid j \succ_v i} (1 - \alpha_j) \Big) = \sum_{i \in Y} \Big(p_i v_i \prod_{j \in Y \mid j \succ_v i} (1 - p_j) \Big) \tag{13}$$

for an agent type v. Let $\succ$ be a preference ordering of $[n]$ consistent with $\succ_v$, and let $t = (t_1, t_2, ..., t_n)$ be the ticket choice vector constructed by setting each t_i equal to the number of copies of item i in $\hat{Y}$. Expanding (13), we have

$$f(\hat{Y}) = \sum_{i \in [n]} \Big(v_i (1 - (1 - p_i)^{t_i}) \prod_{j \succ i} (1 - p_j)^{t_j} \Big) = v \cdot \chi((t, \succ), p) = w(v, (t, \succ), p). \tag{14}$$

Thus, given v, the MIA determines a strategy $(t, \succ)$ that maximizes $w(v, (t, \succ), p)$.

Let s_Y denote the strategy consisting of the ticket choice vector corresponding to $Y \subseteq S$ and a preference order $\succ$ reflecting the true preferences of a fixed agent type v. On our portfolio, the MIA proceeds as follows:

1. Begin with an empty portfolio $Y_0 = \emptyset$.

2. Choose any item
$$i \in \arg \max_{i \in S \setminus Y_{j-1}} w(v, s_{Y_{j-1} + \{i\}}, p). \tag{15}$$

3. If
$$w(v, s_{Y_{j-1} + \{i\}}, p) - w(v, s_{Y_{j-1}}, p) > c(n), \tag{16}$$
set $Y_j := Y_{j-1} + \{i\}$ and go to Step 2.

Our algorithm is equivalent to the MIA with the additional requirement that the process of strategy augmentation stop after k steps. This is ensured by defining a cost function c such that $c(|Y|) = 0$ if $Y \leq k$ and $c(|Y|) = \infty$ if $|Y| > k$ (referred to in Chade and Smith (2006) as the *fixed sample size* k case). The optimality of the greedy algorithm follows from the optimality of the resulting instance of downward recursive portfolio choice. $\square$

B.2 Existence of Equilibria

In this section, we complete the proofs that every strategy profile Φ implies a unique consistent win probability vector p (Lemma 2) and that every k-ticket lottery has an equilibrium (Proposition 1).

Proof of Lemma 2. Given a strategy profile Φ, our setting maps directly to the two-sided matching market defined by Azevedo and Leshno (2016). In their framework, agent types are identified by a preference ranking of items $\tilde{\succ}$ and a score e_i that determines priority for each acceptable item. Each strategy $s = (t, \succ)$ in our model corresponds to a measure over types $(\tilde{\succ}, e)$ as follows:

1. $\tilde{\succ}$ ranks items with $t_i > 0$ according to $\succ$.

2. items with $t_i = 0$ are unacceptable according to $\tilde{\succ}$.

3. e_i is distributed as the first order statistic of t_i uniform random draws from $[0, 1]$. Scores are independent across items.

Allocation in Azevedo and Leshno (2016) is summarized by a vector of cutoffs that lists the minimum score required to receive each item i. Under this mapping, the vector c of market clearing cutoffs is precisely $\vec{1} - p$, and the consistency condition is equivalent to the market clearing condition (Definition 2 of Azevedo and Leshno (2016).)

Corollary A1 in the appendix of Azevedo and Leshno (2016) implies that we can find a set of market clearing cutoffs c (equivalently, a vector of win probabilities p). Theorem A1 implies that the set of market clearing cutoffs c forms a lattice with the supremum and infimum operators in $[0, 1]^n$. By Theorem A2, the measure of agents matched to each item i is the same under any market clearing c (consistent p).

For each item i, the distribution of scores e_i has full support for each strategy s that places at least one ticket in the lottery for item i. As a result, any isolated decrease in a win probability for an item that a positive measure of agents receive results in strictly fewer agents matching. Thus the minimal and maximal elements of the lattice of consistent c coincide, implying that it contains a unique element. □

Proof of Proposition 1. Fix k, a set of items, an agent mass, and an agent type measure. We define a function $f : [0, 1]^n \to [0, 1]^n$ as follows. For any win probability vector p, denote by Φ^p the strategy profile determined by Lemma 1 that is unique up to a set of agents with measure zero.

We define $f(p)$ to be the vector of win probabilities determined by Φ^p, which is unique by Lemma 2.

By inspection, a fixed point of f corresponds to a k-ticket equilibrium. We show that f is continuous, from which the existence of a fixed point follows from Brouwer's theorem.

For $s \in S_k$, we have

$$(\Phi^p)^{-1}(s) = \{v : \forall s' \in S_k \setminus \{s\},\ v \cdot \chi(s,p) > v \cdot \chi(s',p)\}. \tag{17}$$

Because $\chi(s,p)$ is constant in v, $(\Phi^p)^{-1}(s)$ is an open convex polytope in $\mathbb{R}^n$ (or the empty set).

Let (j) denote the item on which a given agent v uses her j^{th} ticket according to the algorithm specified in Lemma 1. Lemma 1 implies that the win probability of the items on which v bids is weakly decreasing. Using this fact to unwind the algorithm, we have that $\Phi^p(v)$ is determined by the unique sequence of items $(1),(2),...,(k)$ such that $\forall i \in [n], j \in [k]$,

$$\sum_{l=1}^{j} \big(v_{(l)}p_{(l)} \prod_{m=l+1}^{j} (1-p_{(m)})\big) \geq v_i p_i + (1-p_i)\Big(\sum_{l=1}^{j-1} \big(v_{(l)}p_{(l)} \prod_{m=l+1}^{j-1} (1-p_{(m)})\big)\Big). \tag{18}$$

Because each constraint is polynomial in p, and because η is continuous, $\eta((\Phi^p)^{-1}(s))$ is continuous in p for all s.

Let $X(f(p))$ be the matrix in which each column is the probabilistic allocation for a strategy under $f(p)$, that is,

$$X_{(i,s)}(f(p)) := \chi_i(s, f(p)), \tag{19}$$

for all $i \in [n]$, $s \in S_k$. Let $H(p)$ be the vector containing the measure of agents who adopt each strategy under Φ^p, that is,

$$H_s(p) := \eta((\Phi^p)^{-1}(s)), \tag{20}$$

for all $s \in S_k$. If in fact $f(p)_i < 1$ for all $i \in [n]$, $f(p)$ is the unique vector which satisfies

$$X(f(p)) \cdot H(p) = \mu, \tag{21}$$

as this equation is equivalent to the consistency constraints. However, the consistency constraints allow the possibility that some items go unallocated if $f(p)_i = 1$. We generalize (21) by defining the map A as follows.

$$A_i(f(p),p) := \Big(\sigma\big(X_i(f(p))H(p)-\mu_i\big)+(1-f_i(p))\Big)\Big(X_i(f(p))H(p)-\mu_i\Big), \tag{22}$$

where $\sigma(y)=\max(0,y)$. By inspection, $A(f(p),p)=0$ if and only if $f(p)$ is consistent with Φ^p.

$A(f(p),p)$ is continuous in $f(p)$ and p, and by Lemma 2, p determines a unique $f(p)\in[0,1]^n$ such that $A(f(p),p)=0$. Thus $f(p)$ is continuous in p, and an equilibrium point exists by Brouwer's theorem. □

B.3 1-Ticket Lotteries Have Unique Equilibria

The following lemma demonstrates that 1-ticket lotteries have a unique win probability vector p. Because p is unique, every agent receives the same expected value in every equilibrium. Note that the result is very similar to a result of Gale (1976), who considered the case where agents may receive more than one item.

Lemma 4 (Uniqueness of 1-ticket Equilibria)**.** *if* (Φ,p) *and* (Φ',p') *are 1-ticket equilibria, then* $p=p'$.

Proof. Let (Φ,p) and (Φ',p') be 1-ticket equilibria. For contradiction, assume $p\neq p'$. Without loss of generality, for some item i, we have $p_i>p'_i$. Because every agent has a single ticket, we can write a closed-form expression for p_i. Let $t(v)$ denote the ticket choice vector of agent type v according to Φ and $t'(v)$ denote the ticket choice vector of agent type v according to Φ'. For each item i, we have

$$p_i=\min\{1,\frac{\mu_i}{\eta(\{v:t_i(v)=1\})}\}, \tag{23}$$

and the analogous statement holds for $t'(v)$.

Thus $p_i>p'_i$ implies the existence of a set of agents V with $\eta(V)>0$ such that for every $v\in V$, $t_i(v)=0$ and $t'_i(v)=1$. Let j be the item that maximizes $\eta(\{v\in V:t_j(v)=1\})$. Because agents in the set $\{v\in V:t_j(v)=1\}$ switched from j to i, it must be the case that for all $v\in V$,

$$p_jv_j>p_iv_i>p'_iv_i>p'_jv_j. \tag{24}$$

Rearranging, we have

$$\frac{p_j}{p'_j} > \frac{p_i}{p'_i}. \tag{25}$$

This corresponds to the intuition that if item i is relatively fully allocated in (Φ', p') compared to (Φ, p), some other item j must be more relatively fully allocated in (Φ', p') to create additional demand for item i. Repeating this argument implies the existence of a sequence of items, each more relatively fully allocated than the last. Because the number of items is finite, this presents a contradiction, and thus $p = p'$. □

B.4 k-Ticket Lotteries Are Not Unique for $k > 1$

Although every 1-ticket lottery has a unique equilibrium, k-ticket lotteries do not necessarily have unique equilibria when k is greater than 1. The following example demonstrates a 2-ticket lottery with multiple equilibria.

Consider a lottery with two items and $\mu_1 = \mu_2 = 1$. For simplicity, we consider a setting with two discrete groups of agents.[7] The first group of agents has type $(1, 0)$. Regardless of other agents' behaviors, these agents will always place both of their tickets in the lottery for item 1. The mass of this group is $1/(2\epsilon - \epsilon^2)$, which ensures that if only these agents apply for item 1, $p_1 = \epsilon$. We set $\epsilon > 0$ to be small enough that the win probability p_1 is (multiplicatively) close to ϵ regardless of the behavior of other agents. The second group of agents has type $(\frac{7}{32\epsilon}, 1)$ and mass $4/3$. This scenario is summarized in Table 1.

Agent Group	Mass	Type
Group 1	$1/(2\epsilon - \epsilon^2)$	$(1, 0)$
Group 2	$4/3$	$(\frac{7}{32\epsilon}, 1)$

Table 1: A 2-ticket lottery that yields 2 equilibria.

In the first equilibrium, group 1 agents enter both tickets in the lottery for item 1 and group 2 agents put both tickets in the lottery for item 2. The resulting vector of win probabilities is $p = (\epsilon, 1/2)$. In this equilibrium, a group 2 agent receives expected values of $14/32 - 7\epsilon/32$ from putting both tickets in the lottery for item 1, $24/32$ from putting both tickets in the lottery for item 2, and $23/32 - \epsilon/2$ from splitting her tickets, so the strategy profile is optimal given p.

In the second equilibrium, group 1 agents enter both tickets in the lottery for item 1 and group 2 agents split their tickets between the two lotteries. The resulting vector of win probabilities

[7] This assumption can be relaxed by replacing the agent point masses with Gaussians distributed tightly around the original points.

is $p = (\epsilon', \frac{3}{4-4\epsilon'})$, for some $\epsilon' < \epsilon$ with $\epsilon' \approx \epsilon$. In this equilibrium, a group 2 agent receives expected values of $\frac{\epsilon'}{\epsilon}(14/32 - 7\epsilon'/32) \approx 14/32$ from putting both tickets in the lottery for item 1, $\frac{24(1-\epsilon')-9}{16(1-\epsilon')^2} \approx 30/32$ from putting both tickets in the lottery for item 2, and $\frac{\epsilon'}{\epsilon}(7/32) + 3/4 \approx 31/32$ from splitting her tickets, so the strategy profile is optimal given p.

C Proof of Theorem 1

In this section, we prove that for almost every agent, the ratio of their expected value at any k-ticket equilibrium to their expected value under CEEI goes to 1 as the mass of agents increases. The proof treats μ, η, and k as fixed. When x is an allocation corresponding to a specified k-ticket equilibrium, we write p^x to denote the vector of win probabilities at this equilibrium.

We begin by proving two properties of equilibria which hold for large agent masses. First, for sufficiently large M, we prove that the maximum and minimum win probabilities for each item are $\Theta(1/M)$ at equilibrium. Second, we show that as demand for each item increases, the η-measure of the set of agents who bid on more than one item goes to 0. As a result, as M increases, the ratio of expected value at any k-ticket equilibrium to expected value at the unique 1-ticket equilibrium goes to 1 for almost every agent. Finally, we observe that the 1-ticket lottery produces the same allocation as CEEI when all items are fully allocated.

Claim 1 (Equilibrium win probabilities are $\Theta(1/M)$)**.** *Let μ, η, and k be fixed as above. There exist positive constants c_1, c_2 such that for every $M \geq max_j \mu_j$, every $x \in X_M$, and every $i \in [n]$,*

$$\frac{c_1}{M} \leq p_i^x \leq \frac{c_2}{M}. \tag{26}$$

Proof. Fix $x \in X_M$. We have for each item i that

$$\frac{\mu_i}{kM} \leq p_i^x. \tag{27}$$

This follows immediately from the fact that at most kM tickets can possibly be entered into a single lottery, and at least μ_i of these are drawn if $p_i^x < 1$.

We prove the upper bound by inductively constructing a set of items with win probabilities bounded by constant multiples of $1/M$. First, observe that in any equilibrium, there exists an item

i_1 such that $p^x_{i_1} \leq \frac{\sum_i \mu_i}{M}$. Otherwise, the quantity of items allocated by x would exceed $\sum_i \mu_i$, violating the equilibrium requirements. Furthermore, by assumption, some agents have nonzero value for a second item j. Thus there exist constants $1 > \epsilon_1, \gamma_1 > 0$ such that

$$\eta(\{v : \exists j \neq i_1, v_j > \epsilon_1 v_{i_1}\}) > \gamma_1. \tag{28}$$

As a result, there exists an item i_2 such that $p^x_{i_2} \leq \frac{\sum_i \mu_i}{\epsilon_1 \gamma_1 M}$. Otherwise, by Lemma 1, a mass $\gamma_1 M$ of agents would place their first ticket on items other than i_1. Observe that this ensures each agent is allocated with probability at least $\frac{\sum_i \mu_i}{\epsilon_1 \gamma_1 M}$ in the 1-ticket case, and the total allocation probability of each agent is strictly increasing in the number of tickets when p^x is held constant. This presents a contradiction, as the quantity of items allocated to the set of agents with mass $\gamma_i M$ would exceed $\sum_i \mu_i$. Repeating this argument for $m = 2, 3, ..., n$, we construct a series of constants $1 > \epsilon_m, \gamma_m > 0$ such that

$$\eta(\{v : \exists j \neq i_1, ..., i_{m-1}, v_j > \epsilon_m \max\{v_{i_l}\}_{l=1}^{m-1}) > \gamma_m, \tag{29}$$

and a series of items i_m such that

$$p^x_{i_m} \leq \frac{\sum_{i=1}^n \mu_i}{M \prod_{j=1}^{m-1} \epsilon_j \gamma_j}. \tag{30}$$

□

Claim 2 (The set of ticket splitters goes to 0)**.** *For $x \in X_M$, let S_x be the set of* ticket splitters*, agents whose optimal strategy under Lemma 1 bids on more than one item. We have*

$$\lim_{M \to \infty} \sup_{x \in X_M} \eta(S_x) = 0. \tag{31}$$

Proof. By definition, every agent in S_x adopts an equilibrium strategy in which she bids on at least two items. Let v be an agent in S_x who places her first $t-1$ tickets in the lottery for item i and her t^{th} ticket in the lottery for item j. By Lemma 1, we have

$$p_i^x v_i > p_j^x v_j. \tag{32}$$

$$p_j^x v_j + (1 - p_j^x)(1 - (1 - p_i^x)^{t-1}) v_i > p_i^x v_i + (1 - p_i^x)(1 - (1 - p_i^x)^{t-1}) v_i. \tag{33}$$

Rearranging yields

$$\frac{p_i^x}{p_j^x} > \frac{v_j}{v_i} > \frac{p_i^x}{p_j^x} + (1 - \frac{p_i^x}{p_j^x})(1 - (1 - p_i^x)^{t-1}), \tag{34}$$

and making use of the fact that $(1 - (1 - p_i^x)^{t-1}) \leq k p_i^x$ for $t - 1 < k$, we have

$$S_x \subseteq \{v : \frac{v_i}{v_j} \in [\frac{p_i^x}{p_j^x}(1 - k(p_i^x - p_j^x)), \frac{p_i^x}{p_j^x}]\}. \tag{35}$$

Claim 1 implies p_i^x / p_j^x is upper-bounded by a constant, so Claim 2 follows from the absolute continuity of η as the interval shrinks to zero. □

Lemma 5 (k-ticket utility converges to 1-ticket utility)**.** *Let x_M^1 denote the allocation corresponding to the unique 1-ticket equilibrium with agent mass M. For almost every agent v, we have*

$$\lim_{M \to \infty} \sup_{x \in X_M} \frac{x(v) \cdot v}{x_M^1(v) \cdot v} = \lim_{M \to \infty} \inf_{x \in X_M} \frac{x(v) \cdot v}{x_M^1(v) \cdot v} = 1. \tag{36}$$

Proof. To prove Lemma 5, we must understand the structure of the 1-ticket lottery. Consider the outcome of a 1-ticket lottery in a market with parameters μ, η, and M. By Lemma 4, this 1-ticket lottery has a unique equilibrium win probability vector p which in turn determines a unique expected value for any agent that plays an optimal strategy. Suppressing the dependence on η, we write $w^1(\mu, M, v)$ to denote the expected value of an agent v at 1-ticket equilibrium. We proceed to show that w^1 is continuous in μ and M.

Let $V_i(p)$ denote the measure of the set of agents who bid on good i under the optimal strategy given by Lemma 1 when the win probability vector is p. In the 1-ticket lottery, for all i, we have

$$\mu_i = M p_i V_i(p) \tag{37}$$

when M is sufficiently large. By Lemma 4, each agent v bids on a good $i \in \arg\max_{i \in [n]} v_i p_i$

at equilibrium. $V_i(q)$ is thus continuous in p, as it is the measure of a polytope whose sides are defined by linear equations in p. As a result, μ is continuous in p in the 1-ticket lottery. We write $\mu^1(p) := \mu^1(p, M, \eta)$ to indicate this.

By Lemma 4, μ determines a unique value of p satisfying the system of equations defined by (37). Combined with the fact that μ is continuous in p, this implies that p is a continuous function of μ. For fixed μ and sufficiently large M, scaling M by a constant c means that (37) is satisfied by the win probability vector p/c,[8] and thus p is a continuous function of M as well. Finally, the expected value of any agent v at 1-ticket equilibrium is $\max_{s\in S_1} w(v, s, p)$, which is a continuous function of p. Thus w^1 is continuous in μ and M.

Define $i(p, v) := \arg\max_{i\in[n]} v_i p_i$.[9] By Lemma 1, the optimal strategy for an agent v is to place her first ticket in the lottery for good $i(p, v)$, after which the marginal benefit for each additional ticket weakly decreases. Thus for all M, $x \in X_M$, we have

$$(1 - (1 - p^x_{i(p,v)})^k) v_{i(p,v)} \leq x(v) \cdot v \leq k p^x_{i(p,v)} v_{i(p,v)} = w^1(M, \mu^1(kp^x), v). \tag{38}$$

In (38), the first inequality follows because $x(v) \cdot v$ is a weakly better strategy than placing all k tickets in the lottery for good $i(p, v)$, the second inequality follows because $x(v) \cdot v$ is weakly less than k times the marginal benefit of the first ticket placed, and the final equality follows because $kp^x_{i(p,v)} v_{i(p,v)}$ is exactly the expected value of v in a 1-ticket lottery with win probability vector kp^x.

From this observation and Claim 1 it follows that

$$\lim_{M\to\infty} \sup_{x\in X_M} \frac{x(v) \cdot v}{w^1(M, \mu^1(kp^x), v)} = \lim_{M\to\infty} \inf_{x\in X_M} \frac{x(v) \cdot v}{w^1(M, \mu^1(kp^x), v)} = 1. \tag{39}$$

As $w^1(M, \mu, v) = x^1_M(v) \cdot v$, it remains to prove that

$$\lim_{M\to\infty} \sup_{x\in X_M} \frac{w^1(M, \mu^1(kp^x), v)}{w^1(M, \mu, v)} = \lim_{M\to\infty} \inf_{x\in X_M} \frac{w^1(M, \mu^1(kp^x), v)}{w^1(M, \mu, v)} = 1. \tag{40}$$

As w^1 is continuous in μ and M, it is sufficient to show that for all $i \in [n]$,

$$\lim_{M\to\infty} \sup_{x\in X_m} \frac{\mu^1_i(kp^x)}{\mu_i} = \lim_{M\to\infty} \inf_{x\in X_m} \frac{\mu^1_i(kp^x)}{\mu_i} = 1. \tag{41}$$

[8]Note that $V_i(p) = V_i(cp)$ as long as $cp_i \leq 1$ for all $i \in [n]$.

[9]Given p, we will ignore the set of agents who have multiple optimal strategies and for whom $i(p, v)$ is undefined, as this set has measure 0 and will not affect equilibrium parameters.

For sufficiently large M, we have

$$\mu_i^1(kp^x) = Mkp^x V_i(p^x) \tag{42}$$

by (37). Moreover, for any M, and any $x \in X_M$, $i \in [n]$, we have

$$M(1-(1-p_i^x)^k)V_i^k(p^x) \le \mu_i \le M(1-(1-p_i^x)^k)V_i^1(p^x), \tag{43}$$

where $V_i^k(p^x)$ denotes the measure of the set of agents whose optimal strategy under p^x is to place all k tickets into the lottery for item i, and $V_i^1(p^x)$ denotes the measure of the set of agents whose optimal strategy is to place at least 1 ticket in the lottery for item i. Equation (43) holds because the lefthand side is the mass of item i allocated to agents who place every ticket in the lottery for item i and the righthand side is the amount of item i which would be distributed if every agent who bid on item i placed every ticket in the lottery for item i.

Dividing (42) by (43) yields

$$\lim_{M\to\infty} \sup_{x\in X_m} \frac{kp_i^x}{(1-(1-p_i^x)^k)} \le \lim_{M\to\infty} \sup_{x\in X_m} \frac{\mu_i^1(kp^x)}{\mu_i} \le \lim_{M\to\infty} \sup_{x\in X_m} \frac{kp_i^x}{(1-(1-p_i^x)^k)} \frac{V_i^1(p^x)}{V_i^k(p^x)} \tag{44}$$

for sufficiently large M. The expressions on left and right both evaluate to 1, as by Claim 1 we have $\lim_{M\to\infty}\sup_{x\in X_m} kp^x/(1-(1-p_i^x)^k) = 1$, and $\lim_{M\to\infty}\sup_{x\in X_m} V_i^1(p^x)/V_i^k(p^x) = 1$ by Claim 2. The analogous claim for the infimum holds by identical reasoning, so (41) follows. □

Lemma 6 (1-ticket lotteries converge to CEEI)**.** *Let x_M^1 be an allocation corresponding to a 1-ticket equilibrium where $p_i < 1$ for all $i \in [n]$. For almost every agent v, $x_M^1(v) = ce_M(v)$.*

Proof. A market defined by μ, η and $M \ge \sum_i \mu_i$ corresponds to the instance of CEEI in which the set of items is $[n] + \{o\}$, where o indicates an outside option with value 0 and quantity $M - \sum_i \mu_i$.

When each agent is given a budget of 1, the price vector which assigns the price $1/p_i^{x_M^1}$ to each item and 0 to the outside option clears the market. To see this, observe that each agent has sufficient budget to buy a probability share $p_i^{x_M^1}$ in any item $i \in [n]$. For almost every agent $v \in V$, there exists a unique item $i \in K$ that maximizes $p_i^{x_M^1} v_i$, and spending her entire budget on this item maximizes her expected value. Thus almost every agent purchases a share $p_i^{x_M^1}$ in the same item on which she bids in the 1-ticket equilibrium, and picks up a share $1 - p_i^{x_M^1}$ of the outside option for free. By Lemma 1, this allocation is equivalent to $x^1_M(v)$. □

Theorem 1 follows directly from Lemmas 5 and 6, as every item is fully allocated for sufficiently large values of M.

D Proofs from Section 5

D.1 k-Ticket Lotteries are Not Trade Efficient

Proposition 6. *For $k > 1$, if η has full support on $\mathbb{R}^n$, any equilibrium with goods i and j such that $p_j < p_i < 1$ corresponds to an allocation that is not trade efficient.*

Proof. Let $V_{i,j}$ denote the set of agent types v such that

$$p_i v_i \geq p_j v_j \tag{45}$$

$$p_j v_j + (1 - p_j)p_i v_i \geq p_i v_i + (1 - p_i)p_i v_i \tag{46}$$

$$\forall l \notin \{i, j\}, v_l < 0. \tag{47}$$

By Lemma 1, agents in $V_{i,j}$ place their first ticket in the lottery for good i and their remaining tickets in the lottery for good j. Solving the two equations, we find that

$$V_{i,j} = \{v : \frac{p_j v_j}{p_i v_i} \in [1 - (p_i - p_j), 1], v_l < 0 \text{ for } l \notin \{i, j\}\}. \tag{48}$$

Because η has full support, $\eta(V_{i,j}) > 0$, and $V_{i,j}$ contains sets of agents with nonzero measure who have different relative values for items i and j and a chance of receiving nothing. In this case, there exists an exchange rate at which agents benefit from exchanging shares of i and j. □

D.2 Proof of Theorem 2

Let (Φ, p) be a k-ticket equilibrium and define the corresponding allocation $x(v) := \chi(\Phi(v), p)$. Let x' be a trade reallocation of x. We will show that there is a set V with $\eta(V) > 0$ such that for all $v \in V$,

$$\frac{x'(v) \cdot v}{x(v) \cdot v} < 2 - \frac{1}{k}.$$

Define the *cost* of a fractional allocation y to be

$$C(y) = \frac{1}{k}\sum_i y_i/p_i. \tag{49}$$

Because $\chi_i(\Phi(v), p) \le \Phi_i(v)p_i$, we have that under x, every agent receives an allocation with cost at most one:

$$C(x(v)) = \frac{1}{k}\sum_i \chi_i(\Phi(v), p)/p_i \le \frac{1}{k}\sum_i \Phi_i(v) = 1.$$

Futhermore, the total cost of an allocation is invariant under trade reallocation:

$$\int C(x(v))d\eta = \int C(x'(v))d\eta.$$

Thus there exists a subset V with $\eta(V) > 0$, such that for all $v \in V$ we have

$$C(x'(v)) \le C(x(v)) \le 1. \tag{50}$$

For $v \in V$, we will provide an upper bound on the ratio $\frac{x'(v)\cdot v}{x(v)\cdot v}$ by providing an upper bound for the numerator and a lower bound for the denominator.

For a given agent v, the set

$$Y_v := \arg \max_{\{y\,:\,C(y)\le 1\}} y \cdot v \tag{51}$$

consists of the solutions to the agent decision problem in an instance of CEEI in which each item i has an equilibrium price of $1/kp_i$. As such, every agent has an optimal strategy which purchases at most two items (Hylland and Zeckhauser (1979)). Fix $v \in V$ and select $y \in Y_v$ with the minimum number of nonzero entries $y(v)$. It follows from (50) and the definition of y that $x'(v) \cdot v \le y \cdot v$.

First, suppose that $y_j > 0$ for a single value $j \in [n]$. Then $y_j \le \min(kp_j, 1)$, and $y \cdot v \le \min(kp_j, 1)v_j$. Meanwhile, $x(v) \cdot v \ge (1 - (1 - p_j)^k)v_j$, since the chosen strategy must be at least as good as using all k tickets on item j. It follows that

$$\frac{x'(v)\cdot v}{x(v)\cdot v} \le \frac{y\cdot v}{x(v)\cdot v} \le \frac{\min(kp_j, 1)}{1-(1-p_j)^k} \le \frac{1}{1-(1-1/k)^k}, \tag{52}$$

where the final inequality follows because the expression above is maximized at $p_j = 1/k$. Finally, we have that

$$\frac{1}{1-(1-1/k)^k} \leq 2 - \frac{1}{k},$$

as the expressions coincide for $k = 1$, the lefthand side is strictly lesser when $k = 2$, and finally $1 - (1 - 1/k)^k \geq 1 - 1/e \geq \frac{1}{2-\frac{1}{k}}$ when $k \geq 3$.

Next, suppose that y has two nonzero components. Without loss of generality, suppose that these components are y_i and y_j and that $p_i \leq p_j$. First, we observe that

$$p_i < 1/k < p_j. \tag{53}$$

To see this, note that $p_i \leq p_j \leq 1/k$ implies that $1/kp_i \geq 1/kp_j \geq 1$, in which case $y \cdot v$ can be weakly improved by an allocation that selects a single item from $\arg\max_{l \in \{i,j\}} p_l v_l$. Similarly, $1/k \leq p_i \leq p_j$ implies $1 \geq 1/kp_i \geq 1/kp_j$, in which case $y \cdot v$ can be weakly improved by an allocation that selects a single item from $\arg\max_{l \in \{i,j\}} v_l$. Both of these conclusions contradict our choice of y so (53) follows.

Furthermore, if $y_i + y_j < 1$, we could weakly improve $y \cdot v$ by exchanging probability shares to increase the share of $\arg\max_{l \in \{i,j\}} p_l v_l$ until $y_i + y_j = 1$ (or until we reached an allocation with a single item, which would contradict our choice of y.) Thus we can assume

$$y_i + y_j = 1 \tag{54}$$

without loss of generality.

The agent must prefer $x(v)$ to randomizing between using all tickets on item i and using all tickets on item j. Therefore, for any $\alpha \in [0, 1]$ we have

$$x(v) \cdot v \geq \alpha(1 - (1 - p_i)^k)v_i + (1 - \alpha)(1 - (1 - p_j)^k)v_j. \tag{55}$$

We can set α so that the allocation of this strategy is a scaled-down version of y:

$$\alpha(1 - (1 - p_i)^k)/y_i = (1 - \alpha)(1 - (1 - p_j)^k)/y_j \triangleq \frac{1}{\gamma}. \tag{56}$$

It follows from the fact that $x'(v) \cdot v \le y \cdot v$ and (55), (56) that

$$\frac{x'(v) \cdot v}{x(v) \cdot v} \le \frac{y \cdot v}{\frac{1}{\gamma} y \cdot v} = \gamma. \tag{57}$$

We now provide a upper bound for γ. Using (54) and solving (56) for α, we can solve for γ:

$$\gamma = \frac{y_i}{1-(1-p_i)^k} + \frac{y_j}{1-(1-p_j)^k} \tag{58}$$

Solve for the choice of y satisfying $y_i + y_j = 1$ and $C(y) = 1$, and substitute into (58). After simplification we get the following expression for γ, which we write in terms of the functions f and h:

$$f(p_i, p_j) := \frac{p_j - 1/k}{p_j - p_i} h(p_i) + \frac{1/k - p_i}{p_j - p_i} h(p_j) = \gamma, \tag{59}$$

where

$$h(p) := \frac{pk}{1-(1-p)^k}. \tag{60}$$

By Lemma 7, $f(p_i, p_j) = \gamma$ is maximized by taking $p_i \to 0$ and $p_j \to 1$, in which case it converges to $2 - \frac{1}{k}$.

Lemma 7. *For $0 \le p_i < 1/k < p_j$, the function f defined in* (59) *is decreasing in p_i and increasing in p_j.*

Proof. We can express $f(p_i, p_j)$ as

$$f(p_i, p_j) = h(p_j) - (p_j - 1/k) \frac{h(p_j) - h(p_i)}{p_j - p_i}. \tag{61}$$

We claim that $\frac{h(p_j)-h(p_i)}{p_j-p_i}$ is increasing in p_i on $[0, p_j]$. This follows because

$$\frac{d}{dp_i} \frac{h(p_j) - h(p_i)}{p_j - p_i} = \frac{(p_j - p_i)h'(p_i) + h(p_j) - h(p_i)}{(p_j - p_i)^2}, \tag{62}$$

and the numerator is positive by the convexity of h on $(0, 1)$ (see Lemma 8). Therefore, for $p_j \ge 1/k$, the expression in (61) is decreasing in p_i.

Analogously, we can express $f(p_i, p_j)$ as

$$f(p_i, p_j) = h(p_i) + (1/k - p_i) \frac{h(p_j) - h(p_i)}{p_j - p_i}. \tag{63}$$

Because h is convex on $(0,1)$, $\frac{h(p_j)-h(p_i)}{p_j-p_i}$ is increasing in p_j for $p_j > p_i$. Thus the expression above is increasing in p_j if $p_i < 1/k$.

□

Lemma 8. *The function h defined in* (60) *is convex on the interval* $(0,1)$.

Proof. It is equivalent to show that $g(x) = \frac{1}{k}h(1-x) = \frac{1-x}{1-x^k}$ is convex. Differentiating twice, we see that

$$g''(x) = \frac{kx^{k-2}((k-1)x^{k+1} - (k+1)x^k + (k+1)x - k + 1)}{(x^k-1)^3}. \tag{64}$$

Thus $g''(x) \geq 0$ if and only if the quantity

$$-(k-1)x^{k+1} + (k+1)x^k - (k+1)x + k - 1 \tag{65}$$

is greater than or equal to 0. The expression in (65) is equal to 0 when $x = 1$, so to prove the convexity of $g(x)$ it suffices to show that (65) is weakly decreasing on $(0,1)$. The derivative of (65) is

$$(k+1)((1-k)x^k + kx^{k-1} - 1). \tag{66}$$

We claim that this quantity is weakly less than 0. Equation 66 is 0 at $x = 1$, and differentiating reveals that it is weakly increasing in x on $(0,1)$. □

E Proofs from Section 5

The following lemma is useful in the proofs of Propositions 4 and 5.

Lemma 9. *At equilibrium in the two-item setting, almost every agent adopts one of three strategies:*

1. *Agents with $v_2 < p_1$ place k tickets in the lottery for item 1 and receive $x(v) = (1-(1-p_1)^k, 0)$.*

2. *Agents with $p_1 < v_2 < 1$ place $k-1$ tickets in the lottery for item 1, place 1 ticket in the lottery for item 2, and receive $x(v) = (1-(1-p_1)^{k-1}, (1-p_1)^{k-1})$.*

3. *Agents with $v_2 > 1$ place k tickets in the lottery for item 2 and receive $x(v) = (0,1)$.*

Proof. As $F(1) \geq \mu_1$, item 1 is fully allocated. As a result, $\int 1 - x_1(v)\ d\eta = 1 - \mu_1 = \mu_2$, so $p_2 = 1$ in any win probability vector p corresponding to a k-ticket equilibrium in the two-item setting. As a result, agents with $v_2 > 1$ are guaranteed their preferred item when they adopt Strategy 3. For agents with $v_2 < 1$, every strategy that places less than $k-1$ tickets in the lottery for item 1 is dominated by Strategy 2. Finally, agents who adopt Strategy 1 have expected value $(1-(1-p_1)^{k-1})v_1 + (1-p_1)^{k-1}p_1 v_1$, while agents who adopt Strategy 2 have expected value $(1-(1-p_1)^{k-1})v_1 + (1-p_1)^{k-1}p_2 v_2$. Plugging in $v_1, p_2 = 1$ implies that agents with $v_2 < p_1$ adopt Strategy 1 and agents with $p_1 < v_2 < 1$ adopt Strategy 2. □

E.1 Proof of Proposition 4

Proof of Proposition 4. Lemma 9 implies that in any k-ticket equilibrium in the 2-item setting, plotting expected value in terms of v_2 yields a piecewise linear function (as depicted in Figure 2.) Specifically, all agents with $v_2 < p_1$ adopt Strategy 1 and have expected value $1-(1-p_1)^k$. Agents with $p_1 < v_2 < 1$ adopt Strategy 2, and expected value increases linearly to 1 on the interval $[p_1, v_2]$. Finally, agents with $v_2 > 1$ adopt Strategy 3 and have expected value v_2. Thus the geometry of the expected value curve implies that two such curves cross at most once.

Fix a k-ticket equilibrium in the two-item setting with win probability vector p. By Lemma 9, we have

$$\mu_1 = (1-(1-p_1)^k)F(p_1) + (1-(1-p_1)^{k-1})(F(1) - F(p_1)). \tag{67}$$

Rearranging this equation, we get

$$F(1) - \mu_1 = (1-p_1)^{k-1}(F(1) - p_1 F(p_1)). \tag{68}$$

By definition, the left side of this equation is a nonnegative value constant in p and k. As a result, we have $F(1) - p_1 F(p_1) \geq 0$. Thus the right side of the equation is decreasing in p_1 and k, which implies that in the two-item setting, a smaller k corresponds to a larger p_1 at equilibrium.

Fix $i < j$, and let x^i and x^j be allocations corresponding to i and j-ticket equilibria in the two-item setting. The win probability for item 1 is higher in the i-ticket equilibrium, so agents with small v_2, who adopt Strategy 1 in both equilibria, prefer x^i to x^j. If the expected value curves corresponding to x^i and x^j do not cross, then every agent with $v_2 < 1$ prefers x^i to x^j. If the expected value curves cross, then every agent with v_2 less than the point of indifference prefers x^i to x^j. Thus if v prefers x_i to x_j, the same is true of any agent w with $w_2 < v_2$. □

E.2 Proof of Proposition 5

Proof of Proposition 5. Let x be an allocation corresponding to a k-ticket lottery in the two-item setting. Our proposed CEEI takes place in the same market, except that the two items have quantities defined by the vector

$$\tilde{\mu} := \int x(v)\, dF(v) \leq \mu. \tag{69}$$

We claim that $c := (\frac{1}{1-(1-p_1)^k}, \frac{p_1}{1-(1-p_1)^k})$ is an equilibrium price vector in this market. To see this, consider an agent $v = (1, v_2)$ who wishes to choose (x_1, x_2) to maximize $x_1 + x_2 v_2$ subject to the unit demand constraint $x_1 + x_2 \leq 1$ and the budget constraint $x_1 c_1 + x_2 c_2 \leq 1$. A straightforward calculation verifies that for agents with $v_2 < p_1$, $p_1 < v_2 < 1$, and $v_2 > 1$, the allocations described in Lemma 9 solve this optimization problem and clear the market. □