Open-ended matching with and without markets

Sophie Bade and Michael Mandler*

Royal Holloway College, University of London

This version: March 2022

Abstract

Suppose agents are to be matched to objects and arrive over time without a definite terminal date. Although the set of core matchings can then be empty, an extended version of the top trading algorithm shows that Pareto-optimal weak-core matchings always exist. Optimal matchings face a difficulty however: some of the agents linked by chains of trades may have lifespans that do not overlap, thus obstructing their trades. To address this problem, we let matchings be implemented via competitive markets. Competitive equilibria always exist and any matching in the core can be competitively implemented. Moreover full core equivalence, where allocations are in the core if and only if they can be competitively implemented, holds for a dense set of models. The extended algorithm also yields a strategyproof mechanism, comparably to the finite model.

JEL codes: C71, C78, D50.

Keywords: matching, competitive equilibria, core, top trading cycles.

*Address: Department of Economics, Royal Holloway College, University of London, Egham, Surrey TW20 0EX, United Kingdom. Email: sophie.bade@rhul.ac.uk and m.mandler@rhul.ac.uk. The authors thank the British Academy, the Minerva Foundation, and the University of London for financial support for this research.

1 Introduction

In classical matching problems, such as the assignment of donor kidneys to patients with end-stage renal disease, there are finitely many agents and objects. These bounds are however modeling abstractions that ignore that in time new patients and donors will appear; these arrivals can potentially be matched to some of the original agents. The matching of doctors to work assignments in a hospital emergency room is similar: the participants may be unhappy with their original assignments but an efficient rescheduling may involve different doctors who will arrive at a later date and may also want to switch their assignments. Families in public housing provide a third example: Pareto improvements to an ex ante inefficient allocation may require families to swap homes with not-yet-arrived families. Indeed it is difficult to find cases of matching where all of the concerned parties can coordinate at a single point in time. While the lifespans of institutions and of the relevant agents may not stretch into an endless future, models with results driven by an exogenously imposed terminal date are misleading and at odds with how agents perceive the future.[1]

We therefore consider a matching model à la Shapley and Scarf [18] that is open-ended. The model will have no terminal date: the set of agents and the set of objects ('houses') will both be the set of natural numbers $\mathbb{N}$. Each agent i will have a linear order over the set of objects and own one object which will also have the label i. These preferences can capture a variety of demographic structures. For example, generations can overlap if any agent i will trade his endowment only for objects that appear within some time span around his endowment; implicitly agent i is alive only at these dates.

To address the coordination problem that efficiency can require exchanges among agents who are never alive at the same time, we will follow the original work of Shapley and Scarf. Agents will trade objects in a competitive equilibrium: they will sell their endowments and use the proceeds to buy preferred objects. In a market, the agents linked indirectly by sequences of trades never have to meet or even live simultaneously, they merely need to form accurate expectations of the prices of the objects that appear in the future.

The Gale top trading cycles algorithm supplies the workhorse for matching finitely many

[1] Osborne and Rubinstein [16] argue persuasively that models of infinitely repeated games can more accurately capture the perceptions of agents in a model that will in fact end up being finite.

agents and objects. Each agent points to the owner of his favorite object. The pointings must form at least one cycle, one of which retires from the algorithm with its agents matched to the objects they pointed to. After enough rounds of pointing, all of the agents will be matched to objects. It is easy to show that the matching that results lies in the core: no coalition of agents C can block since the first agent in C to retire from the algorithm that gets a different object with C must be worse off with that object. Moreover, by letting the prices of objects descend with the date at which they exit the algorithm, top trading constructs a competitive equilibrium that implements the matching. Since in addition every competitive equilibrium allocation is a core matching, full core equivalence obtains.

With open-ended matching, pointing need not lead to cycles. We will therefore allow 'chains' to exit the algorithm, for example, rays where one agent i points to a second agent who points to a third, and so on with no repeats, and two-sided chains where every agent has one predecessor and one successor and again there are no repeats. If a chain has a root agent then this agent's endowment is discarded but we allow agents in later chains to point to this object.

Letting agents depart from top trading in configurations other than cycles is a double-edged sword. Some parts of the Gale program will proceed as successfully as they did in the finite model. For example, although the reformulated algorithm will not terminate after finitely many rounds of pointing, it will terminate after a transfinite number of rounds. Top trading therefore still leads to a matching; the finiteness of the number of rounds is immaterial. Just as importantly, rounds remain well-ordered. So, for any given property of interest, there will be a *first* round to retire that satisfies that property. Every welfare and efficiency conclusion in this paper will rest on this principle.

On the negative side, the exit of non-cycles undermines the argument that top trading must lead to a core matching; the core in fact can be empty. But the damage done is modest. Our revision of top trading will always generate a matching that is both Pareto optimal and in the weak core (no coalition can use its endowments to make each of its members strictly better off).

To assess coordination when limited lifespans prevent agents from meeting, we show first that our model always has a competitive equilibrium and second that a core matching (if one exists) can be implemented by a competitive equilibrium. Half of core equivalence therefore holds: core matchings can be achieved in equilibrium. But half does not and the exit of non-cycles from top trading is again to blame. Competitive equilibrium matchings can then be inefficient and lie outside the core. Indeed in some cases competitive matchings can be strictly Pareto dominated. And the bad news even gets a little worse: when the core is empty it may happen that no matching in the weak core can be competitively implemented.

Given the inefficiency of equilibria in the overlapping generations model of general equilibrium theory, these pathologies may not come as a great surprise. The live question is therefore how intractable or robust the negative results are. In the natural topology for our model, there is a dense set of models where agents retire from our pointing algorithm only in cycles and where therefore core matchings exist. For these models full core equivalence holds: the set of core matchings coincides with the set of matchings that can be implemented as competitive equilibria. So, from the glass is half full perspective, inefficient equilibria are possible but there is at least a rich supply of models where equilibria are necessarily efficient.

Beyond competitive equilibria, one great advantage of Gale's algorithm for finite models is that top trading cycles provide a simple proof that a strategyproof direct revelation mechanism can implement the core (Roth [17]). With open-ended matching, that goal is unattainable since the core can be empty. But a Gale-like algorithm can deliver nearly as good a result: an extension of Gale will show that a strategyproof direct mechanism can implement the core when the core is non-empty and otherwise reach Pareto optima in the weak core. There is consequently an incentive-compatible way to lead agents to efficient matchings.

In a finite Shapley-Scarf model, each agent is assigned a distinct object and consequently no unassigned objects are left over: the issue of disposal does not arise. With infinitely many agents and objects, objects can remain unassigned, e.g., when every $i \in \mathbb{N}$ receives object $i + 1$ then no one consumes object 1. To follow as orthodox a path as possible, the main model of the paper assumes free disposal. In section 7, we point out the relatively minor modifications required when disposal is impossible, such as when doctors are assigned to ER slots that must be filled.

Gale's top trading cycles are one of the most elementary and appealing constructions in economic theory, and many of the arguments in this paper descend from Gale. By generalizing top trading cycles to infinitely many agents and objects, we will be able to identify more precisely the analytical power of top trading. Its crucial feature is that the rounds of agents that depart top trading are well-ordered, as opposed to the stronger requirement that there are finitely many rounds. For example, the well-ordering of rounds implies that if a matching μ differs from top trading there must be a first round of departures from top trading with an agent i that is assigned a different object under the two matchings. Since top trading matches i to his favorite among the objects not yet assigned i must be worse off with μ. Top trading thus leads to an efficient matching. Nothing in this simple argument changes if i departs at a round that corresponds to transfinite ordinal rather than at a finite round.

Well-orderings were first defined by Cantor and curiously this paper's other notable tool, the diagonalization we use to prove the existence of competitive equilibria, is also a Cantor construction.

1.1 Related Literature

There is a growing literature on matching over an infinite horizon. In contrast to the present treatment, most contributions assume that the preferences of agents are randomly drawn from a pool of possible preferences. Unver [20], Akbarpour et al. [2], and Anderson et al. [3] study unilateral matching problems where each agent is endowed with one object and has a 'dichotomous' preference that exhibits indifference among all objects an agent prefers to his endowment. The size of the set of agents who are matched to the endowments of other agents then provides a natural measure of welfare. To consider the trade-off between this measure and the time agents spend waiting for a match, these papers posit random processes that govern the entry and exit of agents and the compatibilities of agents with the endowments of others.

Leshno [14], Bloch and Cantala [7], Schummer [19], Arnosti and Shi [4] and Agarwal et al. [1] study unilateral matching problems over an infinite horizon without initial endowments, for example, allocation of the kidneys of deceased donors. They show that the optimal organization of waiting lists for these objects depends on the heterogeneity of agents' preferences.

Comparably to the countable set of agents used here for one-sided matching, Jagadeesan [13] studies bilateral matching with countably many agents. In an interesting contrast to the one-sided case, where we will see that the Gale algorithm can stall, the Gale-Shapley deferred acceptance algorithm applies without change to two-sided matching. Fleiner [10] earlier considered two-sided matching with an arbitrary set of agents. Motivated by adoption, Baccara et al. [5] study a bilateral matching market in which prospective parents and children of two possible types stochastically enter the adoption pool. Following Baccara et al., Doval and Szentes [9] have analyzed a bilateral matching market of stochastically arriving impatient agents with dichotomous preferences.

Gale's top trade cycles were first linked to strategyproof mechanisms by Roth [17]. See Hassidim et al. [12] on the persistence of preference misrepresentation with strategyproof matching.

2 Preferences and matchings

At each date i in $\mathbb{N} = \{1, 2, ...\}$, there is one object i which is owned by an agent who is also labeled i. Each agent i has a linear preference $\succsim^i$ on $\mathbb{N}$ and $\succ^i$ will be the associated strict preference.[2] A **profile** of preferences assigns a linear preference to each agent that we will represent by the profile of associated strict preferences $\succ = (\succ^i)_{i \in \mathbb{N}}$.

The **favorite** of agent i endowed with $\succsim^i$ from a set of objects $H \subset \mathbb{N}$ is the object $j \in H$ such that $j \succsim^i k$ for $k \in H$, the **second favorite** is the object $s \in H$ such that $s \succsim^i k$ for $k \in H \setminus \{j\}$, and so forth. When not stated explicitly, the reader should assume that $H = \mathbb{N}$. We assume throughout that each agent i has a favorite from any $H \subset \mathbb{N}$ with $i \in H$. If a binary relation $\succsim^i$ satisfies this assumption and linearity then $\succsim^i$ (or $\succ^i$) is a **preference**. At times we impose the stronger condition that there is a half lifespan L such that each agent i prefers his endowment to all objects that appear more than L periods from i. Formally, profile $\succ$ has **bounded lifespans** if there exists a $L > 0$ such that, for

[2] A binary relation $\succsim$ is **linear** if it is complete, transitive, and antisymmetric ($x \succsim y$ and $y \succsim x$ imply $x = y$). Though $\succ$, the asymmetric part of $\succsim$, is not complete, we will also call $\succ$ linear.

each agent i and object j, $|i-j| \geq L$ implies $i \succ^i j$. Since agents will always be able to consume their endowments, their rankings of the objects that appear beyond their lifetimes are irrelevant; with bounded lifespans we could therefore let preferences over these objects be undefined.

The model allows a loose interpretation of how agents and objects are associated with time periods. For example, $1, ..., l$ could designate agents and objects that appear at day one, $l+1, ..., 2l$ could designate agents and objects that appear at day two, etc.

A **matching** or **allocation** is a one-to-one map μ from the set of agents $\mathbb{N}$ to set of objects $\mathbb{N}$: for each object $j \in \mathbb{N}$ there is at most one agent i with $\mu(i) = j$. Until section 7, objects can be freely disposed of and thus the image of μ does not have to equal $\mathbb{N}$.

The coalition $S \subset \mathbb{N}$ **blocks** matching μ at $\succ$ if there is a one-to-one map $\nu : S \to S$ such that $\nu(i) \succsim^i \mu(i)$ for all $i \in S$ and $\nu(j) \succ^j \mu(j)$ for some $j \in S$ and **strictly blocks** μ if there is a one-to-one ν such that $\nu(i) \succ^i \mu(i)$ for all $i \in S$. A matching μ is in the **core** of $\succ$ if no coalition can block μ at $\succ$ and is in the **weak core** if no coalition can strictly block μ.

3 Top trading

In a finite matching problem, each agent $i \in \{1, ..., n\}$ owns object i. The Gale top trading cycles algorithm consists of rounds where each agent points to his favorite object among those still available. At least one cycle must form in the first round and the agents in one of these cycles, S^1, receive their favorites and retire from the algorithm. Agents then point anew to their favorites from the remaining objects and a second cycle S^2 retires, and so on. The matching μ that results must lie in the core. For a proof, suppose some coalition B can block μ and let S^a be the first cycle to retire with an agent i who prefers the object j that B assigns to i over the object $\mu(i)$. Since $\mu(i)$ is agent i's favorite from the objects owned by the agents who retire at round a or later, agent j must be in a cycle S^b with $b < a$. Since agent j must be in B and B assigns $\mu(j)$ to j, the owner of $\mu(j)$ must also be in B. Iterating this argument, B must contain all of S^b and must assign to each agent $k \in S^b$ the object $\mu(k)$. But then B assigns object j to both i and some agent in S^b, a contradiction.

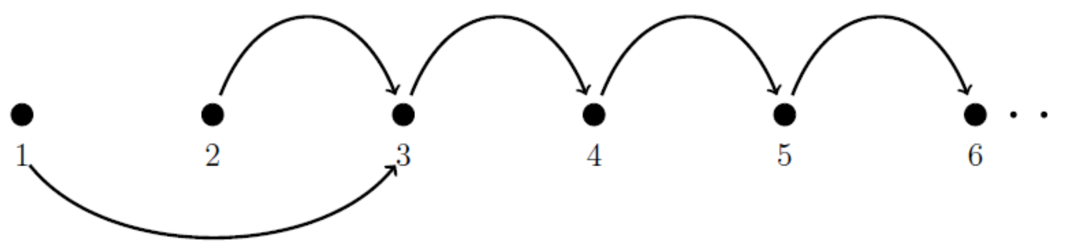

Figure 1: An empty core

This argument has two crucial components: (1) there is a 'first' exiter who is better off with B than with his top-trading allocation and (2) this agent points to an object in a cycle rather than some other type of directed graph. The first component will carry over readily to an infinite setting with well-ordered exits. The appeal to cycles, on the other hand, is doomed: in extending Gale we will have to allow non-cycles to depart from our algorithm. This change will undermine both the existence of core matchings and the optimality of competitive equilibria. Finiteness in contrast plays only an indirect role in Gale's argument. It implies that agents must depart in cycles, but we will see that finiteness is not needed to show that top trading algorithms terminate.

We begin with an example to illustrate that cycles need not form in the Gale algorithm.

Example 1 *Suppose that the favorite of each agent $i \neq 1$ is object $i+1$ and that the favorite of agent 1 is object 3. Letting a solid arrow point from each agent to the agent's favorite, the preference profile is pictured in Figure 1. In this example and all the examples to follow, we can set the agents' remaining preference rankings to be consistent with the bounded lifespan assumption, for instance by letting each agent i's first unspecified favorite (in this case the second favorite) be object i. The preferences in the present example would then be: for each $i \geq 2$, $i+1 \succ^i i \succ^i j$ for $j \notin \{i, i+1\}$, and $3 \succ^1 1 \succ^1 j$ for $j \notin \{1, 3\}$.*

While no cycle appears in Example 1, there are 'rays,' that is, infinite sequences of agents $(s_1, s_2, ...)$ such that, for all $j \geq 1$, the favorite of s_j is s_{j+1} and no agent appears more than once. Example 1 has two maximal rays, $(2, 3, 4, ..)$ and $(1, 3, 4, ...)$. Suppose we mimic the Gale algorithm by assigning each agent in one of these rays his favorite object and then retiring that ray from the algorithm. A single agent, either 1 or 2, would then remain and the second round would assign this agent his endowment.

This removal of a ray rather than a cycle unfortunately voids the argument that the algorithm must generate a core matching: a subset of agents that retires from the algorithm could well join some of the surviving agents to form a blocking coalition. Suppose in Example 1 that the first round assigns objects $(3, 4, 5, ...)$ to agents $(2, 3, 4, ...)$. Then $\{3, 4, ...\}$ joined with $\{1\}$ can block by switching object 1 from agent 2 to agent 1. Example 1 in fact has an empty core. If a matching fails to assign objects $(3, 4, 5, ...)$ to agents $(2, 3, 4, ...)$ then $\{2, 3, 4, ...\}$ can block and if it fails to assign $(3, 4, 5, ...)$ to $(1, 3, 4, ...)$ then $\{1, 3, 4, ...\}$ can block. Since object 3 can be assigned to only one agent, the core must be empty.

While the emptiness of the core is unwelcome, Example 1 relies on agents that block allocations even when they gain nothing by doing so – comparably to the empty-core examples that arise in finite matching when agents can be indifferent among objects. The weak core can therefore identify more sharply which allocations will survive unchallenged. Since however weak-core allocations can be inefficient, our goal will be to show that there are allocations that both lie in the weak core and are Pareto optimal.

When top trading cycles is run on a finite set of agents, the agents in a subset S that retires from the algorithm at some round take their endowments with them, which does no harm since these endowments are allocated to agents in S. In an open-ended model, however, when a ray retires no one in the ray consumes the endowment of the agent at the root of the ray. If that endowment is subsequently unavailable, inefficiency can result, even when each set of retirees is as large as possible and a core allocation exists.

Example 2 *Suppose the favorite of both agents* 1 *and* 2 *is object* 1*, agent* 2*'s second favorite is object* 3*, and the favorite of each agent* $i \geq 3$ *is object* $i+1$*. See Figure* 2 *where solid and dashed arrows continue to point to favorites and second favorites respectively. In round* 1*, the ray that matches each agent* $i \geq 3$ *with object* $i+1$ *forms. If we remove the endowments of this ray then in round* 2 *only the cycle that matches agent* 1 *with object* 1 *can form. In the final round, agent* 2 *is matched with his endowment. If we could offer the discarded object* 3 *to agent* 2 *then we would achieve both a Pareto improvement and a core matching.*

Showing that Pareto-optimal weak-core allocations exist will therefore require some modifications to top trading cycles. The rules for exiting the algorithm must be loose enough

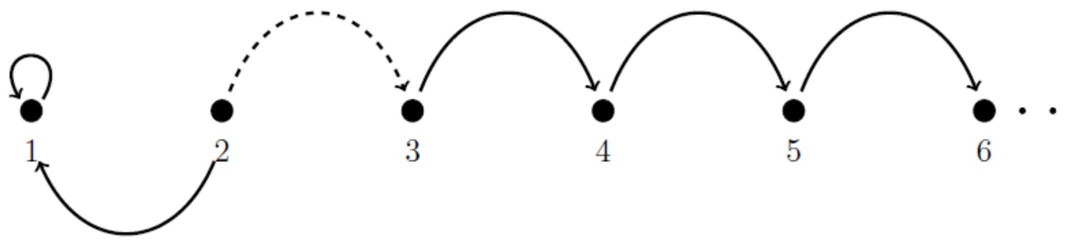

Figure 2: Inefficiency generated by top trading cycles

that at least one of the agents who has not yet been assigned an object can exit – but not so loose that agents can end up with objects that are not their favorites among those still available. In line with Gale, at every round some group will retire and each agent in the group will be assigned his favorite currently available object and give up his endowment. But other configurations beyond cycles will be allowed to retire, current retirees can be matched to objects that previous retirees have discarded, and finally the algorithm may terminate only after transfinitely many rounds.

Let a **chain** S be a subset of $\mathbb{N}$ that are assigned indices – distinct from their original labeling in $\mathbb{N}$ – that consist of a set of consecutive integers. A chain S can have an element with a maximal index, $\max S$, and an element with a minimal index, $\min S$, but is not required to have either. Chains may therefore be infinite. A notational warning: the alternative indices used to define chains are used only in this section.

Each round of the **extended top trading** algorithm begins with a set of unmatched agents $U \subset \mathbb{N}$ and a set of discarded objects D that does not intersect U. Let $s \rightarrow_{U,D} t$ mean that $s \in U$ and t is the favorite object of agent s from $U \cup D$. Given U and D, a chain S is **admissible** if $S \subset U$ and, for all $s \in S$,

(1) if $s \neq \max S$ and has index i then $s \rightarrow_{U,D} t$ where t has index $i + 1$, and

(2) if $s = \max S$ then there is a $t \in \{\min S\} \cup D$ such that $s \rightarrow_{U,D} t$.

So each agent in a chain S except possibly $\max S$ points to another agent in S while $\max S$ either points to $\min S$ or a discarded object. Call the object t that $s \in S$ points to, given by (1) or (2), the **successor** of s determined by S. If an admissible chain S has a $\min S$ agent but lacks a s such that $s \rightarrow_{U,D} \min S$ then $\min S$ is the **root** of S.

If $\max S \to_{U,D} \min S$ then S is a **cycle**. The cycles are finite but a finite admissible chain can also arise when $\max S$ points to a discarded object. An infinite chain can be one of the rays discussed above, where the positive integers supply the indices, or a reverse ray, where the negative integers supply the indices and $\max S$ points to a discarded object, or finally a 'two-sided' chain, where the indices are the entire set of integers.

Given any nonempty U and an arbitrary D, an admissible chain will always exist. Beginning with some $i \in U$, suppose that $i \to_{U,D} j \to_{U,D} ....$ This sequence will (A) eventually repeat, (B) end finitely by reaching an element of D, or (C) form a set of infinitely many distinct agents. Each case leads to at least one admissible chain. For (A), we can extract from $i \to_{U,D} j \dot{\to}_{U,D} ...$ a sequence that begins and ends with the same element of U and has no other repetitions, which defines an admissible cycle. In (B), the penultimate entry in the sequence forms a singleton admissible chain. Case (C) defines an admissible ray. Given a nonempty U and arbitrary D, extended top trading will select one of the admissible chains S and match each $i \in S$ with the successor of i determined by S. We remove S from U and say that S **exits** the algorithm. The set of unmatched agents that enter the next round is therefore $U \setminus S$. If there is a $\max S$ and $\max S \to_{U,D} d$ where $d \in D$ then we remove d from D and if $\min S$ exists and is a root then we add $\min S$ to $D \setminus \{d\}$. This addition and/or removal fixes the set of discarded objects for the next round.

The algorithm begins with $U = \mathbb{N}$ and $D = \varnothing$ and terminates at the first round where the set of unmatched agents that enters the round is empty. With rounds as they have been defined so far, the algorithm need not terminate after finitely many rounds nor in the limit of finitely many rounds. The failure of finite termination is unsurprising and is not caused by the rule that only one admissible chain exits at each date: in the following failure of finite termination, each round produces only one admissible chain.

Example 3 *For each agent $i \geq 2$, i's favorite object is $i-1$ and his second favorite is $i+1$. Agent 1 has object 2 as his favorite. See Figure 3. In the first round of extended top trading, only the cycle between agents 1 and 2 forms. In the second round, only the cycle between 3 and 4 can form, and so on. So, after finitely many rounds, infinitely many agents remain unmatched.*

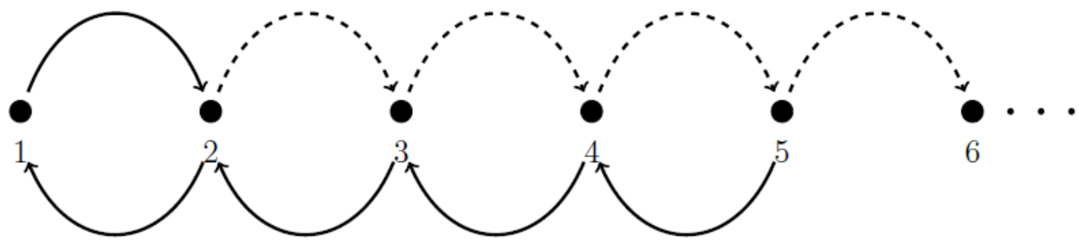

Figure 3: Top trading cycles that do not finitely terminate.

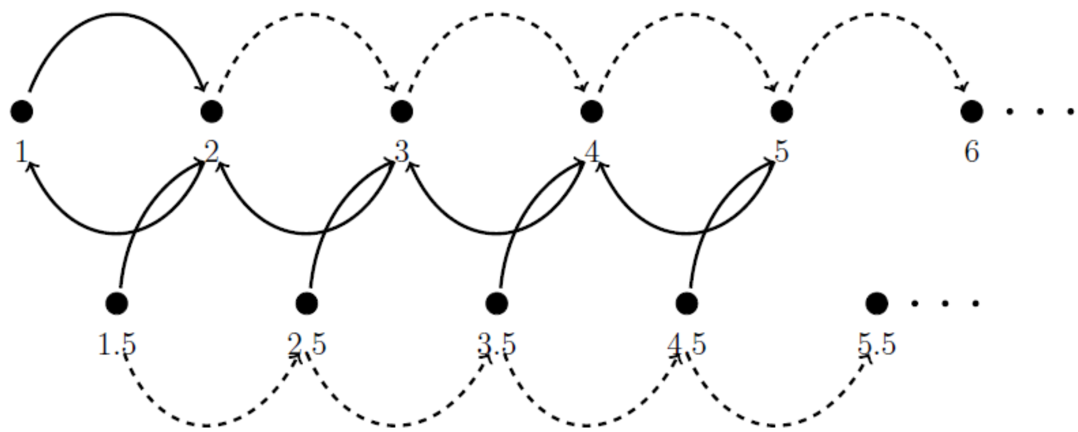

Figure 4: Top trading cycles that terminate transfinitely

In Example 3, each agent i is at least assigned to an admissible cycle after finitely many rounds. In the following example, some agents fail to exit even after all of the finite rounds have passed.

Example 4 *Modify Example 3 by adding new agents with labels* 1.5*,* 2.5*,* 3.5*, ... Each new agent i has $i + \frac{1}{2}$ as his favorite and $i + 1$ as his second favorite. See Figure* 4*. As long as any of the original whole-number agents remains unmatched, all the new fractional agents remain unmatched. Since it takes all of the finite rounds to match the whole-number agents, none of the fractional agents is matched at any finite round.*

Our final change to the Gale algorithm will therefore be to let termination occur transfinitely. To do so, we let the number of rounds equal an ordinal number and define the sets of agents that enter and exit each round of the algorithm recursively.[3] Letting α be an ordinal, U^{α}_{end} will indicate the set of unmatched agents at the end of round α. Given an ordinal β and U^{α}_{end} for each $\alpha < \beta$, the set of unmatched agents that enters round β will be

[3]Though 0 is the first ordinal, our examples will always let 1 be the label of the first round.

given by $U = \bigcap_{\alpha<\beta} U^{\alpha}_{\text{end}}$. Some admissible chain $S \subset U$ then forms and exits and we set $U^{\beta}_{\text{end}} = U \setminus S$.

The sets of matched agents and discarded objects are also defined recursively. Letting M^{α} be the set of matched agents that enters round $\alpha < \beta$, the set of matched agents that enters round β is given by $M^{\beta} = \bigcup_{\alpha<\beta} M^{\alpha}$. Setting $\mu(i)$ to equal the successor of i determined by the departing chain that contains i, $M^{\beta} \setminus \mu(M^{\beta})$ is the set of discarded objects that enters round β.

When β is a successor ordinal and β therefore has an immediate predecessor $\beta - 1$, the unmatched agents that enter round β could be equivalently defined as $U^{\beta-1}_{\text{end}}$. But when β is a limit ordinal – for example ω the first ordinal that succeeds the finite numbers – we must use the $\bigcap_{\alpha<\omega} U^{\alpha}_{\text{end}}$ definition. So, in Example 4 the fractional agents form the set of unmatched agents that enters round ω while each preceding set of unmatched agents contains infinitely many whole-number agents. Similar remarks apply to the definition of M^{β}. See the end of the section for a formal definition of the algorithm and of the recursion step in particular.

Comparably to Gale, extended top trading **terminates** at the first round α such that $U^{\alpha}_{\text{end}} = \varnothing$. To see that the algorithm is guaranteed to terminate when the rounds extend to arbitrary ordinals, let the algorithm proceed through ω_1 rounds, where ω_1 is the first uncountable ordinal. Each round that begins with a nonempty set of unmatched agents will match at least one of them and the set of rounds that precede round ω_1 is uncountable.[4] Hence there must be a round $\beta < \omega_1$ where the set of unmatched agents at the end of the round, U^{β}_{end}, is empty. Keep in mind that we are not constructing a procedure that agents will follow in real time: the goal is only to show that Pareto-optimal weak core matchings exist. Since the algorithm specifies a $\mu(i)$ for each $i \in \mathbb{N}$, the map μ defines a matching and we say it is **generated** by the algorithm.

A 'first agent to exit' argument now shows that any matching μ generated by extended top trading lies in the weak core. As this is the first of many cases where we exploit the well-ordering of rounds, we spell out the details. Apply extended top trading and let $N \subset \mathbb{N}$ be an arbitrary coalition. Since rounds are assigned only to ordinals, the nonempty set of

[4] Since each ordinal β equals the set of ordinals $\{\alpha : \alpha < \beta\}$, ω_1 has uncountably many predecessors. The existence of ω_1 does not depend on the axiom of choice, which we have taken some care to avoid.

rounds where some agent in N departs the algorithm,

$$\{\gamma : N \cap S \neq \varnothing \text{ where } S \text{ is the admissible chain that exits at } \gamma\},$$

is well-ordered, that is, it has a minimal element β. So β is the first round at which some $i \in N$ exits and hence, for any round $\alpha < \beta$, $N \subset U^{\alpha}_{\text{end}}$. Letting $U = \bigcap_{\alpha<\beta} U^{\alpha}_{\text{end}}$ and D, respectively, be the unmatched agents and discarded objects that enter round β, $\mu(i)$ is the favorite of i from $U \cup D$. Since U contains N, agent i cannot strictly prefer any object in N to $\mu(i)$ and so N cannot strictly block μ.

For Pareto optimality, consider an arbitrary matching $\eta \neq \mu$ and the first round β at which the chain S that exits the algorithm at β contains an agent k with $\eta(k) \neq \mu(k)$. Since every agent j who exits prior to β receives $\mu(j) = \eta(j)$, $\eta(k)$ must be an element of $U = \bigcap_{\alpha<\beta} U^{\alpha}_{\text{end}}$ or of the discarded objects D that enter round β. Hence $\mu(k) \succ^{k} \eta(k)$ and η cannot be a Pareto improvement.[5]

We have proved:

Theorem 1 *For any profile $\succ$, there exist allocations that are Pareto optimal and in the weak core of $\succ$.*

That the proof above invokes the first uncountable ordinal ω_1, a difficult concept to grasp, is no accident. We show in the Appendix that the procession of retirements can last as long as any countable ordinal; the least upper bound on the algorithm's termination date is therefore ω_1. While countable, the termination dates can be large countable ordinals: they therefore need not be computable and can be nearly as intangible as ω_1. But if lifespans are bounded then a simple rule on retirements – always retire an admissible chain that contains the agent whose label is minimal among agents in admissible chains – will ensure that the number of rounds that can occur prior to the algorithm's termination is bounded by a comparatively small ordinal, the product ωL. See Appendix A.

[5] Once the existence of a weak-core matching μ^0 is established, it is straightforward to argue directly that the Pareto-optimal weak core is nonempty. Given the matchings $\mu^0, ..., \mu^{i-1}$, let the set of matchings that weakly Pareto improve on μ^{i-1} be denoted $P = \{\phi : \phi(j) \succsim^{j} \mu^{i-1}(j) \text{ for all } j \in N\}$. Observing that agent i must have a favorite from $\{\phi(i) : \phi \in P\}$, set μ^i to be any matching in P such that $\mu^i(i)$ equals this favorite. The matching μ defined by $\mu(j) = \mu^j(j)$ for all $j \in \mathbb{N}$ is then Pareto optimal and if some coalition B could block μ then B could also have blocked μ^0.

Extended top trading summary. Given an ordinal α, U^{α}, D^{α}, and M^{α}_{end} below will indicate, respectively, the sets of unmatched agents and discarded objects that enter round α and the matched agents at the end of round α. Set $U^0 = \mathbb{N}$ and $M^0 = D^0 = \varnothing$ for ordinal 0 (the initial round). Fix an ordinal β and suppose that U^{α}, M^{α}, D^{α}, U^{α}_{end}, M^{α}_{end}, and object $\mu(i)$ for $i \in M^{\alpha}_{\text{end}}$ have been defined for all ordinals $\alpha < \beta$. If $\beta > 0$, set $U^{\beta} = \bigcap_{\alpha<\beta} U^{\alpha}_{\text{end}}$, $M^{\beta} = \bigcup_{\alpha<\beta} M^{\alpha}_{\text{end}}$, and $D^{\beta} = M^{\beta} \backslash \mu(M^{\beta})$. If $\beta \geq 0$, set $S \subset U^{\beta}$ to be a chain that is admissible given U^{β} and D^{β}, for each $s \in S$ set $\mu(s)$ to be the successor of s determined by S, and finally set $U^{\beta}_{\text{end}} = U^{\beta} \setminus S$ and $M^{\beta}_{\text{end}} = M^{\beta} \cup S$. These definitions by transfinite induction imply that U^{α}, D^{α}, and M^{α}_{end} are well-defined for any ordinal α.

4 An implementation difficulty and a competitive solution

As we have seen, a core allocation can exchange objects among infinitely many agents: agent i passes his object to j who passes his object to k ... and so on. If each agent i is born or enters the model no earlier than L periods before i then the agents involved in such an exchange cannot meet to arrange the trade. The same problem can also arise when exchanges form cycles. While an agent i will agree only to trades that give him an object that appears within L periods of i, the cycle of exchanges that contains i might extend more than L periods beyond i. If for example $L = 2$ then the exchanges in the cycle $1 \mapsto 3 \mapsto 5 \mapsto 6 \mapsto 4 \mapsto 2 \mapsto 1$ (where $i \mapsto j$ means i receives object j) would include agents that never live at the same date.

Competitive equilibria can maneuver around this problem. In a market, agents simply sell the objects they are endowed with and buy the objects they most prefer given the prices that they anticipate. Equilibrium obtains when agents satisfy their budget constraints and markets clear. The currency of the market need not be a money that can be traded for other goods; it could be a credit that can be used only for, say, public housing.

A **competitive equilibrium** for $\succ$ consists of a matching μ and a price sequence $p = (p(1), p(2), ...)$ such that each agent can afford the object to which he is matched, any preferred object is unaffordable, and unassigned objects are free: for all $i, j \in \mathbb{N}$, $p(\mu(i)) = p(i)$, if $j \succ^i \mu(i)$ then $p(j) > p(i)$, and if there is no agent $k \in \mathbb{N}$ such that $\mu(k) = j$ then $p(j) = 0$. In interpretation, an agent i who first buys or sells at date j accurately anticipates the prices for the objects that appear later than j, for example selling his endowment at date i in the expectation that the proceeds will pay for object $k > i$.

4.1 Existence of equilibria

Theorem 2 *For any $\succ$ with bounded lifespans, a competitive equilibrium exists for $\succ$.*

A proof in the style of Shapley-Scarf, where the prices of objects descend according to the date at which objects exit from top trading, will not succeed. The non-cycles are to blame: if a discard from an exiting one-sided chain S is later selected by an agent in a different chain then all of the objects in both chains will have to share a common price, and if no one later selects the discard then the endowments of S must be 0-priced, preventing any further descent. Our existence argument will instead piggyback on the good behavior of the finite matching model, where agents and their objects exit only in cycles and in a clear temporal order. For any finite n, apply top trading to the model that consists solely of the first n agents in $\mathbb{N}$; the matching that results can then be supported by prices for objects that decrease with the exit date. Fix some $k \leq n$ and consider the sequence of allocations for the first k agents and the ordering of the prices of the first k objects as n increases. Since the possible allocations and price orderings can assume only finitely many values, there must be a constant subsequence of allocations and price orderings. Restricting attention to this subsequence, we next define a sequence of allocations and price orderings for the first $k+1$ agents and again go to a constant subsequence. The diagonalization that results identifies a matching for all agents and a corresponding price sequence and it is easy to confirm that these form a competitive equilibrium. This argument does not use a uniform bound on lifespans: we could let each agent i have his own L^i. The proof, in Appendix B along with other proofs omitted from the text, bears some similarity to arguments for the existence of equilibria in the overlapping generations model of general equilibrium theory (e.g., Balasko et al. [6]).

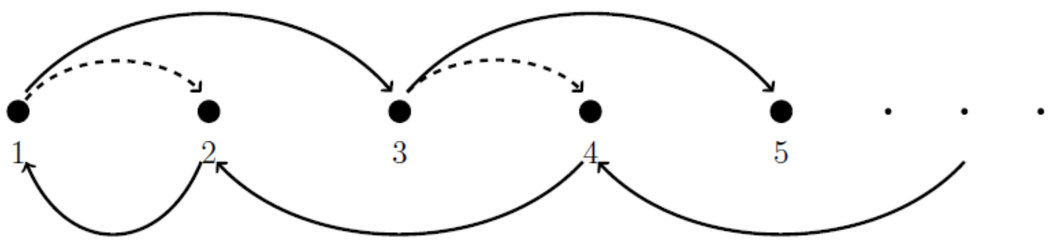

Figure 5: A profile and its limit matching

Examples 5 and 6 below illustrate the equilibria the proof builds. Example 5 underscores that although the trades a specific agent i conducts stabilize as the number of agents that enter into the top trading cycles increases, the competitive equilibrium chain of trades that contains i might appear only in the limit.

Example 5 *The favorite and second favorites of every odd agent i are $i+2$ and $i+1$ respectively, the favorite of each even agent $j \neq 2$ is $j-2$, and the favorite of agent 2 is object 1. The preferences are pictured in Figure 5 where as before a solid arrow points from an agent to his favorite object and a dashed arrow points to the agent's second favorite.*

As in the proof of Theorem 2, run Gale's top trading cycles on the first n agents. When n is even, the matching μ^n that results forms a single cycle and a chain of trades defined by $\mu^n(i) = i+2$ for all odd i except $n-1$, $\mu^n(j) = j-2$ for all even j except 2, $\mu^n(2) = 1$, and $\mu^n(n-1) = n$. When n is odd, agent n forms a singleton cycle and exits and then the first $n-1$ agents form the cycle already described. Though the matching that results is a cycle or a pair of cycles for each n, in the limit the matching is neither a cycle nor set of cycles but the single two-sided chain formed by the solid arrows.

That the equilibria of overlapping-generations economies can be inefficient suggests that the competitive equilibria of the present matching model can be inefficient as well.

Example 6 *We preserve the favorites of the agents in Example 5 and thus the solid arrows in Figure 5 that takes agents to their favorites. But now let each agent i's second-favorite object be i's endowment. The profile of preferences $\succ$ therefore has $i+2 \succ^i i \succ^i j$ for all $j \notin \{i+2, i\}$ when i is odd; $i-2 \succ^i i \succ^i j$ for all $j \notin \{i-2, i\}$ when $i > 2$ is even; and $1 \succ^2 2 \succ^2 j$ for all $j \notin \{1, 2\}$.*

The core of $\succ$ is the solid-arrow matching η that assigns each agent his favorite. Since η strictly Pareto dominates the identity matching ι that gives each agent his endowment, ι is not even in the weak core. But ι can be sustained as a competitive equilibrium. Let p be any price sequence such that $p(j) > p(i)$ when $j = \eta(i)$, for example, by recursively setting $p(1) = 1$ and $p(\eta(i)) = 2p(i)$ for each i. Since no agent i can afford $\eta(i)$, each i will stick with his second-favorite object, his endowment. So (ι, p) is a competitive equilibrium and is in fact the equilibrium built in the proof of Theorem 2.

The price sequence in this last example is not (and could not be) summable, which is why the standard proof of the first welfare theorem cannot be applied.

4.2 Competitive implementation of the core

Our main optimality result is that any matching in the core can be supported as a competitive equilibrium. Given $\succ$, a matching μ can be **competitively implemented** if there is a price sequence p such that (μ, p) is a competitive equilibrium.

Theorem 3 *For any profile $\succ$, if μ is in the core of $\succ$ then μ can be competitively implemented.*

Theorem 3 and Example 6 together show that half and only half of core equivalence obtains: markets can reach the core but may also reach other matchings.

To prove Theorem 3, we begin with a variation on the classical argument for why top trading cycles must generate a core matching. Let μ lie in the core of some profile $\succ$ and suppose extended top trading leads to a distinct matching η. In a finite setting, one would generate a contradiction by considering the first round of top trading cycles at which the departing set of agents S has $\eta(i) \neq \mu(i)$ for some $i \in S$. Since for each $j \in S$ object $\mu(j)$ remains available at this round, and since therefore $\eta(i) \succ^i \mu(i)$, the coalition S could block μ. That exact argument will not work in our setting since S need not be a cycle and some agent in S might therefore point to a discarded object. But a variant will.

Lemma 1 *For any profile $\succ$, if μ lies in the core of $\succ$ then extended top trading must generate μ.*

Proof. Suppose extended top trading generates the matching $\eta \neq \mu$. There is then a first round β such that some agent i exits at β and $\eta(i) \neq \mu(i)$, which implies $\eta(i) \succ^i \mu(i)$. Then for any agent j that exits prior to β, $\eta(j) = \mu(j)$, while for any agent j that exits at round β, object $\mu(j)$ remains available (it is either owned by an unmatched agent or has been discarded) and hence $\eta(j) \succsim^j \mu(j)$. Let S be the set of all agents that exit at or before β. Since η matches each agent in S with the endowment of an agent in S, coalition S can block μ, a contradiction. ■

With this Lemma in hand, a proof of Theorem 3 can roughly follow the Shapley-Scarf strategy of letting the prices of objects decrease with the round at which their owners depart extended top trading. In our setting, prices will decrease as a function of the ordinals assigned to the rounds; agents who depart in later rounds will then be unable to afford the objects that depart in earlier rounds, which they might prefer to their core allocations. But as mentioned in the previous section non-cycles and discards in particular lead to complications: a discarded object cannot have a positive price. The solution is to pre-set a price of 0 for all objects that are discarded or that are assigned to agents endowed with discarded objects, and so on. It turns out that the simplest way to 0-price these objects is to 0-price all objects that do not exit as part of a cycle.

Proof of Theorem 3. Fix an application of extended top trading, which by the Lemma generates μ, and let γ be the round at which the algorithm terminates. Since γ is at most countable there is a strictly decreasing function $f : \gamma \to \mathbb{R}_{++}$. For example, since there is a one-to-one map $g : \gamma \to \mathbb{N}$, we can define f by $f(\alpha) = \left(\sum_{\beta \leq \alpha} 2^{g(\beta)}\right)^{-1}$. If agent i exits the algorithm at some round α as part of a cycle then set $p(i) = f(\alpha)$. If not, set $p(i) = 0$.

To see that (μ, p) forms a competitive equilibrium, suppose to the contrary that there exist $j, s \in \mathbb{N}$ such that $p(s) \leq p(j)$ and $s \succ^j \mu(j)$. If $p(s) > 0$, then the definition of p and $p(j) \geq p(s)$ imply that s is available at the round where j exits the algorithm, which contradicts $s \succ^j \mu(j)$. So suppose $p(s) = 0$. Relabel s as s_0 and let $\overrightarrow{S} = (s_0, s_1, ...)$ be the sequence of agents that satisfies $\mu(s_i) = s_{i+1}$ for $i \geq 0$. Since $p(s_0) = 0$, no initial segment of $\overrightarrow{S}$ can form a cycle of the algorithm. Since μ is one-to-one, any segment of $\overrightarrow{S}$ that omits s_0 also cannot form a cycle. The agents in $\overrightarrow{S}$ therefore do not repeat. If $j = s_n$ for some

$n \geq 0$ then the agents in $(s_0, ..., s_n)$ can block μ, a contradiction since μ is in the core. And if j is not in $\overrightarrow{S}$ then j and the agents in $\overrightarrow{S}$ can block μ.[6] ■

Lemma 1 incidentally shows that the core cannot contain more than one matching: if the core is nonempty then, for any core matching μ, *all* matchings that can be generated by extended top trading must coincide with μ.

Corollary 1 *For any profile $\succ$, the core of $\succ$ contains at most one matching.*

5 Pure-cycle models and full core equivalence

Although core matchings can be implemented in competitive equilibrium, the core may be empty, as we saw in Example 1. And while the weak core is always nonempty, as we saw in Theorem 1, a further pursuit of Example 1 shows that it is possible that neither a matching in the weak core nor a Pareto-optimal matching can be competitively implemented.

Example 7 *Suppose in Example 1 that a weak-core or Pareto-optimal matching can be competitively implemented with prices p. At least one agent must then buy his favorite object and hence there is a least agent j that does so. Since object j is assigned to no agent, $p(j) = 0$. So if $j \geq 3$ then $j - 1$ would buy j. And, since $p(j) = 0$, the object that j buys in equilibrium must also have price 0. So if j is 1 or 2 then $p(3) = 0$ and both 1 and 2 would buy object 3.*

Given Theorem 2, there must be a matching η that can be competitively implemented: if the price sequence is strictly increasing then every agent will consume his endowment, $\eta(i) = i$ for each $i \geq 1$. As we have seen, η is not in the weak core.

Example 7 by no means portrays the general case: for some preference profiles with an empty core, the weak core will be competitively accessible. For instance, suppose agent 1's

[6] At the cost of revising extended top trading, we can prove Theorem 3 without the default assumption that each agent i has a favorite from any set of objects that contains i. Let an agent i that does not have a favorite from some $H \subset \mathbb{N}$ with $\mu(i) \in H$ point instead to $\mu(i)$. Apply this revision to extended top trading and suppose there is a first round α that does not lead to an admissible chain S such that every $i \in S$ is matched with $\mu(i)$. Since each unmatched agent points to an object at round α, some chain S' must form and it will match some $i \in S'$ with an object $j \succ^i \mu(i)$. The same coalition provided in the proof can then block μ.

favorite is object 2 and that the favorites of the remaining agents form a two-sided chain that excludes agent 1. Similarly to Example 1, the core is empty but the prices $p(1) = 0$ and $p(i) = 1$ for $i \geq 2$ will lead to the appealing weak-core allocation where all individuals but agent 1 receive their favorite object.

So, when the core is empty, there may or may not be a satisfactory matching that can be competitively implemented. There is nothing strained about the empty-core cases and they cannot somehow be dismissed, but we can at least place the bad behavior in context. The best-behaved models display **full core equivalence** which will mean that the sets of core allocations and competitively implementable allocations are nonempty and equal. It will come as no surprise that the lapses from this ideal stem from the exit of non-cycles in extended top trading. When only cycles exit the standard argument (given in section 3) for why top trading leads to a core matching will go through, and Theorem 3 implies that this matching can be competitively implemented.

Suppose that extended top trading generates infinitely many blocks of consecutive agents who trade only among themselves and that each block is finite and stretches for at least half a lifespan. These blocks 'clog' the model: an infinite chain cannot form since it would have to include agents in some of these blocks. Extended top trading must then force agents into cycles. We will show that this class of models both displays full core equivalence and is dense in the set of all models: near to any badly behaved model is a well-behaved model.

A profile $\succ$ is **clogged** if there is a $l \in \mathbb{N}$ such that $(\succ^l, \succ^{l+1}, ...)$ has lifespans bounded by some $L > 0$ and, for any matching η generated when extended top trading is applied to $\succ$, there is an infinite set of agents $\mathbb{I} \subset \mathbb{N}$ where for each $i \in \mathbb{I}$ there is a $K \geq L - 1$ such that the restriction of η to $\{i, ..., i + K\}$ defines a bijection. So for instance a profile is clogged if the agents in each $\{i, ..., i + K\}$ prefer objects in this set to all others: since every agent must get an object as good as his endowment, the agents in this set must match with each other.

For a scenario where such a profile might arise, suppose $\succ$ is drawn with some probability law from the state space of preference profiles that have lifespans bounded by $L > 0$. For each agent $i \in \mathbb{N}$ and each preference on $\{i - L, ..., i, ..., i + L\}$, let E_i be the event that i's favorite object is i. If $E_1, E_2, ...$ are independent and there is a positive lower bound for the probability of each E_i then the profile will be clogged with probability 1. Independence assumptions are strong, however, and so this example does not settle the case in favor of clogged profiles.

A profile $\succ$ displays **full core equivalence** if the set of core matchings of $\succ$ is nonempty and equal to the set of competitively implementable matchings.

Theorem 4 *Any clogged profile displays full core equivalence.*

For the question of density, let $\mathcal{P}^i$ be the set of possible preferences for agent i as defined in section 2. Given that the set of objects is discrete, we endow each $\mathcal{P}^i$ with the discrete topology and the **set of profiles** $\prod_{i \in \mathbb{N}} \mathcal{P}^i$ with the product topology.

Theorem 5 *The set of clogged profiles and hence the set of profiles that display full core equivalence are dense in the set of profiles.*

Theorem 5 does not imply that full core equivalence is 'generic.' In fact, any profile $\succ$ including those that display full core convergence can be approximated in the product topology by a sequence of profiles $\succ(n)$ that fail to display full core equivalence. For example, let agent i have the $\succ^i$ specified by $\succ$ for $i = 1, ..., n$ and let agent $n + i$ for $i \geq 1$ prefer his own endowment to any object n or smaller and otherwise have the preferences of agent i in Example 1 translated by n. (That is, for $j, k \geq n$, let $j \succ^{i+n} k$ iff $j - n \succ^i k - n$ for agent i in Example 1.) Although $\succ(n)$ converges to $\succ$, each equilibrium of each $\succ(n)$ is inefficient.

6 Strategyproofness

Gale top trading cycles furnish a simple proof that a strategyproof direct revelation mechanism for finite matching models can implement the core (Roth [17]). With some adjustments, extended top trading can deliver a comparable result, though the proof is no longer as simple.

A **mechanism** is a map that assigns one matching to each profile of preferences. Mechanism M is **strategyproof** if for each profile of preferences $\succ$ and each agent i, the agent weakly prefers the match that results from M when i reports $\succ^i$ to the match that results from M when i reports any other preference, that is, $M(\succ)(i) \succsim^i M(\succ^i_{\text{alt}}, \succ^{-i})(i)$ for any preference relation $\succ^i_{\text{alt}}$ on $\mathbb{N}$.

Extended top trading does not quite define a mechanism since it need not specify a unique matching when multiple admissible chains form at one or more rounds. There are many ways to close this gap; what matters for strategyproofness is that agents cannot manipulate the sequence of retirements to their advantage by pointing to nonfavorite objects. For a set of agents S, let $\min S$ be the minimum-index agent in S according to the original labeling given by $\mathbb{N}$. Define extended top trading to be **constrained** by requiring, for any round α beginning with a nonempty set of unmatched agents U, that: (1) if there exists an admissible cycle at α then retire the admissible cycle S such that $\min S \leq \min S'$ for all cycles S' that are admissible at α and (2) if there does not exist an admissible cycle at α then retire the admissible chain S such that $\min S = \min U$. (Given (1), (2) implies that $\min S$ is a root and object $\min S$ is discarded.) When extended top trading is constrained, the algorithm will generate exactly one matching and therefore defines a mechanism, which we label C.

The crucial features of C are the priority it assigns to cycles and its use of a fixed order of retirement to channel multiple admissible chains through the exit. That the order of retirement tracks the natural ordering of $\mathbb{N}$ is irrelevant; any well-ordering of $\mathbb{N}$ would do.

Theorem 6 *The mechanism C is strategyproof.*

7 Matchings without disposal

One of our introductory illustrations of an open-ended matching problem, the allocation of doctors to emergency rooms, violates free disposal: the hours initially assigned to a doctor, which define his endowment, presumably must be staffed by someone.

When disposal is prohibited, matchings need to map the set of agents onto, not just into, the set of objects. To redefine the core suitably, we say that a coalition $S \subset \mathbb{N}$ blocks the matching μ **without disposal** if there is a bijection $\nu : S \to S$ such that $\nu(i) \succsim^i \mu(i)$ for all $i \in S$ and $\nu(j) \succ^j \mu(j)$ for some $j \in S$. Map μ is in the **no-disposal core** if μ is a bijection on $\mathbb{N}$ and no coalition can block μ without disposal. Strict blocking and the no-disposal weak core are defined accordingly. A **no-disposal equilibrium** (μ, p) is a competitive equilibrium such that μ maps onto $\mathbb{N}$ and we now say that a matching μ can be competitively implemented if there is a p such that (μ, p) is a no-disposal equilibrium.

While our proof of the nonemptiness of the weak core does not apply to the no-disposal weak core, the results on competitive equilibria extend. The matching μ of a competitive equilibrium (μ, p) must now map onto $\mathbb{N}$ but our constructions have satisfied that constraint. Specifically, the equilibria used in the proof of Theorem 2 do not dispose of any objects and hence that result holds for no-disposal equilibria with no adjustments in the proof. Theorems 4 and 5, appropriately reworded to apply to the no-disposal core, also continue to hold, again with no changes in their proofs.

Regarding Theorem 3, any matching μ in the standard core that does not dispose of any object is evidently in the no-disposal core: the set of reallocations that a coalition can potentially use to block shrinks in the no-disposal case and so if no coalition can block with disposal then no coalition can block without disposal. Theorem 3 therefore implies that any such μ can be competitively implemented.

What about arbitrary matchings in the no-disposal core? The following result places no bound on lifespans.

Theorem 7 *For any profile $\succ$, if μ is in the no-disposal core of $\succ$ then μ can be competitively implemented.*

In an equilibrium that leads a no-disposal core matching μ, the objects that agents prefer to their core allocations must be more expensive than their core allocations. The proof shows that if the binary relation R defined by agents pointing to these preferred objects cycled then a coalition that blocks μ could be built from the cycle. Since R must therefore be acyclic, it can be extended to a complete and transitive ordering; prices that represent that ordering can then serve in equilibrium.

A Appendix: slow and quick termination

Proposition 1 *Let α be a countably infinite ordinal. Then there exists a preference profile such that extended top trading* must *terminate at round α and there also exists a profile that satisfies bounded lifespans such that extended top trading* can *terminate at round α.*

Proof. Let f be a bijection from $\mathbb{N}$ to α. Let the preference profile satisfy $j \succ^i i$ if and only if $f(j) < f(i)$, for all $i, j \in \mathbb{N}$. Since agent i can exit only after every j with $f(j) < f(i)$ has exited, for any $\gamma < \alpha$ only the agent k that minimizes $\{f(j) : j$ is unmatched at the beginning of round $\gamma\}$ can exit at round γ. It follows that, for any $\gamma < \alpha$, only agent $f^{-1}(\gamma)$ exits at round γ. The first round that begins with an empty set of unmatched agents is therefore α. For bounded lifespans, let the profile be such that the favorite of each $i \in \mathbb{N}$ is i and let the chain that exits in round $f(j)$ consist of j alone. The first round that begins with an empty set of unmatched agents is again α.[7] ■

That each chain that exits in the above proof consists of one agent who consumes his endowment is inessential. For example, if the agents can be partitioned into a countably infinite set of chains then there will be a bijection from the set of chains to α and we could again let one chain exit at each of α rounds.

The construction used for the 'must terminate' conclusion of Proposition 1 is a little exotic and illustrates how reasonable an assumption bounded lifespans is. Bounded lifespans moreover will open the door to a tweaked algorithm where the bound on the maximal number of rounds prior to termination is small.

Extended top trading is **minimal** if, for any round α at which some chain S exits and any admissible chain S' that forms at α, $\min S \leq \min S'$ (using the original labeling given by $\mathbb{N}$).

If k is a finite ordinal then ωk denotes the ordinal that is order isomorphic to k copies of ω, say $\{0^1, 1^1, 2^1, ...\}, ..., \{0^k, 1^k, 2^k, ...\}$, ordered lexicographically: i^a precedes j^b in the ordering if and only if $a < b$ or ($a = b$ and $i \leq j$).

Proposition 2 *For any preference profile with bounded lifespans, minimal extended top trading must terminate at or before round ωL.*

Proof. Apply minimal extended top trading and let μ be the matching that is generated. Suppose some agents remain unmatched at the beginning of round ωk, where k is finite, and

[7] If α is finite, the 'must terminate at round α' conclusion becomes 'cannot terminate before round α' and the 'can terminate' conclusion stands unaltered. For example, let the favorite of each $i = 1, ..., \alpha - 1$ be i and let the favorite of any $i \geq \alpha$ be $i + 1$. Then each of the first $i - 1$ agents must exit as a singleton, and the remaining agents can exit in any positive finite number of rounds.

that chain S^k exits at round ωk.

Lemma. S^k contains infinitely many agents. *Proof.* If to the contrary S^k is finite then the set of objects that the agents in S^k prefer to their μ-matches, $S^* = \{s \in \mathbb{N} : s \succ^i \mu(i) \text{ for some } i \in S^k\}$, is also finite. For S^k to be admissible at ωk, each object $s \in S^*$ must be assigned by μ to an agent that exits at some round $\gamma_s < \omega k$. Since S^* is finite, $\gamma = \max\{\gamma_s : s \in S^*\}$ is well-defined and $\gamma < \omega k$. Hence S^k is admissible at any round greater than γ. Minimality therefore implies that each chain that exits after γ and before ωk contains an agent with index less than $\min S^k$. Hence S^k must exit on or before round $\gamma + \min S^k$. Since ωk is a limit ordinal, $\gamma + \min S^k < \omega k$, which gives a contradiction. □

To see that, for any $n \geq \min S^k$, S^k must visit the set $\{n+1, \ldots, n+L\}$, suppose instead that $S^k \cap \{n'+1, ..., n'+L\} = \varnothing$ for some $n' \geq \min S^k$. Since $|j - i| > L$ implies $i \succ^i j$, we have $|i - \mu(i)| \leq L$ for each $i \in \mathbb{N}$. Agents drawn from the finite set $S^k \cap \{1, ..., n'\}$ could therefore form an admissible chain at round ωk, contradicting the Lemma.

To conclude, suppose that the algorithm has not terminated prior to beginning of round $\omega(L-1)$, which implies that the infinite chains $S^1, ..., S^{L-1}$ exit at rounds $\omega, \omega 2, ..., \omega(L-1)$. Set $\bar{n} = \max[\min S^1, ..., \min S^{L-1}]$. For any $j \geq \bar{n}$, the set $\{j, j+1, \ldots, j+(L-1)\}$ must contain one agent from S^k for $k = 1, ..., L-1$. Define the infinite set of agents $T = \{i \in \mathbb{N} : i \text{ exits prior to round } \omega \text{ and } i \geq \bar{n}\}$. Then, for any $j \in T$, every agent in $T^j = \{j, j+1, ..., j+(L-1)\}$ has exited by the beginning of round $\omega(L-1)$. Since there are infinitely many distinct T^j sets, no further admissible infinite chains can form. No agent can remain unmatched at the beginning of round ωL: if there were unmatched agents the Lemma would imply that no admissible finite chain could form. ■

The bound provided in Proposition 2 is tight: it is not difficult to build profiles with bounded lifespans where any application of extended top trading, whether minimal or not, cannot terminate before ωL rounds have passed.

B Appendix: remaining proofs

Proof of Theorem 2. Each positive integer n defines a model where the sets of agents and objects are both $\{1, ..., n\}$ and the preference for i is the restriction of $\succ^i$ to $\{1, ..., n\} \times \{1, ..., n\}$. Gale's top trading cycles (TTC's) generates a matching μ^n for this model. As in extended top trading, fix an order in which cycles exit the algorithm by selecting, for each round r, some admissible cycle $S^r \subset \{1,...,n\}$ to remove from the set of unmatched agents that enters round r. Define an 'exits earlier than' binary relation $\succsim^n$ on $\{1,...,n\}$ by $l \succsim^n m$ if and only if there are rounds r and r' such that $l \in S^r$, $m \in S^{r'}$, and $r \leq r'$.

We use the following notation. If $\tilde{\mathbb{N}} \subset \mathbb{N}$ with $\tilde{\mathbb{N}} \supset \{1,...,k\}$ is the domain of a map η and $\succcurlyeq \subset \tilde{\mathbb{N}} \times \tilde{\mathbb{N}}$ is a binary relation, let $\eta_{[k]}$ be the restriction of η to $\{1,...,k\}$ and let $\succcurlyeq_{[k]}$ denote $\succcurlyeq \cap (\{1,...,k\} \times \{1,...,k\})$.

Define the sequence $\langle E^n \rangle$ by $E^n = (\mu^n, \succsim^n)$ and, for $k \leq n$, define $E^n[k] = (\mu^n_{[k]}, \succsim^n_{[k]})$. Since due to bounded lifespans $\left\langle E^k[k], ..., E^n[k], ... \right\rangle$ can assume only finitely many values, this sequence or any of its subsequences must have a constant subsequence.

Let $\left\langle E^{n_1(j)} \right\rangle$ be a subsequence of $\langle E^n \rangle$ (indexed by j) such that $E^{n_1(j)}[1]$ is constant as a function of j. For the diagonalization, suppose for each $i \in \{1,...,k\}$ that we are given a subsequence $\left\langle E^{n_i(j)} \right\rangle$ of $\langle E^n \rangle$ such that $n_i(1) \geq i$ and $E^{n_i(j)}[i]$ is constant. Since $n_i(1) \geq i$, each $E^{n_i(j)}$ assigns objects to and orders of exit of at least agents $1,...,i$. Set $\left\langle E^{n_{k+1}(j)} \right\rangle$ to be a subsequence of $\langle E^{n_k(j)} \rangle$ such that $n_{k+1}(1) \geq k+1$ and $E^{n_{k+1}(j)}[k+1]$ is constant.

Define the matching μ by $\mu(k) = \mu^{n_k(1)}(k)$ for $k \in \mathbb{N}$ and define $\succsim$ on $\mathbb{N}$ by $l \succsim m$ iff $l \succsim^{n_k(1)}_{[k]} m$ for some $k \geq l, m$. Since, for each $k \in \mathbb{N}$, $\succsim^{n_k(1)}_{[k]}$ is complete and transitive on $\{1,...,k\}$ and due to the diagonalization $\succsim_{[k]} = \succsim^{n_k(1)}_{[k]}$, $\succsim$ is complete and transitive. Since in addition $\mathbb{N}$ is countable there exists a $p : \mathbb{N} \to \mathbb{R}_{++}$ such that $p(l) \geq p(m)$ iff $l \succsim m$.

To show that (μ, p) is an equilibrium, we first show that μ is a bijection. Fix some $i \in \mathbb{N}$. Due to bounded lifespans and since $\mu^{n_{i+L}(1)}$ is bijective, there exists a $j \leq i + L$ with $\mu^{n_{i+L}(1)}(j) = i$. Due to the diagonalization, $\mu^{n_t(1)}(j) = i$ for $t \geq j$ and $\mu^{n_j(1)}(t) = \mu^{n_t(1)}(t)$ for $t < j$. Setting $t = j$ shows that μ is onto. Setting $t > j$, the fact that $\mu^{n_t(1)}$ is bijective implies $\mu^{n_t(1)}(t) \neq i$. For $t < j$, the fact that $\mu^{n_j(1)}$ is bijective implies $\mu^{n_t(1)}(t) \neq i$. Hence μ is one-to-one. Since μ is therefore bijective, p satisfies the requirement 'if there is no $k \in \mathbb{N}$ with $\mu(k) = j$ then $p(j) = 0$.' Consider an agent i and a good k with $p(k) \leq p(i)$. Since $i \succsim k$, the diagonalization implies $i \succsim^{n_{\max\{i,k\}}(1)} k$ and $\mu^{n_{\max\{i,k\}}(1)}(i) = \mu(i)$ and thus $\mu(i) \succsim^i k$. ■

Proof of Theorem 4. Assume that $\succ$ is clogged, let η be a matching generated when extended top trading is applied to $\succ$, and let $\mathbb{I}$ be the infinite set of agents and K the parameter then given by cloggedness. Fix some $s_1 \in \mathbb{N}$ and define the sequence $(s_1, s_2, ...)$ where $s_{j+1} = \eta(s_j)$ for all $j \geq 1$. To conclude that this sequence cycles, suppose that it does not. Then there must be a $m \in \mathbb{N}$ such that, for all $n \geq m$, $s_n \geq l$. Let $i \in \mathbb{I}$ satisfy $i > s_m$. Since $(\succ^{s_m}, \succ^{s_m+1}, ...)$ has bounded lifespans and η restricted to $\{i, ..., i+K\}$ is bijective, the sequence $(s_1, s_2, ...)$ is bounded by i and therefore must cycle. We conclude that every admissible chain of η is a cycle. The argument at the start of section 3 or Shapley and Scarf [18] then implies that η is in the core of $\succ$. By Theorem 3, η can be competitively implemented. The set of core matchings is therefore nonempty and contained in the set of competitively implementable matchings.

Conversely, suppose that μ can be competitively implemented using the price sequence p. Suppose that at all rounds of extended top trading prior to r the cycle S' that exits has

$\mu(S') = \eta(S')$. We show that the cycle S that exits at r then has $\mu(S) = \eta(S)$ which implies that $\mu = \eta$. Assign the label s_1 to some element of $\text{argmax}_{s \in S}\, p(s)$ and let $s_{i+1} = \eta(s_i)$ for $i = 1, ..., n-1$, where $n = |S|$. Since S is a cycle, $s_1 = \eta(s_n)$. For induction purposes, assume that $p(s_1) = \cdots = p(s_k)$ where $k < n$. Since $s_1 \in \text{argmax}_{s \in S}\, p(s)$, $p(s_j) \geq p(s_i)$ for $1 \leq j \leq k$ and $1 \leq i \leq n$. Hence, for $1 \leq j \leq k$, agent s_j when facing p can afford s_{j+1}, which is the favorite of s_j from the set of unmatched objects that enter round r. Since by assumption $\mu(s_k) \neq l$ for any l that exits prior to r, $\mu(s_k) = s_{k+1}$ and so, by the definition of a competitive equilibrium, $p(s_{k+1}) = p(s_k)$. The induction therefore gives $p(s_1) = \cdots = p(s_n)$. Hence $\mu(s_i) = s_{i+1} = \eta(s_i)$ for $i = 1, ..., n-1$ and $\mu(s_n) = s_1 = \eta(s_n)$, that is, $\mu(S) = \eta(S)$.

■

Proof of Theorem 5. Let $\mathcal{P}$ be the set of profiles and, for $i \in \mathbb{N}$, let $\succ^{i\prime}$ be a preference relation such that i is the favorite of i. We define

$$\mathcal{F} = \{\succ \in \mathcal{P} : \succ^i = \succ^{i\prime} \text{ for all but finitely many } i \in \mathbb{N}\}.$$

Let $\succ \in \mathcal{F}$. Then there exists a l such that $\succ^i = \succ^{i\prime}$ for all $i \geq l$ and hence any matching μ generated by extended top trading has $\mu(i) = i$ for $i \geq l$. Setting $L = 1$, profile $(\succ^l, \succ^{l+1}, ...)$ has lifespans bounded by L which, with $\mathbb{I} = \{l, l+1, ...\}$ and $K = 1$, implies $\succ$ is clogged.

Since each $\mathcal{P}^i$ is given the discrete topology, any nonempty $\mathcal{O} = \prod_{i \in \mathbb{N}} \mathcal{O}^i \subset \mathcal{P}$ is open in the product topology if and only if $\mathcal{O}^i = \mathcal{P}^i$ for all but finitely many i. So, given a nonempty open $\mathcal{O}$, there is a finite $N \subset \mathbb{N}$ such that $\mathcal{O}^j \subset \mathcal{P}^j$ for $j \in N$ and $\mathcal{O}^k = \mathcal{P}^k$ for $k \in \mathbb{N} \backslash N$. Any $\succ$ with $\succ^j \in \mathcal{O}^j$ for $j \in N$ and $\succ^k = \succ^{k\prime}$ for $k \in \mathbb{N} \backslash N$ is then an element of both $\mathcal{F}$ and $\mathcal{O}$. So $\mathcal{F} \cap \mathcal{O} \neq \varnothing$ for any nonempty open $\mathcal{O}$ and therefore $\mathcal{F}$ is dense in $\mathcal{P}$. Since, by Theorem 4, the set of profiles that display full core equivalence contains the set of clogged profiles which, as shown above, contains $\mathcal{F}$, both the clogged profiles and the profiles that display full core equivalence are also dense in $\mathcal{P}$. ■

Proof of Theorem 6. Suppose by way of contradiction that there is an agent $i \in \mathbb{N}$, a profile $\succ$, and a preference relation $\succ^i_{\text{alt}}$ such that $C(\succ^i_{\text{alt}}, \succ^{-i})(i) \succ^i C(\succ)(i)$. Let the maps A and A_{alt} assign to each ordinal α the admissible chains, the objects discarded, and the sets of matched and unmatched agents that enter and exit at α when the constrained algorithm is applied to $\succ$ and $(\succ^i_{\text{alt}}, \succ^{-i})$ respectively. Let A match i at round β. Since A at round β matches i to his favorite from the set of unmatched objects and $C(\succ^i_{\text{alt}}, \succ^{-i})(i) \succ^i C(\succ)(i)$, A_{alt} must match i at some round $\beta' < \beta$.

We first show that there does not exist an admissible cycle at round β' of A_{alt}. Suppose there is such a cycle. Then condition (1) in the definition of constrained top trading implies that an admissible cycle S exits at round β' of A_{alt}. Recall from section 3 that an admissible chain can be indexed by a set of consecutive integers, say $S = \{s_1, s_2, \ldots, s_n\}$ in the present case, where, for each j, object s_{j+1} is the favorite of s_j from the unmatched objects available at β'. Since S is a cycle and S contains i, we may set $s_n = i$ without loss of generality. So $C(\succ^i_{\text{alt}}, \succ^{-i})(i) = s_1$. Since the exits from A and A_{alt} must be identical at each round prior to β' and s_n does not exit from A prior to β, s_{n-1} cannot exit from A at any γ such that $\beta' \leq \gamma < \beta$. By induction, s_1 also does not exit from A at any γ with $\beta' \leq \gamma < \beta$. But then, since $s_1 \succ^i C(\succ)(i)$ and s_1 is unmatched at β, A cannot match i to $C(\succ)(i)$ at β, a contradiction.

Suppose A has an admissible cycle S at β'. Since A and A_{alt} have the same exits prior to β' and A_{alt} has no admissible cycle at β', S must contain i. Since any two admissible cycles are disjoint at any round, S must be the only admissible cycle of A at β'. So, by condition (1), S must exit A at β', which contradicts the assumption that A matches agent i at round β. So A also does not have an admissible cycle at β'.

Since A and A_{alt} have the same exits prior to round β', A and A_{alt} have the same sets of unmatched agents at β' and hence the same least unmatched agent, say j. Since both A and A_{alt} have no admissible cycles at β', both A and A_{alt} must retire noncycle chains at β' that contain a common root agent j (and so they both discard object j at β'). Since A_{alt} matches agent i at β', the chain retired by A_{alt} at β' contains the ordered set of agents $\{t_1, ..., t_n\}$ where $t_1 = j$, $t_n = i$, and the favorite unmatched object of each t_k with $k < n$ is t_{k+1}. Since $\succ$ and $(\succ^i_{\text{alt}}, \succ^{-i})$ list the same preferences for every agent except i, the chain retired by A at β' also contains $\{t_1, ..., t_n\}$. Hence A matches agent $t_n = i$ at β', a contradiction. ■

Proof of Theorem 7. Recall from section 3 that each chain can be indexed by set of consecutive integers. Given a chain S, let the subscript of $s_i \in S$ denote the index assigned to s_i. We may then represent μ as a disjoint set of chains $\mathcal{S}$ where, for each $S \in \mathcal{S}$ and $s_i \in S$, $\mu(s_i) = s_{i+1} \in S$ and, when $s_i = \max S$, $s_{i+1} = \min S$. Since disposal is prohibited, we may assume that $\mathcal{S}$ contains only cycles and two-sided chains. We first show any $s_i \in S \in \mathcal{S}$ prefers s_{i+1} to all other objects in S. If to the contrary s_i prefers some $s_j \in S$ to s_{i+1} then there is a subset of S that can block μ: the subset consists of s_i, who consumes s_j, and each agent $s_k \in S'$, who continues to consume $\mu(s_k)$, where S' equals $\{s_k \in S : j \leq k < i\}$ when $j < i$ and $\{s_k \in S : k < i \text{ or } k \geq j\}$ when $i < j$. It is therefore compatible with equilibrium to set $p(i) = p(j)$ whenever i and j lie in the same chain S. Call this common price $p(S)$.

We show below that the following binary relation R on $\mathcal{S}$ is acyclic: $S'RS$ if and only if there exists an agent $i \in S$ and an object $j \in S'$ such that $j \succ^i \mu(i)$. The transitive closure of an acyclic R on a set $\mathcal{S}$ can, by Szpilrajn, be extended to a transitive and asymmetric order R^* such that, for all $S, S' \in \mathcal{S}$, $S \neq S'$ implies (SR^*S' or $S'R^*S$), where 'extended' means that SRS' implies SR^*S'. Consequently, as at the end of the proof of Theorem 3, we may set $p(S)$ for $S \in \mathcal{S}$ so that SR^*S' implies $p(S) > p(S')$.

Turning to the acyclicity of R, suppose to the contrary that there is a $\{S^1, ..., S^n\} = \mathcal{S}' \subset \mathcal{S}$ such that $S^1R \cdots RS^nRS^1$. We say that i is *linked to* j if $j \succsim^i \mu(i)$ and that i is *linked via* $T \subset \mathbb{N}$ *to* j if T is a finite ordered set $(r^1, ..., r^t)$, i is linked to r^1, r^k is linked to r^{k+1} for $k = 1, ..., r^{t-1}$, and r^t is linked to j. If S is a cycle then, for each pair $r^i, r^j \in S$, r^i is linked via $(r^{i+1}, r^{i+2}, ..., r^{j-1})$ to r^j. Hence if we suppose that each $S^i \in \mathcal{S}'$ is a cycle then there must be a i in some $S^k \in \mathcal{S}'$, a j in some $S^{k'} \in \mathcal{S}'$ with $j \succ^i \mu(i)$, and a finite ordered set $T \subset \mathbb{N}$ that begins with j such that i is linked via T to i. Since $j \succ^i \mu(i)$, $\{i\} \cup T$ can block μ. Alternatively suppose there is an infinite chain $S \in \mathcal{S}'$, which must be two-sided. Then there must be an infinite chain $S^* \in \mathcal{S}'$ and cycles $S[1], ..., S[t]$ in $\mathcal{S}'$ such that $SRS[1]R \cdots RS[t]RS^*$. Hence there is a $i \in S$, a $k \in S[1]$ with $k \succ^i \mu(i)$, a $j \in S^*$, and some finite ordered set $T \subset \mathbb{N}$ that begins with k such that i is linked via T to j. Let $i \leq_{\bar{S}} j$ mean that, for chain $\bar{S}$, the index assigned to i is less than or equal to the index assigned to j. If $S \neq S^*$ or $i \leq_S j$, the coalition that consists of i, all l in S with $l \leq_S i$, T, j, and all l

in S^* with $j \leq_{S^*} l$ can block μ. If $S = S^*$ and $j <_S i$ then the coalition that consists of j, all $k \in S$ with $j <_S k <_S i$, i, and T can block μ. ■

References

[1] Agarwal N., I. Ashlagi, M. Rees, P. Somaini, and D. Waldinger, 2019, 'An empirical framework for sequential assignment: the allocation of deceased donor kidneys,' *mimeo.*

[2] Akbarpour M., S. Li, and S. Oveis Gharan, 2020, 'Thickness and information in dynamic matching markets' *Journal of Political Economy* forthcoming.

[3] Anderson R., I. Ashlagi, D. Gamarik, and Y. Kanoria, 2017, 'Efficient dynamic barter exchange,' *Operations Research* 65: 1446-1459.

[4] Arnosti N. and P. Shi, 2020, 'Design of lotteries and waitlists for affordable housing allocation,' *Management Science* forthcoming.

[5] Baccara M., S. Lee, and L. Yariv, 2020, 'Optimal dynamic matching,' *Theoretical Economics* 15, 1221-1278.

[6] Balasko Y., D. Cass, and K. Shell, 1980, 'Existence of competitive equilibrium in a general overlapping-generations model,' *Journal of Economic Theory* 23: 307-322.

[7] Bloch F. and D. Cantala, 2017, 'Dynamic assignment of objects to queuing agents', *American Economic Journal: Microeconomics* 9: 88-122.

[8] Delacrétaz D., S. Kominers, and A. Teytelboym, 2020, 'Matching Mechanisms for Refugee Resettlement,' *mimeo.*

[9] Doval L. and B. Szentes, 2019, 'On the efficiency of queueing in dynamic matching markets', *mimeo.*

[10] Fleiner T., 2003, 'A fixed-point approach to stable matchings and some applications,' *Mathematics of Operations Research* 28: 103-126.

[11] Gale D. and L. Shapley, 1962, 'College admissions and the stability of marriage,' *American Mathematical Monthly* 69: 9-14.

[12] Hassidim A., A. Romm, and R. Shorrer, 2021, 'The limits of incentives in economic matching procedures,' *Management Science* 67: 951-963.

[13] Jagadeesana R., 2018, 'Lone wolves in infinite, discrete matching markets,' *Games and Economic Behavior* 108: 275-286.

[14] Leshno J., 2017, 'Dynamic matching in overloaded waiting lists,' *mimeo.*

[15] Li S., 2017, 'Obviously strategyproof mechanisms,' *American Economic Review* 107: 3257-87.

[16] Osborne M. and A. Rubinstein, 1994, *A Course in Game Theory*, The MIT Press: Cambridge, Massachusetts.

[17] Roth A., 1982, ‘Incentive compatibility in a market with indivisibilities,’ *Economics Letters* 9: 127-132.

[18] Shapley L. and H. Scarf, 1974, ‘On Cores and Indivisibility,’ *Journal of Mathematical Economics* 1: 23-37.

[19] Schummer J., 2019, ‘Influencing Waiting Lists’ *mimeo*.

[20] Unver U., 2010, ‘Dynamic kidney exchange,’ *Review of Economic Studies* 77: 372-414.