Matching and equilibrium over an infinite horizon*

Sophie Bade and Michael Mandler†

Royal Holloway College, University of London

This version: May 2021

Abstract

Suppose that agents are to be matched to objects and arrive over time without a definite terminal date. In an optimal matching, the agents linked by chains of trades might have lifespans that fail to intersect, thus obstructing the execution of these trades. To overcome this problem, we let matchings be implemented via competitive markets. Competitive equilibria always exist and any matching in the core can be competitively implemented. The set of core matchings can be empty but a transfinite variant of top trading cycles shows that a Pareto-optimal weak-core matching always exists. Finally if there is minimum positive probability that an agent's favorite object is his endowment then, with probability 1, core allocations exist and all competitive equilibria lie in the core. The full core equivalence of the finite matching model is then achieved.

JEL codes: C71, C78, D50.

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

*Thanks

†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

1 Introduction

In classical matching problems, such as the assignment of kidney donors to patients with end-stage renal disease, the set of objects and the set of agents are both finite. The finite bounds are however modeling abstractions that ignore that in time additional patients will be diagnosed and additional donors will appear; these new arrivals can potentially be matched to some of the original agents. The assignment of workers, such as doctors to hospital emergency rooms, also violates the fixed bounds of a finite model. The ER needs to be staffed in the immediate future but both the doctors and the hospitals know that a new day will come, with more patients that need medical treatment and, due to retirements and new hires, a somewhat different set of available doctors. Generations thus overlap and efficient allocations might require letting a current doctor swap shifts with a not-yet-hired doctor. It is indeed difficult to think of examples of matching where all of the concerned parties could coordinate at a single point in time. While the lifespans of institutions and perhaps even of all relevant agents may not stretch into an endless future, a model with results driven by an exogenously imposed terminal date would be misleading and at odds with how agents perceive the future.[1]

We therefore consider a matching model à la Shapley and Scarf [15] with 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 has a linear order over the set of objects and is assigned the same index i as the object i that the agent initially owns. Interpreting $\mathbb{N}$ as dates, the overlap of generations is embedded in the preferences of agents: if an agent i prefers an object j over his endowment then j must appear within some fixed time span around i. Implicitly agent i is alive only during this time or at least these are the only dates at which i can contract with other agents.[2]

Infinite-horizon matching can therefore present a coordination problem: efficiency may require exchanges among agents that are never alive at the same time and that consequently cannot agree to those exchanges. The solution we propose closely follows the original work

[1] Osborne and Rubinstein [13] 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.

[2] For several results, we will make do with the weaker assumption that an agent facing an arbitrary set of objects always has a favorite.

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 provides the workhorse for matching finitely many agents with objects. Each agent i points to the agent j that owns i's most preferred object. With finitely many agents, the pointings must form at least one cycle and hence, if we retire this cycle of agents by assigning them the object to which they are pointing, repeated rounds of pointing will eventually match all agents and objects. The matching that results lies in the core of the model and, by letting the objects in a cycle share a common price that descends as the cycles retire, the algorithm constructs a competitive equilibrium that implements the matching. Since in addition every competitive equilibrium allocation is also a core matching, full core equivalence obtains.

With infinitely many agents and objects, the standard argument for each of these conclusions breaks down. In addition to cycles, pointing can generate a ray in which one agent i points to a second agent who points to a third, and so on, and no agent in this sequence points to any preceding agent. Pointing can also lead to a two-sided chain where every agent has one predecessor and one successor and there are no repeats. While one can mimic the Gale algorithm by retiring in some round a set of agents that forms a ray or a two-sided chain rather than a cycle, the argument for why the algorithm will terminate no longer applies. Textbook methods for finding a core matching can therefore fail and indeed may be doomed to fail: the core can be empty, in contrast to models of finite matching. In addition, the matchings generated by the competitive equilibria can be inefficient and therefore lie outside the core, again a contrast to the finite model. Competitive matchings can even be strictly Pareto dominated. Given the inefficiency of equilibria in the overlapping generations model of general equilibrium theory, this possibility does not come as a great surprise but it is a problem we will need to address.

Our positive work begins with a change in the rules of the Gale algorithm. First, when the pointings of a set of agents S forms either a finite or infinite chain then S can be retired from the algorithm. When we allow disposal, a chain can have a root agent whose object is discarded and agents in chains can point to previously discarded objects. Second, we allow the algorithm to terminate transfinitely. With these rules, we will always find a matching in the 'weak' core – the set of allocations such that no coalition S can block by using its own endowments to make every agent in S strictly better off.

To see which matchings can be achieved when limited lifespans will prevent agents linked by chains of trades to meet, we show first that our model always has a competitive equilibrium and second that any core matching can be implemented by a competitive equilibrium. Half of core equivalence therefore holds: core matchings can be achieved in equilibrium but equilibrium matchings can lie outside the core. From the glass is half full perspective, inefficient equilibria are possible but there is at least a rich supply of models with efficient equilibria.

Matchings in the weak core unfortunately cannot always be competitively implemented. But if we take a random draw of models and assume there is a minimum positive probability that each agent's favorite object will be his endowment then the standard core will be nonempty with probability 1. So, restricted to generic models, a desirable matching can always be reached. The same restriction rules out competitive equilibrium matchings that lie outside the core. Full core equivalence is therefore restored.

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 in any example where workers are assigned to shifts that must be filled.

As Gale's top trading cycles are among the most elementary and appealing arguments in economic theory, we are pleased to report that most of the arguments in this paper descend from Gale. The only notable exceptions are the Cantor diagonalization used to prove the existence of competitive equilibria and an acyclic pointing mechanism that we use to construct competitive equilibria when a Gale-style mechanism of sequential exit is unavailable (Theorem 6).

We do however aim to shed light on what drives matching arguments that rely on Gale's top trading cycles. The finiteness of the number of agents would appear to be the reason why a top trading cycles algorithm terminates and why it leads to core allocations. By allowing algorithms to terminate transfinitely and weakening the definition of the core, we will see that finiteness is not essential for either conclusion.

2 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 [17], 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 [10], Bloch and Cantala [7], Schummer [16], Arnosti and Shi [4] and Agarwal et al. [1] study unilateral matching problems over an infinite horizon without initial endowments, for example, the allocation of public housing or the kidneys of deceased donors. They show that the optimal organization of waiting lists for these objects depends on the heterogeneity of the agents' preferences.

Bilateral matching markets over an infinite horizon have also begun to draw attention. 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 [8] have analyzed a bilateral matching market of stochastically arriving impatient agents with dichotomous preferences.

3 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.[3] A profile of all agents' preferences is denoted $\succ = (\succ^i)_{i \in \mathbb{N}}$.

Let the **favorite** of agent i from a set $H \subset \mathbb{N}$ be the object $j \in H$ such that $j \succsim^i k$ for $k \in H$, let the **second favorite** be the object $s \in H$ such that $s \succsim^i k$ for $k \in H \setminus \{j\}$, and so forth. When not specified explicitly, the reader should assume that $H = \mathbb{N}$. Our default assumption will be that each agent i has a favorite from any $H \subset \mathbb{N}$ with $i \in H$. At times we impose the somewhat 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, we will say that **lifespans are bounded** if there exists a $L > 0$ such that, for each agent i and object j, $|i - j| > L$ implies $i \succ^i j$. Since agents can always consume their endowment, their rankings of the objects that appear beyond their lifetimes are irrelevant; consequently when lifespans are bounded we could let agents be indifferent among these goods.

Our assumptions permit a loose interpretation of how agents and objects are associated with dates. For example, $1, ..., l$ can designate agents and objects that appear at calendar date one, $l + 1, ..., 2l$ can designate agents and objects that appear at calendar date two, etc.

A **matching** or **allocation** is a map μ from the set of agents $\mathbb{N}$ to set of objects $\mathbb{N}$ such that 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. The image of μ therefore need not equal $\mathbb{N}$.

Define $\nu : S \to S$ to be a **submatching** if $S \subset \mathbb{N}$ and ν is one-to-one. The coalition $S \subset \mathbb{N}$ **blocks** matching μ at $\succ$ if there is a submatching ν 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 submatching 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] 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.

4 Top trading cycles

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 S^1 must form in the first round and the agents in one of these cycles 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^i be the first cycle to retire with an agent i^* who does better with the object j^* that B assigns to i^* than with $\mu(i^*)$. Since $\mu(i^*)$ is the favorite of i^* from the objects owned by the agents who retire at S^i or later, agent j^* must be in a cycle S^j with $j < i$. Since j^* is therefore in B and B must assign $\mu(j^*)$ to j^*, the owner of $\mu(j^*)$ must also be in B. Iterating this argument, B must contain all of S^j. But then B assigns object j^* to both i^* and some agent in S^j, a contradiction.

What happens if we apply classical top trading cycles to our infinite setting? The first and most obvious fact is that no cycles might form.

Example 1 *Suppose for each agent i besides 1 that i's favorite object is $i + 1$ and that agent 1's favorite object is 3. Letting a solid arrow point from each agent to the agent's favorite object, the preference profile is pictured in Figure 1. In this example and all the examples to follow, we can set agent i's remaining preference rankings to be consistent with the bounded lifespan assumption, for instance by letting 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 $(i_1, i_2, ...)$ such that, for all $j \geq 1$, the favorite of i_j is i_{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 the ray from the algorithm. A single agent, either 1 or 2, would remain and so the second round would assign this agent his endowment.

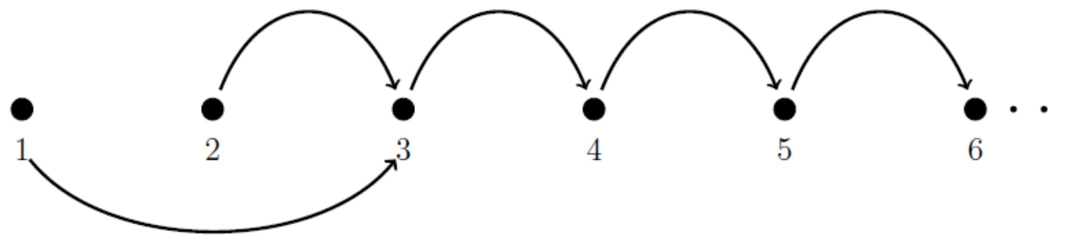

Figure 1: An empty core

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. If in Example 1, the first round assigns objects $(3, 4, 5, ...)$ to agents $(2, 3, 4, ...)$ then a subset of $\{2, 3, 4, ...\}$ can form part of a blocking coalition: $\{1, 3, 4, ...\}$ 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 heavily on agents who agree to block allocations even when they gain nothing by doing so. The weak core can therefore identify more sharply which allocations will survive unchallenged. But since 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.

A seemingly harmless feature of top trading cycles stands in the way of Pareto optimality: when a subset of agents retires from the algorithm it leaves with its endowments. In a finite model, this property follows from the fact that the objects the retiring subset consumes must coincide with the objects it is endowed with. But in an infinite model, a ray that retires from the algorithm can depart with an object that no agent consumes. Inefficiency can therefore result even when a core allocation exists, as the following example illustrates.

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 a solid arrow continues to point from agents to their favorite objects and a dashed arrow points from*

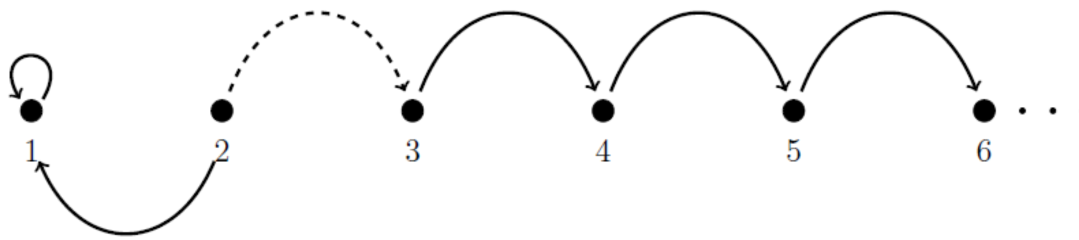

Figure 2: Inefficiency generated by top trading cycles

an agent to his second favorite. In round 1*, the ray that matches each agent* $i \geq 3$ *with object* $i+1$ *forms. If we remove the objects owned by 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 could achieve both a Pareto improvement and a core matching.*

To show via top trading cycles that Pareto-optimal weak-core allocations exist will therefore require some modifications. We will follow Gale in several respects: groups of agents will retire in sequence with their favorite currently available objects as their assignments and retirees are not allowed to seize objects from agents who are not currently exiting. Three changes will be necessary: more groups in addition to cycles will be allowed to retire from the algorithm, current retirees can point to objects that previous retirees have discarded, and the algorithm can terminate after transfinitely many rounds.

Let a **chain** S be a subset of $\mathbb{N}$ indexed by a set of consecutive integers. A chain S can have at most one element with a maximal index, $\max S$, and at most one element with a minimal index, $\min S$, but is not required to have either. Chains may therefore be infinite.

Each round of the **modified top trading cycles** algorithm begins with a set of unassigned agents $N \subset \mathbb{N}$ and a (possibly empty) set of discarded objects D. Let $s \rightarrow_{N,D} t$ mean that $s \in N$ and t is the favorite object of agent s from $N \cup D$. Given N and D, a chain S is **admissible** if $S \subset N$ and, for all $s \in S$,

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

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

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. If $\max S \rightarrow_{N,D} \min S$ 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 can supply the indices, or a 'reverse' ray, where the negative integers can 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 N and an arbitrary D, an admissible chain will always exist. Beginning with some $i \in N$, suppose that $i \rightarrow_{N,D} j \rightarrow_{N,D} \ldots$. This sequence will eventually repeat, end finitely by reaching an element of D, or form a set of infinitely many distinct agents. Each case leads to an admissible chain. In the first, we can extract from $i \rightarrow_{N,D} j \dot{\rightarrow}_{N,D} \ldots$ a sequence that begins and ends with the same element of N and has no other repetitions. An admissible cycle is thus defined. In the second, the penultimate entry in the sequence forms a singleton admissible chain. The third case defines an admissible ray. Given a nonempty N and arbitrary D, modified top trading cycles will select one of the admissible chains S, match each $i \in S$ with i's successor in S, and remove S from N. Each round thus defines a set of **survivors** $N \backslash S$ which will be strictly contained by N when N is nonempty. The survivors provide the unassigned agents for the next round. If there is a $\max S$ and $\max S \rightarrow_{N,D} d$ where $d \in D$ then we remove d from D and if S is finite but not a cycle (so $\max S \nrightarrow_{N,D} \min S$) then we add $\min S$ to $D \backslash \{d\}$. This addition and/or removal fixes the set of discarded objects for the next round.

Modified top trading cycles begins with $N = \mathbb{N}$ and $D = \varnothing$ and terminates once the set of survivors is empty. Before addressing whether a matching that results from this algorithm is Pareto optimal or lies in the weak core, we must consider how and when the algorithm terminates. There is of course no reason for modified top trading cycles to terminate after finitely many rounds. But this slowness should not be attributed to the requirement that only one chain can exit in each round of the algorithm. If we were to let as many chains as possible exit simultaneously, the following example shows that the algorithm could still leave infinitely many agents unassigned to chains at the end of each date.

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 modified top trading*

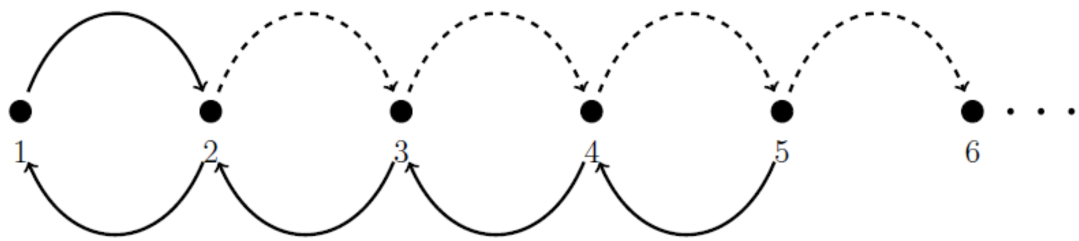

Figure 3: Top trading cycles that do not finitely terminate.

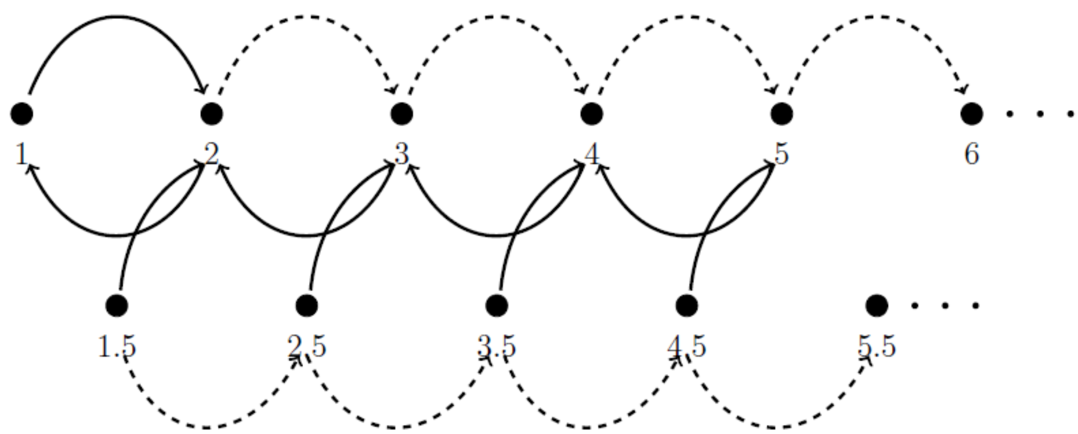

Figure 4: Top trading cycles that terminate transfinitely

cycles, 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 unassigned.*

In Example 3, each agent i is at least assigned to a chain after finitely many rounds. Somewhat less obviously, modified top trading cycles need not reach an empty set of survivors following the completion of every finite round.

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 agents remains unmatched, all the new agents remain unmatched. Since it takes all of the finite rounds to match the original agents, none of the new agents is matched until the algorithm reaches a transfinite round.*

The final change to the Gale algorithm will therefore be that in modified top trading cycle the number of rounds can equal an arbitrary ordinal number.[4] Letting β be an ordinal, the

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

set of survivors N^β of round β is defined recursively. Given N^α for each $\alpha < \beta$, the set of unassigned agents that enters round β will equal $\bigcap_{\alpha<\beta} N^\alpha$. If this set is not empty then some admissible chain S forms and we set $N^\beta = \left(\bigcap_{\alpha<\beta} N^\alpha\right) \setminus S$. When β is a successor ordinal and β therefore has an immediate predecessor $\beta - 1$, the set of unassigned agents that enters round β reduces to $N^{\beta-1}$. But when β is a limit ordinal – for example ω the first ordinal that succeeds the finite numbers – we must use the definition provided. So on the first round ω that follows the rounds associated with the finite numbers, the set of unassigned agents that enters is given by $\bigcap_{\alpha<\omega} N^\alpha = \bigcap_{i=0}^\infty N^i$. In Example 4, the 'new' agents form the set of unassigned agents that enters round ω.

The set of assigned agents at the end of round β is $N_a^\beta = \mathbb{N} \setminus N^\beta$. Let μ^β indicate the algorithm's assignments for these agents: for $s \in N_a^\beta$, let $\mu^\beta(s)$ equal s's immediate successor in the admissible chain S selected by the algorithm that has $s \in S$. The set of discarded objects at the end of round β is therefore $D^\beta = N_a^\beta \setminus \mu^\beta(N_a^\beta)$.

Comparably to Gale, modified top trading cycles **terminates** at the first round α such that $N^\alpha = \varnothing$ and the matching **generated** by the algorithm is μ^α. To see that termination is guaranteed when the rounds extend to arbitrary ordinals, let the algorithm proceed through ω_1 rounds, where ω_1 is the first uncountable ordinal. Since each round that begins with a nonempty set of unassigned agents N eliminates at least one agent from N and since the set of predecessors of ω_1 is uncountable, the set of survivors must be empty following some round β with $\beta < \omega_1$. Keep in mind that we are not constructing a procedure that agents will follow in real time: we are showing only that the weak core is nonempty.

To confirm that any matching μ generated by modified top trading cycles is in the weak core, let $N^* \subset \mathbb{N}$ be an arbitrary coalition. Since rounds are assigned only to ordinals, the set of rounds α such that N^α does not contain N^* is well-ordered, that is, it has a minimal element β. So, for any round $\gamma < \beta$, $N^\gamma \supset N^*$ and β is the first round of the algorithm where there is a $i \in N^*$ that is not also in N^β. Let N and D be the sets of unassigned agents and discarded objects that enter round β. Since $\mu(i)$ is the favorite of i from $N \cup D$ and $N^* \subset N$, i cannot strictly prefer any object $j \in N^*$ to $\mu(i)$. Hence N^* cannot strictly block μ.

For Pareto optimality, suppose some matching η Pareto improves on μ. There must then be a first chain S that exits the algorithm that contains an agent i with $\eta(i) \neq \mu(i)$ and thus $\eta(i) \succ^i \mu(i)$. But, since every agent j who exits prior to S receives $\mu(j) = \eta(j)$, $\eta(i)$ must then be available when S forms. Hence the algorithm would not be assigning every agent in S his favorite available object, a contradiction.

We have therefore proved:

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

That the proof above invokes the first uncountable ordinal ω_1, a notoriously difficult concept to grasp, is no accident. As we show in the Appendix, in the absence of restrictions on the order of retirement when more than one admissible chain forms, the procession of retirements can last as long as any countable ordinal; ω_1 is therefore the least upper bound on the algorithm's termination date. While countable, the termination dates can be large countable ordinals: they therefore need not be computable and are also difficult to grasp. But with a simple rule on retirements – always retire first the admissible chain with the agent that is assigned the smallest natural number – the number of rounds that can occur prior to the algorithm's termination is bounded by a comparatively small ordinal, the product ωL. See the Appendix.

5 An implementation difficulty and the competitive solution

As Example 1 illustrates, a core allocation can involve exchanges of 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 entire set of agents involved in such an exchange cannot meet to arrange the trade. Exchanges that form cycles can present the same problem. 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 $1 \mapsto 3 \mapsto 5 \mapsto 4 \mapsto 2 \mapsto 1$ (where $i \mapsto j$ means that i receives object j) would involve agents that never live at the same date.

Competitive equilibria can maneuver around this implementation 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.

A **competitive equilibrium** for $\succ$ consists of a matching μ and a price sequence $p : \mathbb{N} \rightarrow \mathbb{R}_+$ 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, e.g., i sells his endowed object at date i in the expectation that the proceeds will pay for object $k > i$.

5.1 Existence of equilibria

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

The proof 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]). For any finite n, the matching model that consists solely of the first n agents in $\mathbb{N}$ has a competitive equilibrium that can be found via Gale top trading cycles. 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 goods 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. Using a Cantor diagonalization, we can build a matching for all agents and a corresponding price sequence and it is easy to confirm that these form a competitive equilibrium. The proof below does not use a uniform bound on lifespans: we could let each agent i have his own L^i.

Proof of Theorem 2. For each positive integer n, we can apply Gale's top trading cycles

(TTC's) to the model that consists of agents and objects $1, ..., n$ where we restrict each $\succ^i$ to $\{1, ..., n\}$, thus generating a matching μ^n. Fix an order in which cycles exit the algorithm by choosing arbitrary order of exit, one cycle per round, if more than one cycle forms at some stage. For round r, let $S^r \subset \{1, ..., n\}$ be the cycle that is removed from the set of unassigned agents that enters round r. A competitive equilibrium that supports μ^n is then defined by any price vector $p^n = (p^n(1), ..., p^n(n)) \gg 0$ such that $p^n(i) > p^n(j)$ if and only if there are rounds r and r' such that $i \in S^r$, $j \in S^{r'}$, and $r' > r$. While there are many such price vectors, each represents the same ordering $\precsim^n$ on $\{1, ..., n\}$.

Let E^n denote $(\mu^n, \precsim^n)$. For any n and positive integer $k \leq n$, let $E^n[k]$ denote $(\mu^n[k], \precsim^n [k])$ where $\mu^n[k] = (\mu^n(1), ..., \mu^n(k))$ and $\precsim^n[k]$ is the subrelation of $\precsim^n$ defined by restricting $\precsim^n$ to $\{1, ..., k\}$. Since the sequence $E[k] = \langle E^1[k], ..., E^n[k], ...\rangle$ can assume only finitely many values, $E[k]$ or any of its subsequences must have a constant subsequence.

Let $\langle E^{1_1}, ..., E^{1_n}, ...\rangle$ be a subsequence of $\langle E^1, ..., E^n, ...\rangle$ such that $\langle E^{1_1}[1], ..., E^{1_n}[1], ...\rangle$ is constant and, for each $k \geq 1$, let $\left\langle E^{(k+1)_1}, ..., E^{(k+1)_n}, ...\right\rangle$ be a subsequence of $\langle E^{k_1}$, ..., $E^{k_n}, ...\rangle$ such that $\langle E^{(k+1)_1}[k+1]$, ..., $E^{(k+1)_n}[k+1], ...\rangle$ is constant. Let $\langle E\langle 1\rangle, ..., E\langle k\rangle, ...\rangle$ denote the Cantor diagonalization sequence: $E\langle k\rangle = E^{k_k}$ for each k. Let $\mu_{\langle k\rangle}$ denote the first k entries of the matching in $E\langle k\rangle$ and let $\precsim_{\langle k\rangle}$ denote the restriction of the price ordering in $E\langle k\rangle$ to $\{1, ..., k\}$. For all $j \geq k$, $\langle E^{j_1}[k]$, ..., $E^{j_n}[k], ...\rangle$ is the same constant sequence and so, for $j \geq k$, the first k entries in $\mu_{\langle j\rangle}$ equal $\mu_{\langle k\rangle}$ and the restriction of $\precsim_{\langle j\rangle}$ to $\{1, ..., k\}$ equals $\precsim_{\langle k\rangle}$.

Define the matching μ by setting, for each k, the first k entries of μ to equal $\mu_{\langle k\rangle}$. We set the price sequence p recursively. First let $p^1 = 1$. Then, given p^k, let $p^{k+1} = (p^k(1), ..., p^k(k), p^{k+1}(k+1))$ where $p^{k+1}(k+1)$ equals $p^k(j)$ if $k+1 \approx_{\langle k+1\rangle} j$ and otherwise set $p^{k+1}(k+1)$ to lie between $\min\{p^k(r) : r \succ_{\langle k+1\rangle} k+1\}$ and $\max\{p^k(r) : r \prec_{\langle k+1\rangle} k+1\}$. Define p by setting its first k entries to equal p^k.

To see that (μ, p) is an equilibrium, observe first that since $\mu_{\langle i+L\rangle}$ is determined by TTC's and since only agents born between $i - L$ and $i + L$ can have i as an immediate successor in a TTC chain, $\mu_{\langle i+L\rangle}$ maps exactly one agent in $\{i - L, ..., i + L\}$ to i. Again because each $\mu_{\langle j\rangle}$ is determined by TTC's and since no agent born after $i + L$ can have i as an immediate successor in a TTC chain, no $\mu_{\langle j\rangle}$ with $j > i + L$ maps a different agent to i. The matching

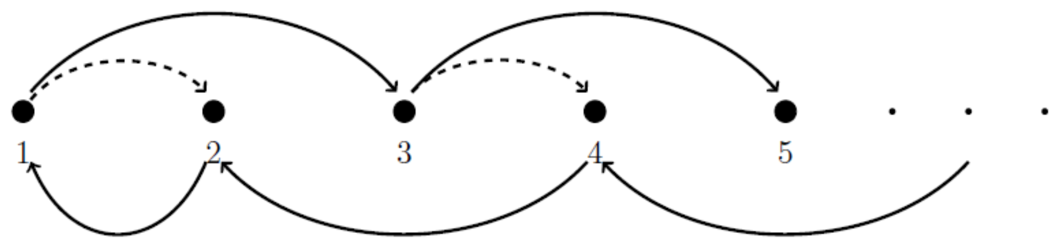

Figure 5: A profile and its limit matching

μ therefore matches each object to exactly one agent (with no disposal). Second, since the prices $p(1), ..., p(i + L)$ represent $\succsim_{\langle i+L \rangle}$ and $\succsim_{\langle i+L \rangle}$ represents an equilibrium that supports μ^{i+L}, object i is the optimizing choice for the agent that $\mu_{\langle i+L \rangle}$ maps to i when the agent faces p. ■

Examples 5 and 6 below illustrate the equilibria this proof builds. A notable feature of the constructions is 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 not stabilize: it might appear only in the limit. Example 5 underscores this point.

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 favorites 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.*

We apply the proof of Theorem 2 by running Gale's top trading cycles on the first n agents. When n is even, the matching μ generated has $\mu(i) = i + 2$ for all odd i except $n - 1$, $\mu(j) = j - 2$ for all even j except 2, $\mu(2) = 1$, and $\mu(n - 1) = n$. The chain of trades thus forms a single cycle. When n is odd, agent n forms a singleton cycle and then the first $n - 1$ agents form a cycle. 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.

Given that the equilibria of overlapping-generations economies can be inefficient, it is no great surprise that the competitive equilibria of the present matching model can lead to inefficient allocations 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 takes each agent to 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, the p defined recursively by setting $p(1) = 1$ and $p(\eta(i)) = 2p(i)$ for each i. Then no agent i can afford $\eta(i)$ and will instead stick with his second-favorite object, his endowment. So (ι, p) is a competitive equilibrium and is as well the equilibrium built in the proof of Theorem 2.

5.2 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 classical argument for why the modified top trading cycles algorithm must generate a core matching. Let μ lie in the core of some profile $\succ$ and suppose modified top trading cycles 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 S has $\eta(i) \neq \mu(i)$ for some $i \in S$. Since each $j \in S$ can still point to $\mu(j)$ at this round, and since therefore $\eta(i) \succ^i \mu(i)$, the coalition S could block μ. To adapt this argument to our setting, we need to deal with the wrinkle that some agent in S might point to a discarded object; we must therefore augment the blocking coalition to include the owner t of the discarded object, the owner of $\eta(t)$, and so on.

Lemma 1 *If μ lies in the core of some profile $\succ$ then modified top trading cycles must generate μ.*

Proof. Suppose there is a first round β of modified top trading cycles such that the admissible chain S^β that exits at β contains an agent s that is not assigned his core matching. So, letting N and D denote respectively the sets of unassigned agents and discarded objects that enter round β, $s \nrightarrow_{N,D} \mu(s)$. We can then find a coalition that can block μ. Before doing so, given some agent t_1, define $\overrightarrow{\mu}(t_1)$ to be the infinite sequence of agents $(t_1, t_2, ...)$ that satisfies $\mu(t_j) = t_{j+1}$ for $j \geq 1$. When no agent j in $\overrightarrow{\mu}(t_1)$ has $\mu(j) = t_1$, the agents in $\overrightarrow{\mu}(t_1)$ must be distinct since μ is one-to-one. Object t_1 can therefore be assigned to an agent outside $\overrightarrow{\mu}(t_1)$ if the agents in $\overrightarrow{\mu}(t_1)$ consume their core matchings. The blocking coalition B will consist of the agents in S^β and, if the favorite of agent $\max S^\beta$ from $N \cup D$ is some $t_1 \in D$, the agents that form $\overrightarrow{\mu}(t_1)$. If there is such a t_1 then all of the agents in $\overrightarrow{\mu}(t_1)$ retired from the algorithm prior to β. Since in addition β is the first round where the algorithm does not assign an agent to his core matching, no agent j in $\overrightarrow{\mu}(t_1)$ has $\mu(j) = t_1$. To confirm that B can block μ, let ν be the submatching that assigns each $i \in S^\beta$ his favorite from $N \cup D$ and, when there is a t_1, leaves each t_i in $\overrightarrow{\mu}(t_1)$ with $\mu(t_i)$. Since by assumption each object k allocated prior to β equals $\mu(j)$ for some j that retired prior to β, $\mu(i) \in N \cup D$ for all $i \in S^\beta$. Moreover, by assumption s has a favorite from $N \cup D$ that differs from $\mu(s)$. Hence B can block μ using ν, a contradiction. ■

To apply this Lemma to Theorem 3, we can roughly follow the Shapley and Scarf strategy of letting the prices of objects be determined by the round at which their owners depart modified top trading cycles. With prices that decrease as a function of the ordinals assigned to the rounds, agents who depart in later rounds will be unable to afford the objects that depart in earlier rounds, which they might prefer to their core allocations. But a wrinkle appears here too: we cannot assign a positive price to any discarded object. The solution is to pre-set a price of 0 for all objects that are discarded or that are assigned to agents that own discarded objects, and so on.

Proof of Theorem 3. Let S^α be a chain formed by modified top trading cycles at some round α. By the Lemma, for each $i \in S^\alpha$ the allocation of i is $\mu(i)$. If S^α is not a cycle then S^α is either infinite or $\mu(\max S^\alpha)$ was discarded in some round that precedes α and then $\overrightarrow{\mu}(\mu(\max S^\alpha))$ (defined in the proof of the Lemma) must consist of infinitely many distinct agents. With either possibility, no agent i that in a chain S^β with $\beta \neq \alpha$ can prefer any $s \in S^\alpha$ to $\mu(i)$ since then the agents in $\overrightarrow{\mu}(i)$ could block by assigning themselves the objects in $\overrightarrow{\mu}(s)$. It is therefore compatible with equilibrium to set $p(j) = 0$ for all $j \in S^\alpha$. Let $\mathcal{C}$ be the set of cycles generated by the algorithm. We need a price $p(S)$ for each $S \in \mathcal{C}$ which we assign to each $j \in S$ such that whenever some agent i in some $S \in \mathcal{C}$ prefers an object in $S' \in \mathcal{C}$ to $\mu(i)$ then $p(S') > p(S)$. Since each i prefers $\mu(i)$ to objects in all cycles that retire after S and all noncycle chains, it is sufficient for prices to decrease as a function of the ordinals assigned to the rounds. Recall that since the set of agents is countable the set of survivors of modified top trading cycles must be empty following some round γ where γ is at most countable. Hence there is a one-to-one map $f : \gamma \to \mathbb{N}$. The following prices will then do: for each $S^\alpha \in \mathcal{S}$, set $p(S^\alpha) = \left(\sum_{\beta \leq \alpha} 2^{f(\beta)}\right)^{-1}$.[5] ■

The Lemma 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 modified top trading cycles must coincide with μ.

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

6 Almost sure competitive implementation

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

[5] At the cost of revising modified top trading cycles, 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 modified top trading cycles and suppose there is a first round α that does not lead to an admissible chain S such that every $i \in S$ is assigned $\mu(i)$. Since each unassigned agent points to some object at round α, some chain S' must form and it will assign some $i \in S'$ an object $j \succ^i \mu(i)$. The same coalition B provided in the proof can then block μ.

in Theorem 3, the following example shows that it is possible that neither any matching in the weak core nor any Pareto-optimal matching can be competitively implemented.

Example 7 *We return to Example 1. In any weak core matching μ, at least one agent $i \geq 3$ must receive his favorite object, $\mu(i) = i + 1$, since otherwise agents $3, 4, 5, \ldots$ would strictly block. Any Pareto optimal matching shares the same feature. Letting i^* be the least $i \geq 3$ that gets his favorite object, we show that there is no supporting competitive equilibrium. If $i^* \geq 4$ then $p(i^*) = 0$ since no agent buys i^*. But $p(i^*) = 0$ would allow $i^* - 1$ to afford i^*. So suppose instead that $i^* = 3$. If $p(3) = 0$ then both 1 and 2 would buy 3. If $p(3) > 0$ then either $j = 1$ or $j = 2$ must buy 3. But then $p(j) = p(3) > 0$ to ensure that j can afford 3 while $p(j) = 0$ since no one buys j.*

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

To avoid the unsatisfactory outcomes of Example 7 and kindred cases, we consider a mild domain restriction on the preferences of agents that ensures a nonempty core and hence that some Pareto-optimal matching can be competitively implemented.

Let $\succ$ be drawn with probability law $\mathbb{P}$ from a state space Ω that equals the set of preference profiles that satisfy our assumptions and the bound L on lifespans. For each agent $i \in \mathbb{N}$ and each linear preference $\succ_r^i$ on $\{i - L, \ldots, i, \ldots, i + L\}$, we assume that the states such that $\succ^i$, when restricted to pairs in i's lifespan, equals $\succ_r^i$ forms a measurable event.

For each $i \in \mathbb{N}$, let $E_i = \{\succ \in \Omega : i \succ^i j \text{ for all } j \in \mathbb{N}\}$ be the event that i's favorite object is i. We assume that $E_1, E_2, \ldots$ are independent and that there is a positive lower bound for the probability of each E_i, that is, a $b > 0$ such that $\mathbb{P}(E_i) \geq b$ for each i.

Theorem 4 *With probability 1, the core of $\succ$ is nonempty.*

Proof. For any $k \in \mathbb{N}$, the event $F_k = E_k \cap \cdots \cap E_{k+L}$ has strictly positive probability and therefore the event F such that $F \subset F_k$ holds for infinitely many values of k has probability 1. For $\succ$ in F and any agent j_1, the sequence $(j_1, \ldots, j_n)$ such that j_{i+1} is the favorite object of j_i must include a cycle for n sufficiently large. The Gale top trading cycles algorithm must therefore partition $\mathbb{N}$ into cycles, thus identifying an element of the core of $\succ$. ■

Corollary 2 *With probability* 1*, there exists a Pareto-optimal matching for* $\succ$ *that can be competitively implemented.*

Theorem 5 *With probability* 1*, every matching for* $\succ$ *that can be competitively implemented lies in the core of* $\succ$*.*

Proof. Assume that $\succ$ lies in the event F defined in the proof of Theorem 4. Let (μ, p) be a competitive equilibrium and suppose S can block μ using some submatching ν. Then, for some $i^* \in S$, $\nu(i^*) \succ^{i^*} \mu(i^*)$. Since $\succ$ is in F, there is a $k^* > i^*$ such that $F \subset F_{k^*}$ holds. Define $N^{k^*} = S \cap \{1, ..., k^* - 1\}$. Since the assumption of bounded lifespans implies that ν is a bijection on N^{k^*}, we can apply a variant of the classical Shapley argument that the existence of a blocking coalition leads to two inconsistent requirements: $\sum_{j \in N^{k^*}} p(\nu(j))\nu(j) > \sum_{j \in N^{k^*}} p(j)j$ (since $p(\nu(j))\nu(j) \geq p(j)j$ for all $j \in N^{k^*}$ and $p(\nu(i^*))\nu(i^*) > p(i^*)i^*$) and $\sum_{j \in N^{k^*}} p(\nu(j))\nu(j) = \sum_{j \in N^{k^*}} p(j)j$ (since ν is a bijection on N^{k^*}). ■

Theorems 3 and 5 together show that full core equivalence obtains with probability 1. While we gave a direct proof of Theorem 4, the result also follows from Theorems 2 and 5: a competitive equilibrium necessarily exists and, with probability 1, any competitive equilibrium lies in the core.

7 Matchings without disposal

One of our introductory illustrations of an infinite-horizon matching problem, the allocation of doctors to emergency rooms, violates free disposal: the ER has a pre-set staffing need at all hours.

To define the core when disposal is prohibited, we say that a coalition $S \subset \mathbb{N}$ blocks the matching μ **without disposal** if there is a submatching $\nu : S \to S$ that maps onto S such that $\nu(i) \succsim^i \mu(i)$ for all $i \in S$ and $\nu(j) \succ^j \mu(j)$ for some $j \in S$. Matching μ is in the **no-disposal core** when μ maps onto $\mathbb{N}$ 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 it is usually easy to build equilibrium matchings that satisfy that constraint. Specifically, the equilibria constructed 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 objects is evidently in the no-disposal core: any coalition that blocks without disposal also blocks with disposal and hence the set of candidate submatchings that must be ruled out as blocking coalitions shrinks in the no-disposal case. Any such μ therefore continues to have a competitive equilibrium.

What about arbitrary matchings in the no-disposal core?

Theorem 6 *For any profile $\succ$ including profiles with arbitrary lifespans, if μ is in the no-disposal core of $\succ$ then μ can be competitively implemented.*

Proof. Recall from section 4 that each chain is 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, there are no rays in $\mathcal{S}$. 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 continue 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 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'. 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 $i \leq_S j$ mean that, for chain 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 $\leq_S$-predecessors of i in S, T, j and all $\leq_{S^*}$-successors of j in S^* 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. ■

A Appendix: lengthy and slow termination

Proposition 1 *For any ordinal α that is at most countable, there exists a profile $\succ$ such that modified top trading cycles terminates at round α.*

Proof. Suppose α is countably infinite. Let the favorite of each $i \in \mathbb{N}$ be i, let f be a bijection from $\mathbb{N}$ to α, and, for any $i \in \omega$, let the chain that exits in round $f(i)$ consist of i alone. Termination then occurs at round α. If α is finite, let the favorite of each

$i = 1, ..., \alpha - 1$ be i and let the favorite of $i \geq \alpha$ be $i + 1$. Then α admissible chains form in round 1 and they can exit in arbitrary order. ■

That most of the chains in the above proof consist of one agent consuming his endowment is inessential. For example, if the agents can be partitioned into a countably infinite set of chains and α is countably infinite then there will be a bijection from the set of chains to α and hence we may again let one chain exit at each of α rounds.

Modified top trading cycles if **simple** if, for any round α at which some chain S exits and any admissible chain S' that forms at α, $\min S \leq \min S'$.

If k is a finite ordinal then ωk denotes the ordinal that is order isomorphic to k copies of ω, say $\{1^1, 2^1, ...\}, ..., \{1^k, 2^k, ...\}$, 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 $\succ$, simple modified top trading cycles terminates at or before round ωL.*

Proof. Apply simple modified top trading cycles to profile $\succ$ and let μ be the matching that is generated. Suppose some agents remain unmatched after ωk rounds, where k is finite, and 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 $l_s < \omega k$. Since S^* is finite, $l = \max\{l_s : s \in S^*\}$ is well-defined and $l < \omega k$. Hence S^k is admissible at any round greater than l. Simplicity therefore implies that each chain that exits after l and before ωk contains an agent with index less than $\min S^k$. Hence S^k must exit on or before round $l + \min S^k$. Since ωk is a limit ordinal, $l + \min S^k < \omega k$, giving us 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 ωk, contradicting the Lemma.

To conclude, suppose that the algorithm has proceeded through $\omega(L-1)$ rounds, which implies that the infinite chains $S^1, ..., S^{L-1}$ have exited 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 $T = \{i \in \mathbb{N} : i$ exits in one of the first ω rounds 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 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 following a further ω rounds: 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 where any application of modified top trading cycles, whether simple or not, cannot terminate before ωL rounds have passed.

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, 2019, 'Optimal dynamic matching,' *Theoretical Economics*, forthcoming.

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

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

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

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

[12] Ma J., 1995, 'Strategyproofness and the strict core in a market with indivisibilites,' *International Journal of Game Theory* 23: 75-83.

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

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

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

[16] Schummer J., 2019, 'Influencing Waiting Lists' *mimeo.*

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