Patience ensures fairness

Florian Brandl* and Andrew Mackenzie† ‡

This draft: January 22, 2024

Abstract

We revisit the problem of fairly allocating a sequence of time slots when agents may have different levels of patience (Mackenzie and Komornik, 2023). For each number of agents, we provide a lower threshold and an upper threshold on the level of patience such that (i) if each agent is at least as patient as the lower threshold, then there is a *proportional* allocation, and (ii) if each agent is at least as patient as the upper threshold and moreover has weak preference for earlier time slots, then there is an *envy-free* allocation. In both cases, the proof is constructive.

Keywords: fair division, intertemporal choice, proportionality, no-envy

JEL Codes: D63, D71

1 Introduction

Can fairness be achieved without dividing goods into pieces, redistributing money, or randomizing, simply by taking turns? For example, can timeshare owners fairly allocate calendar dates, or can divorced parents identify a fair custody schedule? If all parties are sufficiently impatient, then clearly the answer is no. In this article, we revisit a recent idealized model of taking turns (Mackenzie and Komornik, 2023), and we prove that if all parties are sufficiently patient, then the answer is yes. Moreover, if all parties are sufficiently patient and have weak preference for earlier time slots, then the answer is emphatically yes.

In the model, the *time slots* in $T = \{1, 2, \dots\}$ are to be partitioned into *schedules* for n agents, and each agent i has preferences over schedules that can be represented by a countably additive probability measure u_i. An allocation is *proportional* if each agent measures his own schedule to be worth at least $\frac{1}{n}$ (Steinhaus, 1948), and is moreover *envy-free* if no agent measures another's schedule to be worth more than his own (Tinbergen, 1946; Foley, 1967). These fairness notions are central in the large literature on fair division, which spans perfectly divisible cakes, collections of indivisible objects, and classical exchange economies.

*Department of Economics, University of Bonn, Bonn, Germany. Email: florian.brandl@uni-bonn.de

†Department of Microeconomics and Public Economics, Maastricht University, Maastricht, the Netherlands. Email: a.mackenzie@maastrichtuniversity.nl

‡Florian Brandl acknowledges support by the DFG under the Excellence Strategy EXC-2047. We thank Vilmos Komornik and seminar participants of the Randomization, Neutrality, and Fairness Workshop at the Simons Laufer Mathematical Sciences Institute (Berkeley, October 2023).

We formalize each agent's level of patience using a generalization of the heavy-tails condition of Kakeya (Kakeya, 1914; Kakeya, 1915): for each $k \in [0, \infty)$, we say that u_i is *k-Kakeya* if it always measures the relative value of the future to the present to be at least k, or equivalently, if for each $t \in T$ we have $u_i(\{t+1, t+2, \dots\}) \geq k \cdot u_i(\{t\})$. We let $\mathcal{U}_k$ denote the set of these utility functions, and remark that this is a general class that includes many standard models of discounting; for example, if u_i is geometric discounting with respect to discount factor $\delta_i \in (0,1)$, in the sense that for each $S \subseteq T$ we have $u_i(S) = (1-\delta_i) \sum_{t \in S} \delta_i^{t-1}$, then $u_i \in \mathcal{U}_k$ if and only if $\delta_i \geq \frac{k}{k+1}$. In addition, we sometimes impose weak preference for earlier time slots: we say that u_i is *monotonic* if it always assigns an earlier time slot at least as much value as a later time slot, and we let $\mathcal{U}_{\mathrm{M}}$ denote the set of these utility functions.

In the special case that n agents have identical preferences given by a common discount factor $\delta \in [\frac{n-1}{n}, 1)$, the existence of envy-free allocations follows from a classic result for repeated games: if the agents are sufficiently patient given the number of action profiles, then for each convex combination of stage game payoffs, there is a sequence of action profiles that generates those payoffs (Sorin, 1986; Fudenberg and Maskin, 1991). Indeed, consider the repeated game with the following stage game: the oldest agent selects any agent, every other agent has a single dummy action, and the selected agent receives payoff 1 while the others receive payoff 0. By the classic result, there is a sequence of selected agents for which each agent receives payoff $\frac{1}{n}$, and as agents have identical preferences, the associated allocation is envy-free.

There are a variety of procedures that yield constructive proofs of the classic result. Translated into fair division, (i) the Fudenberg-Maskin procedure (Fudenberg and Maskin, 1991) iteratively selects any agent whose cumulative utility is lowest, (ii) the more flexible Sorin procedure (Sorin, 1986) iteratively selects any agent who still requires at least $\frac{1}{n}$ of the remaining utility, and (iii) the iterative application of Rényi's Greedy Algorithm (Rényi, 1957) iteratively asks each agent to construct a schedule for himself that achieves his target utility, by iteratively taking each remaining time slot unless it would make his schedule too valuable. All of these procedures—and many others—work more generally whenever agents have identical Kakeya preferences and are sufficiently patient, but none of them work when preferences are not identical (Mackenzie and Komornik, 2023). This reinforces a recurring theme in the literature: allowing agents to have different levels of patience—and thereby introducing the possibility of gains from intertemporal trade—has profound implications.[1]

For the problem we consider—fairly taking turns when agents need not have identical preferences—Mackenzie and Komornik (2023) established that

- if $n = 2$ and both agents have utility functions in $\mathcal{U}_{n-1}$, then envy-free allocations can be constructed using a version of the ancient Divide and Choose procedure;
- if $n = 3$ and all agents have utility functions in $\mathcal{U}_{n-1}$, then proportional allocations can be constructed using either their Iterative Apportionment procedure or their Simultaneous Apportionment procedure; and

[1]For example, allowing one agent to be more patient than another has significant implications for capital accumulation (Ramsey, 1928; Becker, 1980; Rader, 1981), bargaining (Rubinstein, 1982), reputation (Fudenberg and Levine, 1989), repeated games (Lehrer and Pauzner, 1999; Salonen and Vartiainen, 2008; Chen and Takahashi, 2012; Sugaya, 2015), preference aggregation toward a social discount factor (Weitzman, 2001; Jackson and Yariv, 2015; Chambers and Echenique, 2018), and endogeneous discounting (Kochov and Song, 2023)

- in general, for each $\varepsilon > 0$, if all agents have utility functions in $\mathcal{U}_{\frac{1-\varepsilon}{\varepsilon}} \cap \mathcal{U}_{\mathrm{M}}$, then allocations that are ε-approximately envy-free can be constructed using a simple Round-Robin procedure (Caragiannis, Kurokawa, Moulin, Procaccia, Shah, and Wang, 2019).

Observe that there may not be proportional allocations if the agents do not have utility functions in $\mathcal{U}_{n-1}$, as in this case the first time slot alone may be worth more than $\frac{1}{n}$ to everyone; thus (i) when the number of agents is small, these results provide a tight bound on the level of patience required to guarantee that there are fair allocations, and (ii) in general, these results require a level of patience above what may strictly be necessary in order to guarantee that there are approximately fair allocations.

In this article, we require a level of patience above what may strictly be necessary in order to guarantee that there are fair allocations in general. In particular, for a particular list of patience bounds $(\mathsf{p}(n))_{n\in\mathbb{N}}$, we establish that

- in general, if all agents have utility functions in $\mathcal{U}_{2n-1}$, then there are proportional allocations (Theorem 1), and
- in general, if all agents have utility functions in $\mathcal{U}_{\mathsf{p}(n)} \cap \mathcal{U}_{\mathrm{M}}$, then there are envy-free allocations (Theorem 2).

Notably, both proofs are constructive. The former result is established using Iterative Cycle Apportionment, a novel modification of Iterative Apportionment. For both procedures, iteratively the remaining agents construct and assign a schedule, but only in the novel procedure do they automatically skip certain time slots in earlier rounds to reserve them for later rounds. The latter result is established using the Tripartition Algorithm, a novel alternative to the classic Greedy Algorithm. For both procedures, an agent i is given a schedule S that he considers 'very divisible' and a target value $v \in (0, u_i(S))$, and uses the procedure to output a schedule $S^* \subseteq S$ such that (i) $u_i(S^*) = v$, and (ii) i considers $S \setminus S^*$ 'somewhat divisible,' but only with the novel procedure does consensus across all agents that S is 'very divisible' lead to consensus across all agents that *both* S^* and $S \setminus S^*$ are 'somewhat divisible.' Ultimately, the Tripartition Algorithm allows us to import the Aziz-Mackenzie procedure for fairly dividing a perfectly divisible cake (Aziz and Mackenzie, 2016) to our model.

Altogether, our main results formalize the statement that if all parties are sufficiently patient, then fairness can be achieved without dividing goods into pieces, redistributing money, or randomizing, simply by taking turns. In particular, this statement holds with no further hypothesis for the weaker fairness notion of proportionality, and holds under the additional hypothesis that all parties have weak preference for earlier time slots for the stronger fairness notion of no-envy.

2 Model

We consider economies where a countably infinite collection of time slots is to be partitioned into schedules for the agents, and each agent has preferences over schedules that may be represented by a countably additive probability measure.

Definition: An *economy* is specified by a pair (n, u) with $n \in \mathbb{N}$ as follows:

- $N \equiv \{1, 2, \ldots, n\}$ is the set of *agents*;

- $T \equiv \{1, 2, \dots\}$ is the countably infinite collection of *time slots*;
- $\mathcal{S} \equiv 2^T$ is the collection of *schedules*, which is each agent's consumption space;
- $u = (u_i)_{i \in N}$ is the profile of *utility functions*: for each $i \in N$, $u_i : \mathcal{S} \to [0, 1]$ is a countably additive probability measure[2] representing the preferences of i; and
- $\Pi \subseteq \mathcal{S}^N$ is the collection of *(partitional) allocations*: for each $\pi \in \mathcal{S}^N$, we have $\pi \in \Pi$ if and only if (i) for each pair $i, j \in N$, $\pi_i \cap \pi_j = \emptyset$, and (ii) $\cup_{i \in N} \pi_i = T$.

Whenever we refer to an arbitrary economy, we assume all of the above notation.

We are interested in the existence of allocations that satisfy the following normative axioms, and in constructing these allocations whenever possible.

Definition: Fix an economy and let $\pi \in \Pi$. We say that π satisfies

- *no-envy* if for each pair $i, j \in N$, $u_i(\pi_i) \geq u_i(\pi_j)$;
- *proportionality* if for each $i \in N$, $u_i(\pi_i) \geq \frac{1}{n}$; and
- *efficiency* if there is no $\pi' \in \Pi$ such that
  (i) for each $i \in N$, $u_i(\pi'_i) \geq u_i(\pi_i)$, and
  (ii) for some $i \in N$, $u_i(\pi'_i) > u_i(\pi_i)$.

It is easy to verify that no-envy implies proportionality. Our results involve economies where each agent is sufficiently patient given the number of agents, which we articulate using a generalization of a condition due to Kakeya (Kakeya, 1914; Kakeya, 1915):

Definition: Fix an economy and a utility function u_0. For each $k \in [0, \infty)$, we say that u_0 is

- *k-Kakeya* if for each $t \in T$, we have $u_0(\{t+1, t+2, \dots\}) \geq k \cdot u_0(\{t\})$; and
- *monotonic* if $u_0(\{1\}) \geq u_0(\{2\}) \geq \dots$.

We let $\mathcal{U}_k \subseteq [0, 1]^{\mathcal{S}}$ and $\mathcal{U}_{\mathrm{M}} \subseteq [0, 1]^{\mathcal{S}}$ denote the sets of k-Kakeya and monotonic utility functions, respectively.

Observe that an agent with a higher-level Kakeya utility function is, in a particular sense, more patient.

3 Results

Before proceeding, we make a brief remark about notation: we use i to index agents, t to index dates, and we use integer interval notation: for each pair $a, b \in \mathbb{Z}$, we let $[\![a, b]\!]$ denote $\{a, a+1, \dots, b\}$.

[2]That is, we have (i) $u_i(T) = 1$, and (ii) for each $S \in \mathcal{S}$, $u_i(S) = \sum_{t \in S} u_i(\{t\})$.

3.1 Preliminaries

We begin by introducing some simple methods for constructing a sub-schedule from an infinite source schedule that prove useful for our analysis.

Definition: Let $S \in \mathcal{S}$ be infinite, let $t \in S$, and let $\ell \in \mathbb{N}$ be a length. Then

- $S_{t\rightarrow} \equiv \{t' \in S \mid t' \geq t\}$ denotes the *tail of S starting at t* and $S_{t\circ\rightarrow} \equiv S_{t\rightarrow} \setminus \{t\}$ denotes the *tail of S starting after t*;
- $\mathcal{C}_\ell(t|S) \subseteq S$ denotes the *ℓ-cycle starting at t (given S)* constructed by beginning with an empty basket, considering the time slots in $S_{t\rightarrow}$ in sequence, and iteratively adding one time slot and then skipping $\ell - 1$ time slots;[3] and
- $\mathcal{I}_\ell(t|S) \subseteq S$ denotes the *ℓ-interval starting at t (given S)* consisting of the ℓ earliest members of $S_{t\rightarrow}$.

We discuss tails, cycles, and intervals in sequence. First, a tail of a reference schedule S consists of an initial time slot and all future ones in S. Our notion of patience involves tails of T: a utility function is k-Kakeya if for each $t \in T$, the value of the tail of T starting after t is at least k times the value of t alone. We will more generally be interested in other reference schedules with this property.

Definition: Fix an economy. For each $i \in N$, each $k \in \{0, 1, \ldots\}$, and each $S \in \mathcal{S}$, we say that S is *k-divisible for i* if for each $t \in S$, $u_i(S_{t\circ\rightarrow}) \geq k \cdot u_i(\{t\})$.

This terminology is justified by the fact that if a schedule S is 1-divisible for i, then for each target value $v \in [0, u_i(S)]$, an adaptation of the well-known Greedy Algorithm for representing numbers in arbitrary bases (Rényi, 1957) can be used to cut and remove a sub-schedule from S that is worth precisely v. Moreover, if S is k-divisible for $k > 1$, then the Greedy Algorithm can successfully be used in this manner k consecutive times.

Definition: *Greedy Algorithm.* Fix an economy. For each $i \in N$, each $S \in \mathcal{S}$, and each $v \in [0, u_i(S)]$, let $\mathcal{G}_i(v|S) \subseteq S$ denote the *greedy schedule (for agent i given source S and target v)* constructed by beginning with an empty basket, considering the time slots in S in sequence, and adding each time slot to the basket if and only if the value of the basket according to u_i will not exceed v.

Theorem MK (Mackenzie and Komornik, 2023): Fix an economy. For each $i \in N$, each $k \in \mathbb{N}$, each $S \in \mathcal{S}$ that is k-divisible for i, and each $v \in [0, u_i(S)]$, if $S^* = \mathcal{G}_i(v|S)$, then

- $u_i(S^*) = v$, and
- $S \setminus S^*$ is $(k-1)$-divisible for i.

[3]Equivalently, $\mathcal{C}_\ell(t|S)$ can be described using the language of modular arithmetic: it is the result of re-indexing the members of $S_{t\rightarrow}$ with the earliness-preserving bijection from $S_{t\rightarrow}$ to T, then taking the residue class modulo ℓ that includes 1.

To summarize, tails are useful for articulating divisibility in a manner that allows an individual agent to make precise cuts. Cycles, in turn, are useful because agents with monotonic preferences agree about how to rank cycles of a given length, and therefore agree that a tail's first cycle is worth a given fraction of the tail's value.

LEMMA 1: Fix an economy. For each $i \in N$ such that $u_i \in \mathcal{U}_{\mathrm{M}}$, each infinite $S \in \mathcal{S}$, each $t \in S$, and each $\ell \in \mathbb{N}$, we have $u_i(\mathcal{C}_\ell(t|S)) \geq (\frac{1}{\ell}) \cdot u_i(S_{t\rightarrow})$.

The proof is in Appendix 1. Finally, intervals are useful because if a schedule intersects each interval of a given length in a given tail, then agents with monotonic preferences agree that it is at least as valuable as an associated cycle, which in turn they agree is worth at least a given fraction of a later tail's value.

DEFINITION: Let $S \in \mathcal{S}$ be infinite and let $S^* \subseteq S$ be nonempty. For each length $\ell \in \mathbb{N}$, we say that S^* is *ℓ-dense in S* if for each $t \in S$ such that $t \geq \min S^*$, $S^* \cap \mathcal{I}_\ell(t|S) \neq \emptyset$.

LEMMA 2: Fix an economy. For each $i \in N$ such that $u_i \in \mathcal{U}_{\mathrm{M}}$, each $k \in [0, \infty)$, each $S \in \mathcal{S}$ that is k-divisible for i, each $\ell \in \mathbb{N}$ such that $\ell \leq k+1$, and each $S^* \subseteq S$ that is ℓ-dense in S, we have that S^* is $[(\frac{1}{\ell}) \cdot (k - (\ell - 1))]$-divisible for i.

The proof is in Appendix 1.

3.2 Proportionality

When there are three agents who each have 2-Kakeya preferences, a proportional allocation can be constructed using either the ITERATIVE APPORTIONMENT procedure or the SIMULTANEOUS APPORTIONMENT procedure (Mackenzie and Komornik, 2023). In the former, iteratively the remaining agents construct and assign a single schedule; in the latter, the agents begin constructing all three schedules simultaneously. Our contribution is more closely related to the former.

In particular, ITERATIVE APPORTIONMENT belongs to a class of procedures that can each be described as follows. There are n rounds indexed by $[\![1, n]\!]$, and in a given round $r \in [\![1, n-1]\!]$, N_r is the set of remaining agents, T_r is the set of remaining time slots, and $T_r^* \subseteq T_r$ is the set of time slots under consideration. The procedure works as follows:

- *Round r, $r \in [\![1, n-1]\!]$.* The agents in N_r start with an empty basket and consider the time slots in T_r^* in sequence. At each time slot t, each remaining agent places a flag in the basket if and only if he measures the value of the basket with t to exceed $\frac{1}{n}$. If there are no flags, then t is added to the basket and the remaining agents move to the next time slot. If there is one flag, then t is added to the basket, the basket is assigned to the agent who placed the flag, the basket's recipient exits with the basket's time slots, and we move to the next round. If there are multiple flags, then t is skipped and the remaining agents move to the next time slot. If all time slots are considered and there is no time slot with one flag, then the basket with its limit schedule—that is, the union of its schedules across all time periods—is assigned to any agent whose utility for it is highest, the basket's recipient exits with the basket's time slots, and we move to the next round.

- *Round n.* The remaining agent receives the remaining time slots.

As illustrated by the proof of Theorem 6 in Mackenzie and Komornik (2023), for each round $r \in [\![1, n-1]\!]$, if the remaining agents agree that (i) T_r^* is worth at least $\frac{|N_r|}{n}$, and (ii) T_r^* is 1-divisible, then the basket's recipient values the basket at least $\frac{1}{n}$ while the other remaining agents value the basket at most $\frac{1}{n}$. That said, Iterative Apportionment sets $T_r^* = T_r$ in each round, and after the first round, the remaining agents may not agree that what remains is 1-divisible. While Iterative Apportionment nevertheless works in the special case that there are three agents who each have 2-Kakeya preferences, the proof does not generalize.

In order to guarantee that in later rounds there will be consensus that the set of time slots under consideration is 1-divisible, we modify Iterative Apportionment by automatically skipping time slots to reserve them for later consideration, and doing so in a manner that still ensures there is enough value among the time slots currently under consideration. In particular, (i) in the first round, only time slots from the top n-cycle (given T) are considered, (ii) in the second round, only time slots from the top two n-cycles (given T) are considered, and (iii) in general, for each $r \in [\![1, n-1]\!]$ we define $T_r^* \equiv T_r \cap [\cup_{r' \leq r} \mathcal{C}_n(r'|T)]$. We refer to this procedure as Iterative Cycle Apportionment.

In order to guarantee Iterative Cycle Apportionment constructs a proportional allocation, we require that (i) each agent deems each n-cycle to be 1-divisible, and (ii) each agent has monotonic preferences. Moreover, if an economy satisfies only the first requirement, then we can still construct a proportional allocation by applying Iterative Cycle Apportionment to an associated economy with monotonic preferences and then modifying the resulting allocation.

Definition: Fix an economy and a utility function u_0. The *monotonic reordering of* u_0 is the monotonic utility function $u_0^{\mathrm{M}} \in \mathcal{U}_{\mathrm{M}}$ formed by sorting the values of u_0 in non-increasing order: for each $v \in (0, 1]$, $|\{t \in T \mid u_0(\{t\}) = v\}| = |\{t \in T \mid u_0^{\mathrm{M}}(\{t\}) = v\}|$. We refer to $(n, u^{\mathrm{M}}) = (n, (u_i^{\mathrm{M}})_{i \in N})$ as the *monotonic economy induced by* (n, u).

First, we show that patience is preserved under monotonic reordering.

Lemma 3: For each $k \in [0, \infty)$ and each $u_0 \in \mathcal{U}_k$, $u_0^{\mathrm{M}} \in \mathcal{U}_k \cap \mathcal{U}_{\mathrm{M}}$.

The proof is in Appendix 2. Second, we show that a proportional allocation π^{M} for the economy with reordered utility functions can be used to construct a proportional allocation π for the original economy. In particular, in the original economy, we consider the time slots in sequence, and at each time slot t we ask the recipient of t according to π^{M} to select his favorite of the remaining time slots; we then assign any time slots that are not selected to the first agent. Since each agent's utility for π in the original economy is at least as high as his utility for π^{M} in the economy with reordered utility functions, π is proportional.

Lemma 4: For each economy (n, u), if (n, u^{M}) has a proportional allocation, then (n, u) does as well.

The proof is in Appendix 2. Using (i) the preceding two lemmas, (ii) the implication of Lemma 2 that each agent with $(2n-1)$-Kakeya preferences deems each n-cycle to be 1-divisible, and (iii) an argument establishing that Iterative Cycle Apportionment constructs a proportional allocation under the claimed requirements, we establish our first main result.

THEOREM 1: Fix an economy. If for each $i \in N$, $u_i \in \mathcal{U}_{2n-1}$, then there is a proportional allocation.

The proof is in Appendix 2. By Observation 1 of Mackenzie and Komornik (2023), if there is a proportional allocation, then there is a proportional allocation that is efficient; we thus immediately have the following.

COROLLARY 1: Fix an economy. If for each $i \in N$, $u_i \in \mathcal{U}_{2n-1}$, then there is a proportional allocation that is efficient.

3.3 No envy

Unfortunately, ITERATIVE CYCLE APPORTIONMENT need not construct a cycle that is envy-free: an agent who receives his schedule in an earlier round may envy an agent who receives his schedule in a later round. In order to guarantee that there is an envy-free allocation, we require the agents to be so patient that they can construct such an allocation using a procedure for fair division of a perfectly divisible cake.

In particular, we import the AZIZ-MACKENZIE procedure for fairly dividing a perfectly divisible cake (Aziz and Mackenzie, 2016) to our setting. This belongs to the family of Robertson-Webb procedures (Robertson and Webb, 1998) in the sense that it fairly divides $[0, 1]$ using only EVALUATE and CUT queries, where (i) in an EVALUATE query, an agent i is given a sub-interval $[x, y] \subseteq [0, 1]$ and asked to evaluate its worth according to u_i, and (ii) in a CUT query, an agent i is given a point $x \in [0, 1]$ and a value $v \in [0, 1]$, then asked to cut at a point $y \in [x, 1]$ such that $u_i([x, y]) = v$ (if possible).

Fortunately, the AZIZ-MACKENZIE procedure does not rely crucially on the structure of $[0, 1]$, and can be described in a more general manner that covers our setting using abstract EVALUATE and CUT queries as follows: begin with the coarsest partition of the cake, and at each step either (i) give an agent i a member of the current partition, then ask him to evaluate its worth according to u_i; or (ii) give an agent i a member S of the current partition and a target value $v \in [0, u_i(S)]$, then ask i to cut $S^* \subseteq S$ such that $u_i(S^*) = v$ (if possible), and finally update the partition by replacing S with S^* and $S \setminus S^*$ (if the cut was successful).

The challenge for importing such a procedure for perfectly divisible cakes to our setting is that for perfectly divisible cakes, the cuts are always successful, but in our setting that need not be the case. Moreover, even if i is given a set S that he deems k-divisible for $k \geq 1$, the GREEDY ALGORITHM does not suffice for our purposes: though i can cut a set S^* that is worth the target value while guaranteeing that he deems $S \setminus S^*$ to be $(k-1)$-divisible, there is no guarantee that i deems S^* to be divisible at all, and there is no guarantee that another agent deems either S^* or $S \setminus S^*$ to be divisible at all.

To solve this problem, we introduce an alternative to the GREEDY ALGORITHM. Loosely, when cutting S^* from S, the GREEDY ALGORITHM can be used as long as S is 1-divisible, and involves (i) a small cost on the divisibility of $S \setminus S^*$ (relative to the divisibility of S) for the cutter, (ii) an unbounded cost on the divisibility of $S \setminus S^*$ for other agents, and (iii) an unbounded cost on the divisibility of S^* for everybody. Our alternative, which we refer to as the TRIPARTITION ALGORITHM, requires S to be

5-divisible, and involves (i) a large but bounded cost on the divisibility of $S \setminus S^*$ for everybody, and (ii) a large but bounded cost on the divisibility of S^* for everybody.

DEFINITION: *Tripartition Algorithm.* Fix an economy. For each $i \in N$, each infinite $S \in \mathcal{S}$, and each $v \in (0, u_i(S))$, let $\mathcal{T}_i(v|S) \subseteq S$ denote the *tripartition schedule (for agent i given source S and target v)* constructed in two cases as follows. First, if $v \leq (\frac{1}{2}) \cdot u_i(S)$, then do the following.

- Define $t \equiv \min\{t' \in S \mid (\frac{1}{3}) \cdot u_i(S_{t' \to}) \leq v\}$. Observe that this is well-defined because $v > 0$ and $\sum_{t' \in S} u_i(t') \leq 1$.
- Define $S^{\text{sort}} \equiv \mathcal{C}_3(t|S)$, $S^{\text{skip}} \equiv \mathcal{C}_3(t+1|S)$, and $S^{\text{take}} \equiv \mathcal{C}_3(t+2|S)$. Observe that these three schedules partition $S_{t \to}$.
- Define $\mathcal{T}_i(v|S) \equiv S^{\text{take}} \cup \mathcal{G}_i(v - u_i(S^{\text{take}})|S^{\text{sort}})$.

Second, if $v > (\frac{1}{2}) \cdot u_i(S)$, then define $\mathcal{T}_i(v|S) \equiv S \setminus [\mathcal{T}_i(u_i(S) - v|S)]$.

PROPOSITION 1: Fix an economy. For each pair $i, j \in N$ such that $u_i \in \mathcal{U}_{\text{M}}$ and $u_j \in \mathcal{U}_{\text{M}}$, each $k \in [5, \infty)$, each $S \in \mathcal{S}$ that is k-divisible for both i and j, and each $v \in (0, u_i(S))$, if $S^* = \mathcal{T}_i(v|S)$, then

- $u_i(S^*) = v$,
- S^* is $[(\frac{1}{3}) \cdot (k-2)]$-divisible for both i and j, and
- $S \setminus S^*$ is $[(\frac{1}{3}) \cdot (k-2)]$-divisible for both i and j.

The proof is in Appendix 3. Because there is an upper bound on the number of queries required by the AZIZ-MACKENZIE procedure (Aziz and Mackenzie, 2016), and thus an upper bound on the number of cuts that it requires, it follows that if the agents are sufficiently patient, then we can construct an envy-free procedure by importing the AZIZ-MACKENZIE procedure to our setting, making each cut before the final step with the TRIPARTITION ALGORITHM and using the GREEDY ALGORITHM for the final step. Indeed, if the agents are sufficiently patient, then even if the upper bound is reached exclusively through cuts, we will have that (i) before each step of the procedure except possibly the final one, the agents agree that each member of the partition is at least 5-divisible, and (ii) before the final step of the procedure, the agents agree that each member of the partition is 1-divisible.

DEFINITION: *Patience bounds.* Define the divisibility requirement for one cut by $\mathsf{d}(1) \equiv 1$, and for each number of cuts $c \in \mathbb{N}$ such that $c > 1$, define the associated divisibility requirement by $\mathsf{d}(c) \equiv 3 \cdot \mathsf{d}(c-1) + 2$. Finally, for each $n \in \mathbb{N}$, define the associated patience requirement for using the AZIZ-MACKENZIE procedure by

$$\mathsf{p}(n) \equiv \mathsf{d}\left(n^{n^{n^{n^{n}}}}\right).$$

THEOREM 2: Fix an economy. If for each $i \in N$, $u_i \in \mathcal{U}_{\mathsf{p}(n)} \cap \mathcal{U}_{\text{M}}$, then there is an envy-free allocation.

The proof is in Appendix 3.

4 Discussion

We have shown that (i) if agents are sufficiently patient, then there are proportional allocations, and (ii) if agents are sufficiently patient and moreover have weak preference for earlier time slots, then there are envy-free allocations.

In order to establish the former result, we showed that whenever agents have $(2n-1)$-Kakeya preferences, we can map the given economy to an associated economy with monotonic preferences and run Iterative Cycle Apportionment, then use this allocation to determine the order in which agents select objects in the original economy. We remark that this argument works whenever the agents agree that each n-cycle is 1-divisible, and that more elaborate arguments could improve our bound on the sufficient level of patience for this argument to go through, but we have chosen to omit such arguments as they would have little impact on our main message. Indeed, this argument cannot establish the conjecture of Mackenzie and Komornik (2023) that there are proportional allocations whenever agents have $(n-1)$-Kakeya preferences, and there is no particular reason to believe that this argument establishes the tightest possible bound. That said, it does suffice to establish that a bound on the level of patience that is linear in the number of agents guarantees that there are proportional allocations.

In order to establish the latter result, we introduced the Tripartition Algorithm that allows an agent to divide a set that everybody deems divisible into two sets such that (i) he deems one of them to be worth a target value, and (ii) everybody agrees that both sets retain some divisibility. Because the Aziz-Mackenzie procedure has an upper bound on the total number of queries it requires (Aziz and Mackenzie, 2016) and thus on the total number of cuts that it requires, we can import it to our setting using the Tripartition Algorithm to make cuts, provided the agents are sufficiently patient and have weak preference for earlier time slots. In fact, the same logic applies to other bounded procedures for perfectly divisible cakes that can be described using the abstract Evaluate and Cut queries defined in Section 3.3.[4] Unfortunately, however, this argument has its limitations: in order to use the Tripartition Algorithm to import *any* procedure that ultimately constructs an allocation by making at least $n-1$ cuts, we require a level of patience that is at least exponential in the number of agents.[5]

We conclude by highlighting two open questions. First, are there envy-free allocations whenever agents are sufficiently patient, even if they might sometimes prefer a later time slot to an earlier one? Second, are there envy-free allocations that are moreover efficient whenever agents are sufficiently patient?

[4]As an interesting example, fix an *entitlement vector* $e \in [0,1]^N$ such that each e_i is a rational number and $\sum e_i = 1$, and let us say an allocation π is *e-proportional* if for each $i \in N$ we have $u_i(\pi_i) \geq e_i$. In the special case that each entitlement is $\frac{1}{n}$, this is ordinary proportionality. The Cseh-Fleiner procedure (Cseh and Fleiner, 2020) is a bounded procedure for constructing e-proportional allocations for perfectly divisible cakes, and we can use the Tripartition Algorithm to import it to our setting; it follows that there are e-proportional allocations if agents are sufficiently patient and have weak preference for earlier time slots. The required patience level depends on the entitlement vector.

[5]We remark that there is a lower bound on the total number of queries required to construct an envy-free allocation (Procaccia, 2009), but because this does not yield a lower bound on the number of cuts, we cannot use it to draw stronger conclusions about the limitations of our argument for no-envy. That said, bounds on the number of required cuts have been previously investigated in the context of proportionality (Even and Paz, 1984; Edmonds and Pruhs, 2006), and analogous future work for no-envy could indeed provide further insight about this.

Appendix 1

In this appendix, we prove Lemma 1 and Lemma 2. We begin with Lemma 1.

Lemma 1 (Restated): Fix an economy. For each $i \in N$ such that $u_i \in \mathcal{U}_{\mathrm{M}}$, each infinite $S \in \mathcal{S}$, each $t \in S$, and each $\ell \in \mathbb{N}$, we have $u_i(\mathcal{C}_\ell(t|S)) \geq (\frac{1}{\ell}) \cdot u_i(S_{t\to})$.

Proof: Assume the hypotheses. Observe that for each pair $t_0, t_1 \in S$ such that $t_0 < t_1$, the earliness-preserving bijection from $\mathcal{C}_\ell(t_0|S)$ to $\mathcal{C}_\ell(t_1|S)$ maps each time slot to a later one, and thus since $u_i \in \mathcal{U}_{\mathrm{M}}$ we have $u_i(\mathcal{C}_\ell(t_0|S)) \geq u_i(\mathcal{C}_\ell(t_1|S))$. Since $S_{t\to}$ can be partitioned into ℓ cycles of length ℓ, of which $\mathcal{C}_\ell(t|S)$ has the earliest starting time slot, the desired conclusion follows immediately. ■

To conclude this appendix, we prove Lemma 2.

Lemma 2 (Restated): Fix an economy. For each $i \in N$ such that $u_i \in \mathcal{U}_{\mathrm{M}}$, each $k \in [0, \infty)$, each $S \in \mathcal{S}$ that is k-divisible for i, each $\ell \in \mathbb{N}$ such that $\ell \leq k+1$, and each $S^* \subseteq S$ that is ℓ-dense in S, we have that S^* is $[(\frac{1}{\ell}) \cdot (k - (\ell - 1))]$-divisible for i.

Proof: Assume the hypotheses. If $k = 0$, then we are done; thus let us assume $k > 0$, so S is infinite. We first consider the case that $S = T$. To begin, let $t \in S^*$. Since S^* is ℓ-dense in S, thus $S^*_{t\circ\to}$ is as well. To proceed, we establish four inequalities.

- First, since (i) the earliness-preserving bijection from $S^*_{t\circ\to}$ to $\mathcal{C}_\ell(t+\ell|S)$ maps each time slot to one that is at least as late, and (ii) $u_i \in \mathcal{U}_{\mathrm{M}}$, thus $u_i(S^*_{t\circ\to}) \geq u_i(\mathcal{C}_\ell(t+\ell|S))$.
- Second, by Lemma 1 we have $u_i(\mathcal{C}_\ell(t+\ell|S)) \geq (\frac{1}{\ell}) \cdot u_i(S_{(t+\ell)\to})$.
- Third, since (i) $S_{(t+\ell)\to} = S_{(t+1)\to} \backslash \{t+t'\}_{t' \in [\![1,\ell-1]\!]}$, and (ii) $u_i \in \mathcal{U}_{\mathrm{M}}$, thus $u_i(S_{(t+\ell)\to}) \geq u_i(S_{(t+1)\to}) - (\ell - 1) \cdot u_i(\{t\})$.
- Finally, since S is k-divisible for i, thus $u_i(S_{(t+1)\to}) \geq k \cdot u_i(\{t\})$.

By the above inequalities, $u_i(S^*_{t\circ\to}) \geq [(\frac{1}{\ell}) \cdot (k - (\ell - 1))] \cdot u_i(\{t\})$. Since $t \in S^*$ was arbitrary, thus S^* is $[(\frac{1}{\ell}) \cdot (k - (\ell - 1))]$-divisible for i, as desired.

To conclude, if $S \neq T$, then the desired conclusion follows from re-indexing S with the earliness-preserving bijection from S to T and applying the above argument. ■

Appendix 2

In this appendix, we prove Lemma 3, Lemma 4, and Theorem 1. We begin with Lemma 3.

Lemma 3 (Restated): For each $k \in [0, \infty)$ and each $u_0 \in \mathcal{U}_k$, $u_0^{\mathrm{M}} \in \mathcal{U}_k \cap \mathcal{U}_{\mathrm{M}}$.

Proof: By construction, we have $u_0^{\mathrm{M}} \in \mathcal{U}_{\mathrm{M}}$. We claim that $u_0^{\mathrm{M}} \in \mathcal{U}_k$. If $k = 0$, then we are done; thus let us assume $k > 0$.

Let $t \in T$. Since $u_0 \in \mathcal{U}_k$, thus u_0 does not assign zero to any time slot, so by construction we have $u_0^{\mathrm{M}}(\{t\}) > 0$. Define $t^* \equiv \max\{t' \in T \mid u_0(\{t'\}) \geq u_0^{\mathrm{M}}(\{t\})\}$;

this is well-defined as $\sum_{t' \in T} u_0(\{t'\}) = 1$. It follows that (i) by definition of t^* and construction of u_0^{M}, we have $u_0^{\mathrm{M}}(\{t+1, t+2, \dots\}) \geq u_0(\{t^*+1, t^*+2, \dots\})$; (ii) since $u_0 \in \mathcal{U}_k$, we have $u_0(\{t^*+1, t^*+2, \dots\}) \geq k \cdot u_0(\{t^*\})$; and (iii) by construction, we have $u_0(\{t^*\}) \geq u_0^{\mathrm{M}}(\{t\})$; thus altogether we have $u_0^{\mathrm{M}}(\{t+1, t+2, \dots\}) \geq k \cdot u_0^{\mathrm{M}}(\{t\})$. Since $t \in T$ was arbitrary, we are done. ■

Next, we prove Lemma 4.

Lemma 4 (Restated): For each economy (n, u), if (n, u^{M}) has a proportional allocation, then (n, u) does as well.

Proof: Let (n, u) be an economy and assume that $\pi^{\mathrm{M}} \in \Pi$ is proportional for (n, u^{M}). We construct $\pi \in \Pi$ as follows. First, for each $t \in T$, let i_t denote the agent who receives t according to π^{M}. Second, we proceed through the time slots in sequence, and at each time slot t, agent i_t consumes the remaining time slot $f(t)$ that (i) is at least as valuable as all other remaining time slots according to u_{i_t}, and (ii) is earliest among such time slots. Finally, agent 1 consumes any remaining time slots.

We claim that π is proportional for (n, u). Indeed, let $i \in N$, let $t \in \pi_i^{\mathrm{M}}$, and let $T^* \subseteq T$ be a set of $t-1$ time slots to which u_i assigns the highest value. Then

$$\begin{aligned} u_i(\{f(t)\}) &= \max_{t'' \in T \setminus \{f(t') | t' \in [\![1, t-1]\!]\}} u_i(t'') \\ &\geq \max_{t' \in T \setminus T^*} u_i(t') \\ &= u_i^{\mathrm{M}}(\{t\}). \end{aligned}$$

Since $t \in \pi_i^{\mathrm{M}}$ was arbitrary, thus $u_i(\pi_i) \geq u_i^{\mathrm{M}}(\pi_i^{\mathrm{M}}) \geq \frac{1}{n}$. Since $i \in N$ was arbitrary, we are done. ■

To conclude this appendix, we prove Theorem 1.

Theorem 1 (Restated): Fix an economy. If for each $i \in N$, $u_i \in \mathcal{U}_{2n-1}$, then there is a proportional allocation.

Proof: Assume the hypotheses. By Lemma 3, for each $i \in N$, we have $u_i^{\mathrm{M}} \in \mathcal{U}_{2n-1} \cap \mathcal{U}_{\mathrm{M}}$, and thus u_i^{M} assigns a positive utility to each time slot.

We claim that (n, u^{M}) has a proportional allocation. Indeed, we construct the desired allocation using Iterative Cycle Apportionment, which consists of n rounds indexed by $[\![1, n]\!]$. In each round $r \in [\![1, n-1]\!]$, the $n-(r-1)$ remaining agents consider the remaining time slots in the top r n-cycles in order to fill a basket and assign its contents S_r to an agent i_r. In the final round, the remaining time slots are assigned to the remaining agent.

For each round $r \in [\![1, n-1]\!]$, let (i) $N_r \equiv N \setminus \{i_{r'} \in N \mid r' < r\}$ denote the set of remaining agents at the start of round r, (ii) $T_r \equiv T \setminus [\cup_{r'<r} S_r]$ denote the remaining time slots at the beginning of round r, and (ii) $T_r^* \equiv T_r \cap [\cup_{r' \leq r} \mathcal{C}_n(r'|T)]$ denote the time slots under consideration during round r. Before providing further details about the construction of $(S_r)_{r \in [\![1, n-1]\!]}$, we make two observations.

First, we observe that for each round $r \in [\![1, n-1]\!]$ and each $i \in N$, T_r^* is 1-divisible for i. Indeed, since $u_i^{\mathrm{M}} \in \mathcal{U}_{2n-1}$, thus T is $(2n-1)$-divisible for i. Moreover, since

$\mathcal{C}_n(r|T) \subseteq T_r^*$ and $r \in [\![1, n-1]\!]$, thus T_r^* is n-dense in T. Finally, $u_i^{\mathrm{M}} \in \mathcal{U}_{\mathrm{M}}$. Altogether, then, by Lemma 2 we have that T_r^* is $[(\frac{1}{n}) \cdot ((2n-1)-(n-1))]$-divisible for i, as desired.

Second, we observe that for each round $r \in [\![1, n-1]\!]$ and each $i \in N$, we have $u_i^{\mathrm{M}}(T_r^*) \geq (\frac{1}{n-(r-1)}) \cdot u_i^{\mathrm{M}}(T_r)$. Indeed, define $C \equiv \cup_{r' \in [\![r,n]\!]} \mathcal{C}_n(r'|T)$. Since $u_i^{\mathrm{M}} \in \mathcal{U}_{\mathrm{M}}$, thus

$$
\begin{aligned}
u_i^{\mathrm{M}}(T_r^*) &= u_i^{\mathrm{M}}(T_r \setminus C) + u_i^{\mathrm{M}}(\mathcal{C}_n(r|T)) \\
&\geq u_i^{\mathrm{M}}(T_r \setminus C) + \left(\frac{1}{n-(r-1)}\right) \cdot u_i^{\mathrm{M}}(C) \\
&= u_i^{\mathrm{M}}(T_r \setminus C) + \left(\frac{1}{n-(r-1)}\right) \cdot u_i^{\mathrm{M}}(T_r \cap C) \\
&= \left[1 \cdot \left(\frac{u_i^{\mathrm{M}}(T_r \setminus C)}{u_i^{\mathrm{M}}(T_r)}\right) + \left(\frac{1}{n-(r-1)}\right) \cdot \left(\frac{u_i^{\mathrm{M}}(T_r \cap C)}{u_i^{\mathrm{M}}(T_r)}\right)\right] \cdot u_i^{\mathrm{M}}(T_r) \\
&\geq \left(\frac{1}{n-(r-1)}\right) \cdot u_i^{\mathrm{M}}(T_r),
\end{aligned}
$$

as desired. To ease verification, we highlight that the fives lines above use (i) the definition of T_r^*, (ii) the fact that $u_i^{\mathrm{M}} \in \mathcal{U}_{\mathrm{M}}$, (iii) the fact that $C \subseteq T_r$, (iv) simple re-writing, and (v) the fact that $u_i^{\mathrm{M}}(T_r \setminus C) + u_i^{\mathrm{M}}(T_r \cap C) = u_i^{\mathrm{M}}(T_r)$, respectively.

We now describe how baskets are filled and assigned. In each round $r \in [\![1, n-1]\!]$, the $n-(r-1)$ remaining agents consider the members of T_r^* in sequence. At each time slot t, each agent places a flag in the basket if and only if he measures the value of the basket with t to exceed $\frac{1}{n}$. If there are no flags, then t is added to the basket and the remaining agents move to the next time slot. If there is one flag, then let i_r be the agent who placed the flag, let S_r denote the contents of the basket with t, and assign S_r to i_r; in this case the round ends and we move to the next round. If there are multiple flags, then t is skipped and the agents move to the next time slot. If all time slots are considered and there is no time slot with one flag, then let S_r denote the basket's limit schedule (that is, the union of its schedules across all time periods), let i_r be any agent whose utility for S_r is highest, and assign S_r to i_r.

We claim that for each $r \in [\![1, n-1]\!]$, (i) $u_{i_r}^{\mathrm{M}}(S_r) \geq \frac{1}{n}$, and (ii) $i \in N_r \setminus \{i_r\}$ implies $u_i^{\mathrm{M}}(S_r) \leq \frac{1}{n}$. We proceed by induction, handling the base step and inductive step together. Indeed, assume that for each $r' < r$ we have both statements. By construction, we have the second statement for r. If there is a time slot with one flag, then we are done; thus let us assume there is no such time slot. By the inductive hypothesis, for each $i \in N_r$ we have $u_i^{\mathrm{M}}(T_r) \geq \frac{n-(r-1)}{n}$, so by the second observation we have $u_i^{\mathrm{M}}(T_r^*) \geq \frac{1}{n}$. If each time slot has no flag, then we are done; thus let us assume there is a time slot with multiple flags. By the first observation, for each $i \in N_r$ we have that T_r^* is 1-divisible for i; it is straightforward to verify that this implies there cannot be a maximum time slot with multiple flags. Then there is $i \in N_r$ who places an infinite collection of flags. Since $\lim_{t \to \infty} u_i^{\mathrm{M}}(\{t\}) = 0$, thus $u_i^{\mathrm{M}}(S_r) \geq \frac{1}{n}$, so $u_{i_r}^{\mathrm{M}}(S_r) \geq \frac{1}{n}$, as desired. For further details about the argument we omit as straightforward, see the analogous argument in the proof of Theorem 6 in Mackenzie and Komornik (2023).

By the above claim, Iterative Cycle Apportionment constructs a proportional allocation for (n, u^{M}). Thus by Lemma 4, (n, u) has a proportional allocation, as desired. ■

Appendix 3

In this appendix, we prove Proposition 1 and Theorem 2. We begin with Proposition 1.

Proposition 1 (Restated): Fix an economy. For each pair $i, j \in N$ such that $u_i \in \mathcal{U}_{\mathrm{M}}$ and $u_j \in \mathcal{U}_{\mathrm{M}}$, each $k \in [5, \infty)$, each $S \in \mathcal{S}$ that is k-divisible for both i and j, and each $v \in (0, u_i(S))$, if $S^* = \mathcal{T}_i(v|S)$, then

- $u_i(S^*) = v$,
- S^* is $[(\frac{1}{3}) \cdot (k-2)]$-divisible for both i and j, and
- $S \setminus S^*$ is $[(\frac{1}{3}) \cdot (k-2)]$-divisible for both i and j.

Proof: Assume the hypotheses. If we establish that $v \leq (\frac{1}{2}) \cdot u_i(S)$ implies the desired conclusion, then by construction we are done; thus let us assume $v \leq (\frac{1}{2}) \cdot u_i(S)$. Define t, S^{sort}, S^{skip}, and S^{take} as in the construction of $\mathcal{T}_i(v|S) = S^*$; thus $(\frac{1}{3}) \cdot u_i(S_{t\rightarrow}) \leq v$.

First, we claim that $u_i(S^{\text{skip}}) \leq (\frac{1}{2}) \cdot u_i(S_{t\rightarrow})$. Indeed, since $\mathcal{C}_3(t+1|S) = S^{\text{skip}}$ and $u_i \in \mathcal{U}_{\mathrm{M}}$, thus $u_i(S_{t\rightarrow}) \geq u_i(\mathcal{C}_3(t|S)) + u_i(S^{\text{skip}})$ and $u_i(\mathcal{C}_3(t|S)) \geq S^{\text{skip}}$. Altogether, then, $u_i(S^{\text{skip}}) \leq (\frac{1}{2}) \cdot u_i(S_{t\rightarrow})$, as desired.

Second, we claim that $v \leq (\frac{1}{2}) \cdot u_i(S_{t\rightarrow})$. Indeed, if $S_{t\rightarrow} = S$, then we are done; thus let us assume $S_{t\rightarrow} \subsetneq S$ and define $t^* \equiv \max S \setminus S_{t\rightarrow}$. Since S is 5-divisible for i we have $u_i(S_{t\rightarrow}) \geq 5 \cdot u_i(\{t^*\})$, and since $t^* < t$ we have $(\frac{1}{3}) \cdot u_i(S_{t^*\rightarrow}) > v$, so altogether we have

$$\begin{aligned}
\left(\frac{1}{2}\right) \cdot u_i(S_{t\rightarrow}) &= \left(\frac{1}{2}\right) \cdot \left(\frac{u_i(S_{t\rightarrow})}{u_i(\{t^*\}) + u_i(S_{t\rightarrow})}\right) \cdot u_i(S_{t^*\rightarrow}) \\
&\geq \left(\frac{1}{2}\right) \cdot \left(\frac{5}{6}\right) \cdot u_i(S_{t^*\rightarrow}) \\
&> \left(\frac{1}{3}\right) \cdot u_i(S_{t^*\rightarrow}) \\
&> v,
\end{aligned}$$

as desired.

Third, we claim that $u_i(S^{\text{take}}) \leq (\frac{1}{3}) \cdot u_i(S_{t\rightarrow})$. Indeed, since $\mathcal{C}_3(t+2|S) = S^{\text{take}}$ and $u_i \in \mathcal{U}_{\mathrm{M}}$, thus $u_i(S_{t\rightarrow}) = u_i(\mathcal{C}_3(t|S)) + u_i(\mathcal{C}_3(t+1|S)) + u_i(S^{\text{take}})$, $u_i(\mathcal{C}_3(t|S)) \geq S^{\text{take}}$, and $u_i(\mathcal{C}_3(t+1|S)) \geq S^{\text{take}}$. Altogether, then, $u_i(S^{\text{take}}) \leq (\frac{1}{3}) \cdot u_i(S_{t\rightarrow})$, as desired.

Fourth, we claim that S^{sort} is 1-divisible for i. Indeed, since (i) S is 5-divisible for i, and (ii) S^{sort} is 3-dense in S, thus by Lemma 2 we have that S^{sort} is $[(\frac{1}{3}) \cdot (5-2)]$-divisible for i, as desired.

To conclude, by the first two claims we have $u_i(S^{\text{sort}}) + u_i(S^{\text{take}}) = u_i(S_{t\rightarrow}) - u_i(S^{\text{skip}}) \geq (\frac{1}{2}) \cdot u_i(S_{t\rightarrow}) \geq v$, so $u_i(S^{\text{sort}}) \geq v - u_i(S^{\text{take}})$. Moreover, by construction of t and the third claim we have $v - u_i(S^{\text{take}}) \geq (\frac{1}{3}) \cdot u_i(S_{t\rightarrow}) - (\frac{1}{3}) \cdot u_i(S_{t\rightarrow}) = 0$, so altogether we have $v - u_i(S^{\text{take}}) \in [0, u_i(S^{\text{sort}})]$. Finally, by the fourth claim, S^{sort} is 1-divisible for i. Altogether, then, by Theorem MK we have that $u_i(\mathcal{G}_i(v - u_i(S^{\text{take}})|S^{\text{sort}})) = v - u_i(S^{\text{take}})$, so $u_i(S^*) = v$. Finally, (i) $u_i \in \mathcal{U}_{\mathrm{M}}$ and $u_j \in \mathcal{U}_{\mathrm{M}}$; (ii) S is k-divisible for both i and j; and (iii) since $\mathcal{C}_3(t+2|S) \subseteq S^* \subseteq S_{t\rightarrow}$ and $(S \setminus S_{t\rightarrow}) \cup \mathcal{C}_3(t+1|S) \subseteq S \setminus S^*$, thus S^* and $S \setminus S^*$ are both 3-dense in S; thus by Lemma 2, we have that S^* and $S \setminus S^*$ are both $[(\frac{1}{3}) \cdot (k - (3-1))]$-divisible for both i and j, as desired. ■

To conclude this appendix, we prove Theorem 2.

Theorem 2 (Restated): Fix an economy. If for each $i \in N$, $u_i \in \mathcal{U}_{\mathsf{p}(n)} \cap \mathcal{U}_{\mathrm{M}}$, then there is an envy-free allocation.

Proof: Assume the hypotheses. By Proposition 1, for each $c \in \mathbb{N}$, if (i) we begin from the coarsest partition of T and iteratively apply the Tripartition Algorithm to members of the current partition $c-1$ times; and (ii) for each $i \in N$, we have $u_i \in \mathcal{U}_{\mathsf{d}(c)} \cap \mathcal{U}_{\mathrm{M}}$; then (i) before each step of the procedure, the agents agree that each member of the current partition is 5-divisible, and (ii) when the procedure is complete, the agents agree that each member of the current partition is 1-divisible. Since the upper bound on the total number of queries for the Aziz-Mackenzie procedure (Aziz and Mackenzie, 2016) is also an upper bound on the total number of cuts for this procedure, it follows from this bound and the definition of $\mathsf{p}(n)$ that if we run the Aziz-Mackenzie procedure in our setting by using the Tripartition Algorithm at every step except the final one and using the Greedy Algorithm at the final step, then (i) before each step of the procedure, the agents agree that each member of the current partition is 5-divisible, and (ii) before the final step of the procedure, the agents agree that each member of the current partition is 1-divisible. By Proposition 1 and Theorem MK, each cut in the procedure successfully constructs a set with the given target value; thus the procedure in our setting successfully constructs an envy-free allocation as in the perfectly divisible cake setting, as desired. ∎

References

Aziz, Haris and Simon Mackenzie (2016). "A discrete and bounded envy-free cake cutting protocol for any number of agents." In: *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, 416-427.

Becker, Robert (1980). "On the long-run steady state in a simple dynamic model of equilibrium with heterogeneous households." *The Quarterly Journal of Economics* 95, 375-382.

Caragiannis, Ioannis, David Kurokawa, Hervé Moulin, Ariel Procaccia, Nisarg Shah, and Junxing Wang (2019). "The Unreasonable Fairness of Maximum Nash Welfare." *ACM Transactions on Economics and Computation* 7, Article 12, 1-32.

Chambers, Christopher and Federico Echenique (2018). "On multiple discount rates." *Econometrica* 86, 1325-1346.

Chen, Bo and Satoru Takahashi (2012). "A folk theorem for repeated games with unequal discounting." *Games and Economic Behavior* 76, 571-581.

Cseh, Agnes and Tamas Fleiner (2020). "The complexity of cake cutting with unequal shares." *ACM Transactions on Algorithms* 16, Article 29.

Edmonds, Jeff and Kirk Pruhs (2006). "Cake cutting really is not a piece of cake." In: *Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm (SODA '06)*, 271-278.

Even, Shimon and Azaria Paz (1984). "A note on cake-cutting." In: *Discrete Applied Mathematics* 7, 285-296.

Foley, Duncan (1967). "Resource allocation and the public sector." *Yale Economic Essays* 7, 45-98.

Fudenberg, Drew and David Levine (1989). "Reputation and equilibrium selection in games with a patient player." *Econometrica* 57, 759-778.

Fudenberg, Drew and Eric Maskin (1991). "On the dispensability of public randomization in discounted repeated games." *Journal of Economic Theory* 53, 428-438.

Jackson, Matthew and Leeat Yariv (2015). "Collective dynamic choice: the necessity of time inconsistency." *American Economic Journal: Microeconomics* 7, 150-178.

Kakeya, Sōichi (1914). "On the set of partial sums of an infinite series." *Proceedings of the Tokyo Mathematico-Physical Society* 7, 250-251.

Kakeya, Sōichi (1915). "On the partial sums of an infinite series." *Science reports of the Tôhoku Imperial University* 3, 159-163.

Kochov, Asen and Yangwei Song (2023). "Intertemporal hedging and trade in repeated games with recursive utility." *Econometrica* 91, 2333-2369.

Lehrer, Ehud and Ady Pauzner (1999). "Repeated games with differential time preferences." *Econometrica* 67, 393-412.

Mackenzie, Andrew and Vilmos Komornik (2023). "Fairly taking turns." Forthcoming, *Games and Economic Behavior.*

Procaccia, Ariel (2009). "Thou shalt covet thy neighbor's cake." In: *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI'09)*, 239-244.

Rader, Trout (1981). "Utility over time: the homothetic case." *Journal of Economic Theory* 25, 219-236.

Ramsey, Frank (1928). "A mathematical theory of saving." *The Economic Journal* 38, 543-559.

Rényi, Alfréd (1957). "Representations for real numbers and their ergodic properties." *Acta Mathematica Hungarica* 8, 477-493.

Robertson, Jack and William Webb (1998). *Cake-Cutting Algorithms: Be Fair If You Can.* Boca Raton, Florida: CRC Press.

Rubinstein, Ariel (1982). "Perfect equilibrium in a bargaining model." *Econometrica* 50, 97-109.

Salonen, Hannu and Hannu Vartiainen (2008). "Valuating payoff streams under unequal discount factors." *Economics Letters* 99, 595-598.

Sorin, Sylvain (1986). "On repeated games with complete information." *Mathematics of Operations Research* 11, 147-160.

Steinhaus, Hugo (1948). "The problem of fair division." *Econometrica* 16, 101-104.

Sugaya, Takuo (2015). "Characterizing the limit set of perfect and public equilibrium payoffs with unequal discounting." *Theoretical Economics* 10, 691-717.

Tinbergen, Jan (1946). *Redelijke Inkomensverdeling* [*Reasonable Income*, in Dutch]. Haarlem, the Netherlands: De Gulden Pers.

Weitzman, Martin (2001). "Gamma discounting." *American Economic Review* 91, 260-271.