Fairly taking turns

Andrew Mackenzie* and Vilmos Komornik† ‡

This draft: August 12, 2023

Abstract

We investigate the fair division of a sequence of time slots when each agent is sufficiently patient. If agents have identical preferences, then we construct *perfectly equitable* and *efficient* allocations. Otherwise, (i) if there are two agents, then we construct *envy-free* allocations, (ii) if there are three agents, then we construct *proportional* allocations, and (iii) in general, we construct approximately fair allocations. Finally, we investigate achieving approximate fairness at each time period, strategy-proofness, and a notion of computational simplicity.

Keywords: fair division, intertemporal choice

JEL Codes: D63, D71

1 Introduction

1.1 Overview

We consider an idealized model of one of the simplest methods of achieving fairness: taking turns. Everyday examples include children sharing the use of a toy, divorced parents sharing the custody of a child, and friends rotating who chooses the next group activity. Unlike other prominent methods of achieving fairness, taking turns does not require a physical good that can be divided into parts, a currency with which one can compensate another, or a randomization device; instead, time is divided. Remarkably, though time is neither infinitely divisible nor homogeneous in our model, there are often fair allocations.

In our model, the natural numbers $\{1, 2, ...\}$ are *time slots* that are to be partitioned into *schedules*, one for each of n agents. Each agent i has preferences over schedules that may be represented by a utility function $u_i : 2^{\mathbb{N}} \to [0, 1]$, a countably additive probability

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

†College of Mathematics and Computational Science, Shenzen University, Shenzen, China and Département de Mathématique, Université de Strasbourg, Strasbourg, France. Email: vilmos.komornik@math.unistra.fr

‡The work of V.K. was partially supported by the grant No. 11871348 of the National Natural Science Foundation of China (NSFC). We thank two anonymous referees, Steven Brams, Jens Gudmundsson, Ehud Lehrer, Mihai Manea, William Thomson, Christian Trudeau, Yu Zhou, Bill Zwicker, and participants at the 2023 Society for the Advancement of Economic Theory (SAET) Conference.

measure in a large domain which notably includes geometric discounting with respect to a personal discount factor $\delta_i \in (0, 1)$:

$$\text{for each } S \subseteq \mathbb{N},\ u_i(S) = (1 - \delta_i) \sum_{t \in S} \delta_i^{t-1}.$$

Crucially, though most economic analysis assumes that agents share a discount factor (say, equal to a market interest rate), we allow for agents to have *personal* discount factors, which is natural when (i) payoffs are not monetary, due to differences in intertemporal preferences, or (ii) payoffs are monetary, but agents have access to loans at different interest rates.[1] Previous work has emphasized that discount factor differences have important implications for 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 endogenous discounting (Kochov and Song, 2023); we investigate the implications of personal discount factors for fair division.

Our main results allow for a particularly large domain of utility functions because stationarity, one of the defining behavioral features of geometric discounting (Koopmans, 1960), plays no role in any of our arguments. By a classic theorem (Kakeya, 1914; Kakeya, 1915), if a countably additive probability measure $u_i : 2^{\mathbb{N}} \to [0, 1]$ is monotonic with respect to earliness, then its range is the unit interval if and only if for each $t \in \mathbb{N}$, we have $u_i(\{t + 1, t + 2, ...\}) \geq u_i(\{t\})$. Building on this condition, we say that for each $k \in [0, \infty)$, u_i is a *kth-order Kakeya utility function* if

$$\text{for each } t \in \mathbb{N}, u_i(\{t + 1, t + 2, ...\}) \geq k u_i(\{t\}),$$

and we let $\mathcal{U}_k$ denote the class of these functions. Notice that if u_i is geometric discounting with respect to discount factor $\delta_i \in (0, 1)$, then for each $k \in [0, \infty)$, $u_i \in \mathcal{U}_k$ if and only if $\delta_i \geq \frac{k}{k+1}$. More generally, we can compare the relative patience of two agents by the order of their Kakeya utility functions: if $u_i \in \mathcal{U}_k$ and $u_j \in \mathcal{U}_{k'} \backslash \mathcal{U}_k$ for $k > k'$, then i is more patient in the sense that he always measures the relative value of $\{t + 1, t + 2, ...\}$ to $\{t\}$ to be at least k while j does not always do so.

Our main message is that there are fair allocations if agents are sufficiently patient, which many of our results articulate using three classic notions of exact fairness. The strongest, which we refer to as *perfect equity*, requires that every agent measures every schedule to be worth exactly $\frac{1}{n}$ (Steinhaus, 1949; Dubins and Spanier, 1961). We also consider the weaker notion of *no-envy*, which requires that no agent measures another's schedule to be worth more than his own (Tinbergen, 1946; Foley, 1967). Finally, the weakest classic notion we consider is *proportionality*, which requires that every agent measures his schedule to be worth at least $\frac{1}{n}$ (Steinhaus, 1948).

For our most general results where even *proportionality* proves elusive, we turn to three notions of approximate fairness. First, for any margin of error $\varepsilon > 0$, we approximate *perfect equity* using *ε-perfect equity*. Second, we consider two notions that have been focal in the computer science literature: (i) *envy-freeness up to one good*, written as *EF1* (Budish, 2011), and (ii) *envy-freeness up to any good*, written as *EFX* (Caragiannis, Kurokawa, Moulin, Procaccia, Shah, and Wang, 2019). The former allows i to envy j if

[1] Lehrer and Pauzner (1999) give the example of an employer who can borrow at a lower interest rate than an employee.

the envy can be eliminated by removing *some* time slot from j, while the latter allows i to envy j if the envy can be eliminated by removing *any* time slot from j.

Finally, we complement our main message with further insights involving efficiency, fairness throughout the procedure, incentive compatibility, and computational simplicity. In particular, we consider (i) standard *(Pareto) efficiency*, (ii) *eternal* approximate fairness, the requirement that at each point in time we have approximate fairness thus far, (iii) *strategy-proofness*, the requirement that it is a dominant strategy for each agent to honestly report his preferences, and (iv) *myopia*, the requirement that we can always assign the first t time slots using only the utilities for these time slots.

For most of our results, the agents share a domain of admissible utility functions. Notice that there may not even be *proportional* allocations if we allow Kakeya utility functions with order less than $n-1$, as in this case the first time slot alone may be worth more than $\frac{1}{n}$ to everyone; thus we typically require utility functions to belong to $\mathcal{U}_{n-1}$, sometimes with further restrictions. Altogether, we establish the following:

Identical preferences. For identical preferences, if n agents have a common utility function $u_0 \in \mathcal{U}_{n-1}$, then a *perfectly equitable* and *efficient* allocation can be constructed using our Constrained Priorities procedure, which as a special case includes our Iterative Greedy Algorithm procedure (Theorem 1). This result applies whenever n agents share a discount factor $\delta_0 \in [\frac{n-1}{n}, 1)$, which also follows from a classic result for repeated games (Sorin, 1986; Fudenberg and Maskin, 1991); see the discussion in Section 1.2. We also show that under further domain restrictions, the Fudenberg-Maskin procedure (adapted from repeated games for fair division) constructs such an allocation that is moreover *eternally EFX* and *eternally ε-perfectly equitable* (Theorem 2).

Two agents. For $n \geq 2$ agents who may have distinct preferences in $\mathcal{U}_{n-1}$, there may not be any *perfectly equitable* allocations (Example 1), and the promising Iterative Greedy Algorithm may not deliver *proportional* allocations (Example 2). That said, for $n = 2$, if at least one of the agents has a utility function in $\mathcal{U}_1$, then an *envy-free* allocation can be constructed using our version of the classic Divide and Choose procedure (Theorem 3). We also show that even under further domain restrictions, no mechanism that always selects *proportional* allocations is *strategy-proof* (Theorem 4) or *myopic* (Theorem 5).

Three agents. For $n \geq 3$ agents who may have distinct preferences in $\mathcal{U}_{n-1}$, the existence of *envy-free* allocations is an open question. That said, for $n = 3$, if each agent i has a utility function $u_i \in \mathcal{U}_2$, then a *proportional* allocation can be constructed using either our Iterative Apportionment procedure (Theorem 6) or our Simultaneous Apportionment procedure (Theorem 7).

General case. For the general case where n agents may have distinct utility functions in $\mathcal{U}_{n-1}$, the existence of *proportional* allocations is an open question. That said, for each $\varepsilon > 0$, if each agent i has a *monotonic* utility function $u_i \in \mathcal{U}_{\frac{1-\varepsilon}{\varepsilon}}$, then an allocation that is *ε-perfectly equitable*, *EF1*, *eternally ε-perfectly equitable*, and *eternally EF1* can be constructed using a procedure from the literature (Caragiannis, Kurokawa, Moulin, Procaccia, Shah, and Wang, 2019) that we call Round-Robin (Theorem 8). Moreover, by relaxing exact fairness to allow approximate fairness, we escape our earlier impossibility results: so long as we require that all utility functions are monotonic, Round-Robin is both *strategy-proof* and *myopic* (Theorem 9). Finally, if all utility functions are monotonic, then an *EFX* allocation can be constructed using a procedure from the literature (Lipton, Markakis, Mossel, and Saberi, 2004; Plaut and Roughgarden, 2020) that we call Envy Graph (Theorem 10).

1.2 Related literature

For this discussion, we consider three categories of fair division models:

- *divisible cake* models, where a heterogeneous and infinitely divisible cake is to be partitioned and utility functions are atomless probability measures;
- *object* models, where a (usually finite) collection of indivisible objects is to be partitioned and utility functions are probability measures; and
- *classical exchange* models, including the textbook model of consumer choice, where there is a (usually finite) collection of homogeneous and infinitely divisible goods.

Our model is an object model with a preference restriction that gives it some, but not all, of the features of a divisible cake model. If we allowed schedules to be randomly assigned, we would have a classical exchange model where time slot probabilities are goods, and even without randomization, the notion of competitive equilibrium from classical exchange models offers an intriguing direction for future work. We use these three model categories to discuss the related literature for our main axioms.

Perfect equity. Our strongest fairness notion implies the rest of the fairness notions that we consider, as well as others such as *egalitarian equivalence* (Pazner and Schmeidler, 1978). Divisible cakes always have *perfectly equitable* allocations, and in fact this fairness notion is so strong that it is only compatible with *efficiency* when preferences are identical (Dubins and Spanier, 1961).[2]

Unfortunately, in our model, there need not be *perfectly equitable* allocations even when $n = 2$ (see Example 1). That said, when agents share a discount factor, the existence of *perfectly equitable* allocations follows from a classic result for repeated games: if the common discount factor is sufficiently high given the number of action profiles in the stage game, then any convex combination of stage game payoffs can be achieved *without randomization* by selecting action profiles with appropriate frequencies (Sorin, 1986; Fudenberg and Maskin, 1991).[3] Interestingly, the constructive proof of Fudenberg and Maskin (1991) involves a procedure that generates a sequence of action profiles such that each action profile achieves its target discount weight, and the fair division result

[2] The standard proof of existence appeals to Lyapunov's classic result that if $\mu : \Sigma \to \mathbb{R}^n$ is an atomless, bounded, and countably-additive vector-valued measure, then its range is compact and convex (Liapounoff, 1940). It follows directly that there is a slice that all agents agree is worth $\frac{1}{n}$, and iterating it follows that there is a *perfectly equitable* allocation. Unfortunately, this proof is not constructive.

[3] Indeed, simply apply the classic result to any repeated game (with the given common discount factor) whose stage game is such that a fixed agent selects any agent to receive 1 while the others receive 0. We remark that Sorin (1986) proves a slightly different result (with a proof that makes the cited result clear): if the common discount factor is sufficiently high given the number of *agents*, then any convex combination of stage game payoffs can be achieved by selecting *mixed* action profiles with appropriate frequencies.

also follows from simply reinterpreting the action profiles as agents and the target discount weights as target utilities. Translated into fair division, the FUDENBERG-MASKIN procedure iteratively selects an agent whose cumulative utility is lowest.

We show that the existence of *perfectly equitable* allocations extends to the case of identical Kakeya utility functions using our CONSTRAINED PRIORITIES procedure, which is more flexible than the FUDENBERG-MASKIN procedure in that each time slot can be assigned to *any* agent who has not yet received his target utility.[4] As a special case, this includes the ITERATIVE GREEDY ALGORITHM, or the iterative application of Rényi's GREEDY ALGORITHM (Rényi, 1957), which has recently been applied in the context of decision theory due to its remarkable continuity properties (Mackenzie, 2019). We remark that with identical preferences, any *perfectly equitable* allocation is (i) a *competitive equilibrium from equal incomes* outcome (Kolm, 1971; Varian, 1974) when each agent is endowed with his schedule and the price of each time slot is its common value, and (ii) an allocation that maximizes the *Nash product*, or the product of the agents' utilities (Nash, 1950); these properties remain focal in the fair division literature for models where *perfectly equitable* allocations are not available.

NO ENVY. For divisible cakes, there are at least five procedures for constructing *envy-free* allocations when $n = 3$,[5] and at least two for the general case,[6] but unfortunately none of these extends to our model—largely because we cannot 'trim' arbitrary sets. For classical exchange, there are competitive equilibria (McKenzie, 1954; Arrow and Debreu, 1954), those for which the endowment is equal division are *envy-free* (Kolm, 1971; Varian, 1974), and when preferences are linear these allocations maximize the Nash product (Eisenberg and Gale, 1959). For object models, there need not be competitive equilibria, but a natural kind of approximate competitive equilibrium always exists; that said, its approximation error is unbounded as the number of objects grows (Budish, 2011).

In our model, if $n = 2$, then we can construct an *envy-free* allocation using our version of DIVIDE AND CHOOSE, the ancient procedure dating back to at least the eighth or seventh century BC (Lowry, 1987; Brams, Taylor, and Zwicker, 1995).[7] Unfortunately, this procedure does not generalize for $n \geq 3$, and the existence of *envy-free* allocations for the general case remains an open question.

PROPORTIONALITY. The literature on cake division began with the STEINHAUS procedure for constructing *proportional* allocations when $n = 3$ and the more general BANACH-KNASTER procedure for constructing *proportional* allocations when n is arbitrary (Steinhaus, 1948). As with the *envy-free* procedures for cake division, these procedures do not work in our model because we cannot 'trim' arbitrary sets.

When preferences are not identical, the FUDENBERG-MASKIN procedure need not

[4]The constructive proof of Sorin (1986) involves a procedure of intermediate flexibility: translated into fair division, the SORIN procedure iteratively selects an agent who still requires at least $\frac{1}{n}$ of the remaining utility.

[5]In particular, we have the STROMQUIST procedure (Stromquist, 1980), the LEVMORE-COOK procedure (Levmore and Cook, 1981), the WEBB procedure (see Brams, Taylor, and Zwicker, 1995), the BRAMS-TAYLOR-ZWICKER procedure (Brams, Taylor, and Zwicker, 1995), and an application of the AUSTIN procedure for constructing a *perfectly equitable* partition when there are two agents (Austin, 1982). See Brams, Taylor, and Zwicker (1995) for an overview.

[6]In particular, we have the BRAMS-TAYLOR procedure (Brams and Taylor, 1995) and the AZIZ-MACKENZIE procedure (Aziz and Mackenzie, 2016).

[7]In Hesiod's *Theogony*, Prometheus divides an ox into two portions, then Zeus chooses.

yield a *proportional* outcome.[8] Similarly, the Iterative Greedy Algorithm need not yield a *proportional* outcome—even if we let more impatient agents construct their schedules earlier in the procedure (Example 2). That said, we introduce two procedures for constructing *proportional* allocations when $n = 3$ (Iterative Apportionment and Simultaneous Apportionment). The existence of *proportional* allocations for the general case remains an open question.

Approximate fairness. For the general case, a variety of technical issues arise when agents can have different discount factors—for example, the folk theorem for repeated games does not hold (Lehrer and Pauzner, 1999), and even when there are two agents the Pareto frontier can be everywhere discontinuous (Salonen and Vartiainen, 2008). In order to investigate the general case, we consider three notions of approximate fairness: *ε-perfect equity*, *EF1*, and *EFX*. We remark that in our model, if an *EFX* allocation assigns each agent an infinite schedule, then it is *envy-free*, but the two properties are not equivalent in general.

When there are finitely many objects, Nash product maximization delivers an *EF1* allocation that is *efficient* (Caragiannis, Kurokawa, Moulin, Procaccia, Shah, and Wang, 2019), as does a faster algorithm that runs in polynomial time when valuations are bounded (Barman, Krishnamurthy, and Vaish, 2018). Moreover, the Round-Robin procedure delivers an *EF1* allocation (Caragiannis, Kurokawa, Moulin, Procaccia, Shah, and Wang, 2019), and if agents share a common ranking of the objects, the Envy Graph procedure delivers an *EFX* allocation (Lipton, Markakis, Mossel, and Saberi, 2004; Plaut and Roughgarden, 2020). We show that these latter insights extend to our model regardless of the patience of the agents, and that the Round-Robin procedure moreover delivers an *ε-perfectly equitable* outcome if agents are sufficiently patient. Interestingly, when $n = 3$, an *EF1* allocation where each agent consumes an interval can be constructed using a discrete version of the Stromquist procedure (Bilò, Caragiannis, Flammini, Igarashi, Monaco, Peters, Vinci, and Zwicker, 2022); we leave the general investigation of fair allocations where agents consume intervals in our model for future work.

Strategy-proofness.[9] For classical exchange economies, even if agents are restricted to reporting linear preferences, *strategy-proofness* and *efficiency* (i) imply that one agent consumes everything when there are two agents (Schummer, 1997), and (ii) are incompatible with all of our fairness notions when there are more than two agents (Cho and Thomson, 2023). Any result for classical exchange with linear preferences applies directly to cake division (Aziz and Ye, 2014), and moreover the former result remains true even when dividing a circular cake into intervals (Thomson, 2007).

For cake division, early contributions identified positive results if (i) we allow randomization when agents are risk neutral (Mossel and Tamuz, 2010), and (ii) we restrict preferences to a dichotomous marginal valuation domain (Chen, Lai, Parkes, and Procaccia, 2013); for an overview of more recent contributions on the latter topic see Bu,

[8]Indeed, suppose there are two agents who use geometric discounting, one with discount factor 0.51 and one with discount factor 0.99. In the Fudenberg-Maskin procedure, after the impatient agent receives his first time slot (which is either the first time slot or the second time slot, depending on how the tie is broken to assign the first time slot), a long interval of time slots will be assigned to the patient agent while his utility is less than the impatient agent's, after which it will be impossible for the impatient agent to achieve $\frac{1}{2}$.

[9]We remark that this property is usually called *truthfulness* in the fair division literature, while *strategy-proofness* is more common in economics.

Song, and Tao (2023). The question of whether there are *strategy-proof* and *proportional* mechanisms for a natural larger domain was raised in Chen, Lai, Parkes, and Procaccia (2013) and recently answered for the case of two agents: it was first shown that any such mechanism must involve some undesirable waste (Bei, Chen, Huzhang, Tao, and Wu, 2017) and finally shown that there is no such mechanism at all (Bu, Song, and Tao, 2023).

For assigning finitely many objects, *strategy-proofness* and *efficiency* are incompatible with all of our fairness notions (Klaus and Miyagawa, 2001). Moreover, *strategy-proofness* is incompatible with (i) selecting an *envy-free* allocation whenever possible (Lipton, Markakis, Mossel, and Saberi, 2004; see also Caragiannis, Kaklamanis, Kanellopoulos, and Kyropoulou, 2009), (ii) maximizing the minimum utility (Bezáková and Dani, 2005), and (iii) achieving *EF1* without discarding any objects (Amanatidis, Birmpas, Christodoulou, and Markakis, 2017). That said, if we restrict preferences to a dichotomous object valuation domain, then we can simultaneously achieve *group strategy-proofness*, *EF1*, and *efficiency* by maximizing the Nash product (Halpern, Procaccia, Psomas, and Shah, 2020), and in fact these three objectives remain compatible if we relax additivity to submodularity (Babaioff, Ezra, and Feige, 2021; Barman and Verma, 2022).

As observed by Thomson (2007), there are some unique technical challenges for analyzing *strategy-proofness* in fair division models: (i) the restriction to additive utility functions precludes standard arguments for rich preference domains, and (ii) indifference curves can be extremely thick. This latter point is particularly salient in our model: for many utility functions that we consider, each utility $x \in (0, 1)$ is assigned to a continuum of schedules (Erdős, Joó, and Komornik, 1990).[10] Even so, we reinforce the finding from the cake division literature that for two agents, there is no *strategy-proof* and *proportional* mechanism, even on the restricted domain of sufficiently patient geometric preferences.

Before proceeding, we remark that several other problems of fairly taking turns have been previously considered in the literature, such as taking turns winning at a low price through tacit collusion in repeated auctions (Rachmilevitch, 2013), taking turns performing a chore with stochastic private costs in each period (Leo, 2017), taking turns *selecting* objects (Brams and Taylor, 2000), forming a queue using monetary transfers (Dolan, 1978; Chun, 2016), ordering penalty kicks (Brams and Ismail, 2018; Anbarcı, Sun, and Ünver, 2021; Brams, Ismail, and Kilgour, 2023), and ordering tennis serves (Brams, Ismail, Kilgour, and Stromquist, 2018). That said, these problems and their associated models are not closely related to ours.

2 Model

2.1 Environments and economies

In our model, 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. We distinguish between *environments*,

[10] In particular, if u_i is a geometric discounting utility function with discount factor $\delta_i \in (\frac{2}{1+\sqrt{5}}, 1)$, then for each $x \in (0, 1)$, there are a continuum of schedules S such that $u_i(S) = x$ (Theorem 3, Erdős, Joó, and Komornik, 1990). For context, $\frac{2}{3} > \frac{2}{1+\sqrt{5}}$.

where preferences are private information, and *economies*, where preferences are common knowledge.

DEFINITION: An *environment* is specified by a number $n \in \mathbb{N}$ as follows:

- $N \equiv \{1, 2, ..., n\}$ is the set of *agents*.
- $T \equiv \{1, 2, ...\}$ 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.
- $\mathcal{U} \subseteq [0,1]^{\mathcal{S}}$ is the collection of countably additive probability measures on $\mathcal{S}$,[11] which we call the *full domain*. Each $u_0 \in \mathcal{U}$ is a *utility function* representing preferences over schedules. A *utility function profile* $u = (u_i)_{i \in N}$ is a member of $\mathcal{U}^N$.
- $\Pi \subseteq \mathcal{S}^N$ 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) $\cup_{i \in N} \pi_i = T$.

An *economy* (n, u) is an environment and an associated utility function profile. Whenever we refer to an arbitrary environment or economy, we assume all of the above notation.

In an environment, it is common knowledge that the preferences of the agents are given by some utility function profile $u \in \mathcal{U}^N$, but each agent $i \in N$ knows only his own utility function u_i. In an economy, the utility function profile is common knowledge.

2.2 Exact fairness and utility function domains

Most of our analysis concerns economies with complete information rather than environments with incomplete information. In particular, 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

- *perfect equity* if for each pair $i, j \in N$, $u_i(\pi_j) = \frac{1}{n}$;
- *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 *perfect equity* implies *no-envy*, which in turn implies *proportionality*. In general, there are economies where (i) each agent has a utility function in the full domain, and (ii) even *proportional* allocations do not exist—for example, this is the case whenever each agent is extremely impatient in the sense that he assigns all utility weight to the first time slot. That said, our results involve smaller domains which guarantee the existence of fair allocations.

[11]That is, for each $u_0 \in \mathcal{U}$, we have (i) $u_0(T) = 1$, and (ii) for each $S \in \mathcal{S}$, $u_0(S) = \sum_{t \in S} u_0(\{t\})$.

Definition: Fix an environment. A *domain* is a collection of utility functions $\mathcal{D} \subseteq \mathcal{U}$. We define the domains $(\mathcal{U}_k)_{k\in[0,\infty)}$, $\mathcal{U}_{\mathrm{M}}$, and $\mathcal{U}_{\mathrm{G}}$ as follows: for each $u_0 \in \mathcal{U}$ and each $k \in [0, \infty)$, we say that

- $u_0 \in \mathcal{U}_k$ if for each $t \in T$, we have $u_0(\{t+1, t+2, ...\}) \geq k u_0(\{t\})$, in which case we say that u_0 is *kth-order Kakeya*;
- $u_0 \in \mathcal{U}_{\mathrm{M}}$ if $u_0(\{1\}) \geq u_0(\{2\}) \geq ...$, in which case we say that u_0 is *monotonic*; and
- $u_0 \in \mathcal{U}_{\mathrm{G}}$ if there is $\delta_0 \in (0, 1)$ such that for each $t \in T$, we have $u_0(\{t\}) = (1-\delta_0)\delta_0^{t-1}$, in which case we say that u_0 is *geometric* and that δ_0 is the *discount factor (for u_0)*.

We remark that there are axiomatic foundations for Kakeya utility functions.[12] As discussed in the introduction, an agent with a higher-order Kakeya utility function is more patient, and many of our results concern economies for which all agents are sufficiently patient given the number of agents.

2.3 Approximate fairness and eternal approximate fairness

For our most general results, exact fairness proves elusive and thus we turn instead to approximate fairness. We consider notions where an allocation might be declared 'fair enough' if an undesired utility measurement is within a specified margin of error, if an undesired comparison reverses after the removal of *some* time slot, or if an undesired comparison reverses after the removal of *any* time slot.

Definition: Fix an economy, let $\pi \in \Pi$, and let $\varepsilon > 0$. We say that π is

- *ε-perfectly equitable* if for each triple $i, j, j' \in N$, $u_i(\pi_j) \in [u_i(\pi_{j'}) - \varepsilon, u_i(\pi_{j'}) + \varepsilon]$;
- *envy-free up to one object (EF1)* if for each pair $i, j \in N$ such that $\pi_j \neq \emptyset$, there is $t \in \pi_j$ such that $u_i(\pi_i) \geq u_i(\pi_j \backslash \{t\})$; and
- *envy-free up to any object (EFX)* if for each pair $i, j \in N$ such that $\pi_j \neq \emptyset$ and each $t \in \pi_j$, $u_i(\pi_i) \geq u_i(\pi_j \backslash \{t\})$.

While our results only concern the approximate fairness notions defined above, the literature has also considered others; we mention several prominent alternatives in our concluding discussion (Section 4). We use the above notions to articulate both (i) when an allocation is approximately fair, and (ii) when an allocation is *eternally* approximately fair, in the sense that at each point in time, the assignment of the time slots thus far is approximately fair.

[12] In particular, previous authors have investigated when a preference relation over events in a σ-algebra may be represented by a probability measure, usually in the context of representing beliefs (Bernstein, 1917; de Finetti, 1937; Koopman, 1940; Savage, 1954). Our work is most closely related to Mackenzie (2019), who proved that under standard axioms, *monotone continuity* (Villegas, 1964) and an ordinal analogue of our third-order Kakeya condition are together sufficient to guarantee a unique such representation. Without the third-order Kakeya condition, a probability measure representation can be guaranteed with the necessary and sufficient conditions of Chateauneuf (1985). For geometric discounting, a discount factor that is at least $\frac{1}{2}$ must be unique even in our discrete model (Kettering and Kochov, 2021).

Definition: Fix an economy, let $\pi \in \Pi$, and let $\varepsilon > 0$. For each $i \in N$ and each $t \in \{0\} \cup T$, define the *partial schedule for i at t* by $\pi_i\restriction_t \equiv \pi_i \cap \{1, 2, ..., t\}$. Moreover, define the *partial allocation at t* by $\pi\restriction_t \equiv (\pi_i\restriction_t)_{i\in N}$. We say that π is

- *eternally ε-perfectly equitable* if for each $t \in T$ and each triple $i, j, j' \in N$, $u_i(\pi_j\restriction_t) \in [u_i(\pi_{j'}\restriction_t) - \varepsilon, u_i(\pi_{j'}\restriction_t) + \varepsilon]$;
- *eternally EF1* if for each $t \in T$ and each pair $i, j \in N$ such that $\pi_j\restriction_t \neq \emptyset$, there is $t' \in \pi_j\restriction_t$ such that $u_i(\pi_i\restriction_t) \geq u_i(\pi_j\restriction_t \backslash \{t'\})$; and
- *eternally EFX* if for each $t \in T$, each pair $i, j \in N$ such that $\pi_j\restriction_t \neq \emptyset$, and each $t' \in \pi_j\restriction_t$, $u_i(\pi_i\restriction_t) \geq u_i(\pi_j\restriction_t \backslash \{t'\})$.

In other words, each condition requires that for each $t \in T$, $\pi\restriction_t$ must be approximately fair, where approximate fairness is articulated using *ε-perfect equity*, *EF1*, and *EFX*, respectively. This verbal description serves to formally define these approximate fairness notions for partial allocations. We remark that we define partial schedules and partial allocations at 0 simply to facilitate inductive arguments.

2.4 Properties for mechanisms

Our analysis of complete information economies yields many statements with the following format: if there are n agents and each has a utility function in $\mathcal{D}$, then there are fair allocations. In this context, we use the term *procedure* to mean an algorithm that (i) takes as input an economy (n, u) with $u \in \mathcal{D}^N$, and then (ii) performs a sequence of operations in order to construct a fair allocation π as output.

Given such a result, we can further investigate environments where it is common knowledge that the utility function profile belongs to $\mathcal{D}^N$, but where each agent knows only his own utility function. In particular, we investigate direct mechanisms that (i) ask each agent to report his utility function, and then (ii) return an allocation that is fair according to the reports. In this context, we use the term *mechanism* to mean a function that maps each reported utility function profile in $\mathcal{D}^N$ to an allocation. We emphasize that while each procedure yields an associated mechanism—in particular, the mechanism that feeds each reported utility function profile as input to the procedure in order to return the procedure's output—we deliberately avoid using the terms *procedure* and *mechanism* interchangeably in order to avoid confusion.

Our investigation of mechanisms involves both incentive compatibility and computational simplicity.

Definition: Fix an environment and a domain $\mathcal{D}$. A *mechanism* is a function $\mathcal{M} : \mathcal{D}^N \to \Pi$ that comes with the following associated notation: (i) for each $i \in N$, we let $\mathcal{M}_i(u) \in \mathcal{S}$ denote the schedule assigned to i in allocation $\mathcal{M}(u)$, and (ii) for each $t \in T$, we let $\mathcal{M}(t|u) \in N$ denote the agent who consumes time slot t in allocation $\mathcal{M}(u)$. We say that a mechanism $\mathcal{M}$ is

- *strategy-proof* if for each $i \in N$, each $u \in \mathcal{D}^N$, and each $u'_i \in \mathcal{D}$,

$$u_i(\mathcal{M}_i(u)) \geq u_i(\mathcal{M}_i(u'_i, u_{-i})); \text{ and}$$

- *myopic* if for each $t \in T$ and each pair $u, u' \in \mathcal{D}^N$ such that for each $i \in N$ and each $t' \in \{1, 2, ..., t\}$ we have $u_i(\{t'\}) = u'_i(\{t'\})$,

$$\mathcal{M}(t|u) = \mathcal{M}(t|u').$$

Strategy-proofness requires that in the game form where each agent strategically reports his utility function, honesty is always a dominant strategy for each agent. Myopia requires that for each t, we can always assign the first t time slots using only the utilities for these time slots. Notice that myopia has no bite on any subdomain of $\mathcal{U}_{\mathrm{G}}$, as in this case $u_0(\{1\})$ determines u_0; thus this property only has substance on richer domains.

3 Results

3.1 The single-agent technique

Before proceeding to our main results for fair division, it is useful to first consider a simple problem that a single agent in an economy might try to solve without input from the others. In particular, suppose agent i is given 'source' schedule S and target value $v \in [0, u_i(S)]$, then asked to construct schedule $S^* \subseteq S$ such that $u_i(S^*) = v$. Under what conditions can he succeed?

In the special case that $S = T$ and u_i is geometric discounting with discount factor $\frac{1}{2}$, the binary expansion of v provides an easy solution—for example, for $v = \frac{5}{8}$, the binary expansion is 0.101 and the schedule $S^* = \{1, 3\}$ is a solution. Notice that the binary expansion for v is constructed using a greedy algorithm: write the digits in sequence, and write 1 for the next digit if and only if the written value will not exceed v. The natural generalization of this procedure was notably investigated in the context of representations more general than binary expansions (Rényi, 1957), and is an important tool for our single-agent problem.

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,[13] 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.

This procedure may fail to deliver a schedule with the desired value even when there is a solution—for example, if $S = \{1, 2, 3\}$, $u_i(\{1\}) = 0.3$, $u_i(\{2\}) = 0.2$, $u_i(\{3\}) = 0.15$, and $v = 0.35$, then the Greedy Algorithm will construct $\{1\}$, whose value is not v, even though there is solution $\{2, 3\}$. That said, this procedure always succeeds if the source event is sufficiently divisible in the following manner:

Definition: Fix an economy. For each $i \in N$, each $k \in \{0, 1, ...\}$, and each $S \in \mathcal{S}$, we say that S is *k-divisible for i* if for each $t \in S$, $u_i(\{t' \in S | t' > t\}) \geq k u_i(\{t\})$.

[13] At a high level, each of our procedures can be understood in terms of baskets and flags, such that baskets receive time slots and at each step a basket represents the schedule of time slots it has thus far received, with agents placing flags in baskets to determine which basket receives a given time slot or which agent receives a given basket. We describe each of our procedures in this manner.

In particular, it follows from a classic theorem (Kakeya, 1914; Kakeya, 1915) that if preferences are monotonic with respect to earliness, then the agent can construct a subset of source schedule S with any target value in $[0, u_i(S)]$ if and only if he considers S to be 1-divisible, in which case a solution can be constructed using the GREEDY ALGORITHM; see also Volume I, Part One, Problem 131 of Pólya and Szegö (1925) and Jones (2011).

Our first result generalizes these findings; it implies that if S is k-divisible, then the agent can iteratively solve this problem k times, provided the targets $\{v_1, v_2, ..., v_k\}$ satisfy $\sum v_k \leq u_i(S)$. We remark that an ordinal analogue of this technique was applied in the context of decision theory (Mackenzie, 2019).

PROPOSITION 1: 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
- $S \backslash S^*$ is $(k-1)$-divisible for i.

The proof is in Appendix 1.

3.2 Identical preferences

We begin our investigation of fair allocations with the case where n agents have a common utility function $u_0 \in \mathcal{U}_{n-1}$. In this case all allocations are *efficient*, and it follows easily from Proposition 1 that there are *perfectly equitable* allocations. Indeed, consider the ITERATIVE GREEDY ALGORITHM, where in sequence each agent constructs his own schedule: agent 1 takes T as source schedule and $\frac{1}{n}$ as target value to construct π_1, then agent 2 takes $T \backslash \pi_1$ as source schedule and $\frac{1}{n}$ as target value to construct π_2, and so on for the first $n-1$ agents, with the remaining time slots finally assigned to agent n. By Proposition 1, the resulting allocation is *perfectly equitable*.

Remarkably, the agents need not be nearly so systematic to successfully construct a *perfectly equitable* allocation: our more general CONSTRAINED PRIORITIES procedure works as well. For the given economy, the procedure is initialized by specifying a *priority profile*, or a priority order over the agents for each time slot t, in any manner. To begin, there are n baskets, and each is immediately assigned to an agent. The time slots are assigned to baskets in sequence, and when assigning time slot t, a flag is placed in the basket of agent i if and only if the basket's value (according to the common utility function) would exceed $\frac{1}{n}$ should it receive t. Time slot t is assigned to the basket of the agent whose priority for t is highest, under the constraint that it should not be assigned to a basket with a flag, and then all flags are removed and we proceed to the next time slot. In order to show that the procedure is well-defined, we prove that there is always at least one basket with no flag.

Observe that CONSTRAINED PRIORITIES coincides with the ITERATIVE GREEDY ALGORITHM when the priority order of every time slot orders the agents by index. That said, the CONSTRAINED PRIORITIES formulation has an important practical advantage: after the first t time slots are assigned and consumed, it is easy to calculate which agent should receive time slot $t+1$ before it is too late for $t+1$ to be consumed.

Our first main result states that no matter how priorities are set, the associated CONSTRAINED PRIORITIES allocation is *perfectly equitable*. For readability, we follow the presentation style used for the Deferred Acceptance procedure (Gale and Shapley, 1962), in that we formally define Constrained Priorities implicitly by its associated proof, omitting a separate explicit (and notation-heavy) definition. We also pursue this presentation style for the other procedures associated with our main results.

Theorem 1: Fix an economy. If there is $u_0 \in \mathcal{U}_{n-1}$ such that for each $i \in N$, $u_i = u_0$, then for any priority profile, the Constrained Priorities procedure constructs a *perfectly equitable* and *efficient* allocation.

The proof is in Appendix 2. As we claimed earlier, any Constrained Priorities outcome is a competitive equilibrium outcome when each agent is endowed with his schedule and the price of each time slot is its common value, and it maximizes the Nash product. Indeed, the first claim is true because each agent's income is $\frac{1}{n}$ and the price of any schedule is its utility; the second claim is true because $(\frac{1}{n})^n$ maximizes the Nash product on the simplex with n vertices, which is the set of feasible utility vectors. Though we know little about competitive equilibria and Nash product maximization when preferences are not identical in our model, these positive results for the case of identical preferences suggest that these might be promising directions for future research.

We conclude this section by establishing that under the additional assumption that preferences are monotonic, we can moreover achieve eternal approximate fairness by custom-tailoring the priorities to the given economy. In particular, if we generate the priorities as we assign the time slots by always prioritizing an agent whose cumulative utility is lowest, then the result is the Fudenberg-Maskin procedure adapted from repeated games for fair division (Fudenberg and Maskin, 1991), and this procedure performs extremely well.

Theorem 2: Fix an economy and let $\varepsilon \in (0, 1]$. If there is $u_0 \in \mathcal{U}_{n-1} \cap \mathcal{U}_{\mathrm{M}}$ such that for each $i \in N$, $u_i = u_0$, then the Fudenberg-Maskin procedure constructs a *perfectly equitable* and *efficient* allocation that is *eternally EFX*. If, moreover, $u_0 \in \mathcal{U}_{\frac{1-\varepsilon}{\varepsilon}}$, then this allocation is moreover *eternally ε-perfectly equitable*.

The proof is in Appendix 3. We remark that our positive results for this section have relied on the fact that the common utility function is common knowledge: we analyze mechanisms that ask agents to report their preferences only *after* this section in order to allow for the possibility that the reported preferences are not identical.

3.3 Two agents

As the case where there is a single agent is a special case of identical preferences, we next consider the case where there are two agents. Unfortunately, there need not be *perfectly equitable* allocations when agents have different preferences:

Example 1: Let $N = \{1, 2\}$, let u_1 be geometric discounting with discount factor $\delta_1 = \frac{1}{2}$, and let u_2 be geometric discounting with discount factor $\delta_2 > \frac{1}{2}$. Then $u_1, u_2 \in \mathcal{U}_1$ but there is no *perfectly equitable* allocation: the impatient agent finds $\{1\}$ and $\{2, 3, ...\}$ to be the only schedules worth $\frac{1}{2}$, while the patient agent finds the former to be less valuable and the latter to be more valuable.

What about *envy-free* allocations? When there are two agents, every *proportional* allocation is *envy-free*, as an agent who believes he has received at least $\frac{1}{2}$ does not envy the other agent. Given its usefulness for constructing *proportional* allocations when preferences are identical, the ITERATIVE GREEDY ALGORITHM offers a particularly promising approach: simply have one agent construct a schedule for himself that he values $\frac{1}{2}$, leaving the rest for the other agent.

Unfortunately, without identical preferences, the ITERATIVE GREEDY ALGORITHM may fail even when there are two agents with geometric utility functions: if the more patient agent goes first, then he may begin by taking a set of consecutive time slots that the impatient agent values more than $\frac{1}{2}$. This seems easy to address by having the agents construct their schedules in order of *impatience*, and at first glance this procedure appears quite promising for *any* number of agents: if schedule S is k-divisible to impatient agent 1 with discount factor δ_1, then it is also k-divisible to patient agent 2 with discount factor $\delta_2 > \delta_1$, as for each $t \in S$,

$$\frac{u_2(\{t' \in S | t' > t\})}{u_2(\{t\})} = \sum_{t' \in S | t' > t} \delta_2^{t'-t} > \sum_{t' \in S | t' > t} \delta_1^{t'-t} = \frac{u_1(\{t' \in S | t' > t\})}{u_1(\{t\})} \geq k.$$

That said, the following example illustrates that the ITERATIVE GREEDY ALGORITHM may fail no matter how the agents are ordered.

EXAMPLE 2: *The* ITERATIVE GREEDY ALGORITHM *does not guarantee proportionality.* Let $N = \{1, 2\}$, let u_1 be geometric discounting with discount factor δ_1 that is the smallest $\delta \in (0, 1)$ such that

$$(1 - \delta) + \left(\frac{\delta + \delta^2 + \delta^4}{\delta + \delta^2 + \delta^3 + \delta^4 + \delta^5}\right) \delta^5 = \frac{1}{2},$$

or $\delta_1 \approx 0.536$, and let u_2 be geometric discounting with discount factor $\delta_2 = 0.95$. Then $u_2(\mathcal{G}_1(\frac{1}{2}|T)) > \frac{1}{2}$ and $u_1(\mathcal{G}_2(\frac{1}{2}|T)) > \frac{1}{2}$, so for both orders of the agents, the ITERATIVE GREEDY ALGORITHM does not generate a *proportional* allocation. See Appendix 4 for a detailed explanation.

Though the ITERATIVE GREEDY ALGORITHM does not provide a general method for constructing *proportional* allocations, when there are two agents we can construct such allocations that are moreover *envy-free* using DIVIDE AND CHOOSE. For our version of this classic procedure, the divider fills one basket using the GREEDY ALGORITHM for source schedule T and target value $\frac{1}{2}$, then fills a second basket with the remaining time slots. The chooser then takes the basket he prefers, leaving the other for the divider. It follows easily from Proposition 1 that the resulting allocation is *envy-free*, and this holds regardless of the patience of the chooser so long as the divider is sufficiently patient:

THEOREM 3: Fix an economy. If $n = 2$, then for each $i \in N$ such that $u_i \in \mathcal{U}_1$, the DIVIDE AND CHOOSE procedure with i as divider constructs an *envy-free* allocation.

We omit the straightforward proof. This result suggests a natural question: under the given hypotheses, are there *envy-free* allocations that are moreover *efficient*? In general, there may be no such allocation in fair division models, such as when partitioning a circular cake into intervals (Thomson, 2007). That said, the existence of such an allocation is in fact guaranteed in our model under the given hypotheses, which follows from a topological observation about *proportional* allocations for *any* number of agents.

Observation 1: Fix an economy. If there is a *proportional* allocation, then there is a *proportional* allocation that is *efficient*.

This observation follows from standard topological arguments; we sketch these arguments but omit the details. First, $\mathcal{S}^N = (2^T)^N$ is compact in the product topology by Tychonoff's theorem, and it is straightforward to show that Π is compact as a closed subset of $\mathcal{S}^N$. Second, a countably additive probability measure $u_i : 2^T \to [0,1]$ is continuous, and thus the associated mapping $u_i^* : \Pi \to [0,1]$ given by $u_i^*(\pi) \equiv u_i(\pi_i)$ is continuous. Altogether, then, the set Π^* of allocations for which each agent receives a utility of at least $\frac{1}{n}$ is compact as a closed subset of Π, and by hypothesis it is nonempty; thus any strictly monotonic and continuous function of utilities (such as the Nash product) is maximized on Π^* by an allocation that is *proportional* and *efficient*.

Under the given hypotheses, Theorem 3 guarantees an *envy-free* allocation. Any such allocation is *proportional*, so by Observation 1 there is an allocation that is *proportional* and *efficient*. Finally, when there are two agents, any *proportional* allocation is *envy-free*. Altogether, we have established the following.

Corollary 1: Fix an economy. If (i) $n = 2$, and (ii) there is $i \in N$ such that $u_i \in \mathcal{U}_1$, then there is an *envy-free* allocation that is *efficient*.

Given that fair allocations exist, could we design useful mechanisms that ask agents to report their preferences in order to implement fair allocations when preferences are private information? Unfortunately, we conclude this section with two negative results about mechanisms. First, even if the designer knows that all utility functions are geometric, any mechanism that always selects *proportional* allocations necessarily incentivizes agents to misreport their preferences.

Theorem 4: Fix an environment. If $n = 2$, then for the domain $\mathcal{U}_1 \cap \mathcal{U}_{\mathrm{G}}$, there is no *strategy-proof* mechanism that always selects *proportional* allocations.

The proof is in Appendix 5. Second, even if the designer knows that all utility functions are monotonic, *proportional* allocations cannot be computed using myopic calculation algorithms.

Theorem 5: Fix an environment. If $n = 2$, then for the domain $\mathcal{U}_1 \cap \mathcal{U}_{\mathrm{M}}$, there is no *myopic* mechanism that always selects *proportional* allocations.

The proof is in Appendix 6.

3.4 Three agents

For three agents, the existence of *envy-free* allocations is an open question. That said, we introduce two procedures for constructing *proportional* allocations.

In Iterative Apportionment, the three agents collaboratively fill and assign a first basket, then the remaining agents collaboratively fill and assign a second basket, and finally the third agent receives the rest of the time slots. When filling the first basket, the time slots are considered in sequence, and at time slot t each agent places a flag if and only if he believes that the basket's value would exceed $\frac{1}{3}$ should it receive t. If there are no flags, then t is added; if there is one flag, then t is added and the agent who placed the flag receives the basket; if there are multiple flags, then t is skipped. The process is the same when filling the second basket, except that the time slots in the first basket are not considered.

In Simultaneous Apportionment, the three agents begin simultaneously filling three baskets. The time slots are considered in sequence, and at time slot t each agent places a flag in any basket whose value he believes would exceed $\frac{1}{3}$ should it receive t; the time slot is then placed in any basket with a minimal number of flags. If this basket had no flags, then all agents proceed to $t+1$, while if it had at least one flag, then an agent who placed one of its flags immediately receives the basket and exits; there are then two cases. If the taken basket had only one flag, then the two remaining agents continue as before with the remaining two baskets. Otherwise, we prove that the taken basket only had two flags; let i^* denote the remaining agent who placed one of those flags. In this case, we prove that i^* necessarily placed one other flag, and this basket is assigned to him, though it is at this point only partially filled. To conclude, i^* divides the remaining time slots into two parts he considers equal using the Greedy Algorithm, the final agent chooses the part that he prefers, and i^* receives his basket with the unchosen part while the final agent receives the final basket with his chosen part.

Each of these procedures is similar in spirit to a classic procedure for fairly dividing the unit interval when there are no atoms. First, the Banach-Knaster procedure involves iterative apportionment: the n agents iteratively construct and assign the smallest interval that includes the leftmost remaining point and that some agent values $\frac{1}{n}$ (Steinhaus, 1949). Second, the Stromquist procedure involves simultaneous apportionment: (i) for each leftmost interval, each of the three agents specifies the cut point at which he would be indifferent between the middle interval and the rightmost interval, and the median cut point determines a candidate partition, and (ii) the agents simultaneously consider the candidate partitions as the leftmost interval grows until one of the agents declares that he is indifferent between the leftmost interval and every interval in the candidate partition whose closure includes his cut point, at which point the candidate partition is compatible with an *envy-free* allocation (Stromquist, 1980). Notice, however, that both procedures rely heavily on the fact that the family of intervals $\{[0,x]\}_{x\in[0,1]}$ has the following properties: (i) the agents agree on how to rank the members of this family, and (ii) each agent's utility function maps this family onto the unit interval. Neither of our procedures involve such a family of schedules, and yet nevertheless both construct *proportional* allocations.

Theorem 6: Fix an economy. If (i) $n=3$, and (ii) for each $i \in N$, $u_i \in \mathcal{U}_2$, then the Iterative Apportionment procedure constructs a *proportional* allocation.

Theorem 7: Fix an economy. If (i) $n=3$, and (ii) for each $i \in N$, $u_i \in \mathcal{U}_2$, then the Simultaneous Apportionment procedure constructs a *proportional* allocation.

The proof of Theorem 6 is in Appendix 7, and the proof of Theorem 7 is in Appendix 8. By Observation 1, we immediately have the following.

COROLLARY 2: Fix an economy. If (i) $n = 3$, and (ii) for each $i \in N$, $u_i \in \mathcal{U}_2$, then there is a *proportional* allocation that is *efficient*.

3.5 The general case

For four or more agents, the existence of *proportional* allocations is an open question; we discuss why ITERATIVE APPORTIONMENT and SIMULTANEOUS APPORTIONMENT do not generalize in our concluding discussion (Section 4). That said, we can construct approximately fair allocations using two procedures from the literature, which we refer to as ROUND-ROBIN (Caragiannis, Kurokawa, Moulin, Procaccia, Shah, and Wang, 2019) and ENVY GRAPH (Lipton, Markakis, Mossel, and Saberi, 2004; Plaut and Roughgarden, 2020).

In the general ROUND-ROBIN procedure, each agent begins with an empty basket, the agents form a queue according to index, and at each step the agent at the front of the queue adds his most-preferred unassigned time slot to his basket and then proceeds to the back of the queue. If each agent has monotonic preferences, then at each step t the agent at the front of the queue takes time slot t; in this case each agent $i \in N$ receives the schedule $\{nx+i | x \in \{0, 1, ...\}\}$. Remarkably, this simple procedure is both approximately fair and eternally approximately fair if agents are sufficiently patient:

THEOREM 8: Fix an economy and let $\varepsilon \in (0, 1]$. If for each $i \in N$, $u_i \in \mathcal{U}_{\mathrm{M}}$, then the ROUND-ROBIN procedure constructs an allocation that is *EF1* and *eternally EF1*. If, moreover, for each $i \in N$ we have $u_i \in \mathcal{U}_{\frac{1-\varepsilon}{\varepsilon}}$, then this allocation is moreover *ε-perfectly equitable* and *eternally ε-perfectly equitable*.

The proof is in Appendix 9. Observe that by relaxing exact fairness to allow approximate fairness, we have escaped our earlier impossibility results for mechanisms:

THEOREM 9: Fix an environment. For the domain $\mathcal{U}_{\mathrm{M}}$, the ROUND-ROBIN mechanism is both *strategy-proof* and *myopic*.

We omit the straightforward proof. Unfortunately, however, ROUND-ROBIN cannot guarantee *EFX*: *no-envy* and *EFX* are equivalent for allocations where all schedules are infinite, and ROUND-ROBIN generally constructs such an allocation where the first agent is envied by the others.

We conclude by analyzing the ENVY GRAPH procedure for agents with monotonic preferences. In this procedure, each agent begins with an empty basket, and at each step, the *envy graph* is the directed graph with nodes in N such that there is a directed edge from i to j if and only if i envies j. If the envy graph has any cycles, then a cycle is selected and the agents in this cycle exchange baskets, such that each agent receives the basket of the next agent in the cycle and thus becomes better off. This is repeated until the envy graph has no cycles, which necessarily happens after a finite number of exchanges (Lipton, Markakis, Mossel, and Saberi, 2004). At this point, the current time slot is assigned to an agent who is not envied, and as preferences are monotonic we have that the resulting partial allocation is *EFX* (Plaut and Roughgarden, 2020). Each basket's limit schedule is the union of its schedules across all time periods, and there is some matching of baskets to agents that occurs immediately after the time slot is assigned for an infinite collection of time slots; the baskets' limit schedules are assigned to the agents using any such matching. The resulting allocation is *EFX*, regardless of the patience of the agents:

THEOREM 10: Fix an economy. If for each $i \in N$, $u_i \in \mathcal{U}_{\mathrm{M}}$, then the ENVY GRAPH procedure constructs an allocation that is *EFX* (and thus *EF1*).

The proof is in Appendix 10.

4 Discussion

We conclude by discussing several possible directions for future research. We begin with the open question we consider most important:

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

This appears to be a difficult problem for several reasons. First, the promising ITERATIVE GREEDY ALGORITHM does not work. Second, no procedure that works is associated with a myopic mechanism. Third, ITERATIVE APPORTIONMENT works for $n = 3$ but not $n > 3$ because (i) a time slot is only skipped in a given round if multiple agents place a flag, and (ii) there are $n - 1$ rounds; we use $n = 3$ to conclude that whenever a time slot is skipped, at least one of the active agents placed a flag each time it was skipped thus far. Fourth, SIMULTANEOUS APPORTIONMENT works for $n = 3$ but not $n > 3$ because for each time slot, each agent can place at most $n - 1$ flags; we use $n = 3$ to conclude that if each of the n baskets has multiple flags, then it is possible to assign each agent a basket in which he placed a flag. Altogether, then, it seems that this solving this problem will require some new ideas.

If the general existence of *proportional* allocations were established, then there would be several interesting next steps. First, one might investigate the general existence of *envy-free* allocations; this is how the literature on cake division developed. Second, one might investigate the construction of *proportional* and *efficient* allocations—indeed, Observation 1 states that if there are *proportional* allocations, then there are *proportional* allocations that are moreover *efficient*, but there is no reason to believe that any of our procedures produces *efficient* allocations (unless preferences are identical). These are two interesting directions, but there are of course many other intriguing possibilities.

Finally, one might investigate alternative notions of approximate fairness. In our model, it follows from Theorem 1 that assigning each agent his *maximin share* (Hill, 1987; Budish, 2011) is equivalent to *proportionality*, but prominent notions of approximate fairness that we have not considered here include *near jealousy-freeness* (Gourvès, Monnot, and Tlilane, 2014), *proportionality up to one good* (Conitzer, Freeman, and Shah, 2017), *equitability up to the highest utility* (Suksompong, 2019), and *competitive equilibria in nearby Fisher markets* (Barman and Krishnamurthy, 2019). In each case, one might either investigate the approximate fairness notion directly or investigate the associated eternal approximate fairness notion.

Appendix 1

In this appendix, we prove Proposition 1.

PROPOSITION 1 (RESTATED): 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
- $S \backslash S^*$ is $(k-1)$-divisible for i.

PROOF: Assume the hypotheses. The conclusion is trivial if $v \in \{0, u_i(S)\}$, so assume $u_i(S) > v > 0$. By construction, $u_i(S^*) \leq v$.

We first prove $u_i(S^*) = v$. Since $u_i(S) > v \geq u_i(S^*)$, thus $S \backslash S^*$ is nonempty. Then necessarily $S \backslash S^*$ is infinite; else for $t = \max S \backslash S^*$, since S is 1-divisible for i and since t was skipped when constructing S^*, thus

$$\begin{aligned} u_i(S^*) &= u_i(\{t' \in S^* | t' < t\}) + u_i(\{t' \in S | t' > t\}) \\ &\geq u_i(\{t' \in S^* | t' < t\}) + u_i(\{t\}) \\ &> v, \end{aligned}$$

contradicting $u_i(S^*) \leq v$. For each $t \in S \backslash S^*$,

$$\begin{aligned} u_i(S^*) + u_i(\{t\}) &\geq u_i(\{t' \in S^* | t' < t\}) + u_i(\{t\}) \\ &> v, \end{aligned}$$

so $u_i(\{t\}) > v - u_i(S^*)$. Since $\lim_{t \in S \backslash S^*} u_i(\{t\}) = 0$,[14] thus $u_i(S^*) \geq v$, so $u_i(S^*) = v$, as desired.

To conclude, for each $t \in S \backslash S^*$, by construction $u_i(\{t\}) \geq u_i(\{t' \in S^* | t' > t\})$; thus

$$\begin{aligned} u_i(t' \in S \backslash S^* | t' > t\}) &= u_i(\{t' \in S | t' > t\}) - u_i(\{t' \in S^* | t' > t\}) \\ &\geq k u_i(\{t\}) - u_i(\{t\}) \\ &= (k-1) u_i(\{t\}), \end{aligned}$$

so $S \backslash S^*$ is $(k-1)$-divisible for i, as desired. ■

Appendix 2

In this appendix, we prove Theorem 1.

THEOREM 1 (RESTATED): Fix an economy. If there is $u_0 \in \mathcal{U}_{n-1}$ such that for each $i \in N$, $u_i = u_0$, then for any priority profile, the CONSTRAINED PRIORITIES procedure constructs a *perfectly equitable* and *efficient* allocation.

PROOF: Assume the hypotheses. We in fact prove a more general result. Let $v \in [0,1]^N$ such that $\sum v_i = 1$; we use CONSTRAINED PRIORITIES to construct an allocation π such

[14]Because u_i is a countably additive probability measure, $1 = \sum_{t \in T} u_i(\{t\}) = \lim_{t' \in T} \sum_{t=1}^{t'} u_i(\{t\})$, which directly implies $\lim_{t \in T} u_i(\{t\}) = 0$. We use this observation freely throughout our proofs.

that for each $i \in N$, $u_0(\pi_i) = v_i$. The theorem follows for the case that for each $i \in N$, $v_i = \frac{1}{n}$.

To begin, there are n empty baskets, and each is immediately assigned to an agent. The time slots are assigned to baskets in sequence, and when assigning time slot t, a flag is placed in the basket of agent i if and only if its value (according to the common utility function) would exceed the target v_i should it receive t. Time slot t is assigned to the basket of the agent whose priority for t is highest, under the constraint that it should not be assigned to a basket with a flag.

Assume, by way of contradiction, that there is a time slot for which each basket has a flag. Let t be the earliest such time slot, and for each $i \in N$, let π_i^{t-1} be the schedule in the basket of i before t is assigned. Then

$$\begin{aligned} u_0(\{1, 2, ..., t-1\}) + nu_0(\{t\}) &= \sum_{i \in N} u_0(\pi_i^{t-1} \cup \{t\}) \\ &> \sum_{i \in N} v_i \\ &= 1, \end{aligned}$$

so $nu_0(\{t\}) > u_0(\{t\}) + u_0(\{t+1, t+2, ...\})$, contradicting that $u_0 \in \mathcal{U}_{n-1}$.

Thus the procedure constructs an allocation π, so $\sum_{i \in N} u_0(\pi_i) = 1 = \sum_{i \in N} v_i$. By construction, for each $i \in N$, $u_0(\pi_i) \leq v_i$; thus for each $i \in N$, $u_0(\pi_i) = v_i$, as desired. ■

Appendix 3

In this appendix, we prove Theorem 2. Before proceeding, we remark that formally, the FUDENBERG-MASKIN procedure involves some arbitrary tie-breaking, and we prove that the resulting allocation satisfies the desired properties regardless of how this is done.

THEOREM 2 (RESTATED): Fix an economy and let $\varepsilon \in (0, 1]$. If there is $u_0 \in \mathcal{U}_{n-1} \cap \mathcal{U}_{\mathrm{M}}$ such that for each $i \in N$, $u_i = u_0$, then the FUDENBERG-MASKIN procedure constructs a *perfectly equitable* and *efficient* allocation that is *eternally EFX*. If, moreover, $u_0 \in \mathcal{U}_{\frac{1-\varepsilon}{\varepsilon}}$, then this allocation is moreover *eternally ε-perfectly equitable.*

PROOF: Assume the hypotheses and let π be the FUDENBERG-MASKIN allocation. Since the FUDENBERG-MASKIN procedure is a CONSTRAINED PRIORITIES procedure, thus by Theorem 1 we have that π is *perfectly equitable* and *efficient*.

To prove that π is *eternally EFX*, we proceed by induction. For the base step, observe that $\pi\restriction_0$ is *EFX*. For the inductive step, let $t \in T$ be such that $\pi\restriction_{t-1}$ is *EFX*. Since (i) $\pi\restriction_{t-1}$ is *EFX*, (ii) no agent envies i at $\pi\restriction_{t-1}$, and (iii) $u_0 \in \mathcal{U}_{\mathrm{M}}$, thus $\pi\restriction_t$ is *EFX*: any agent who envies i at $\pi\restriction_t$ would not do so if t or any earlier time slot were removed from $\pi_i\restriction_t$. By induction, we are done.

To conclude, let $\varepsilon \in (0, 1]$, assume that $u_0 \in \mathcal{U}_{\frac{1-\varepsilon}{\varepsilon}}$, let $t \in T$, and let $i, j, j' \in N$. Since $u_i \in \mathcal{U}_{\frac{1-\varepsilon}{\varepsilon}}$, thus

$$\begin{aligned} 1 - u_i(\{1\}) &= u_i(\{2, 3, ...\}) \\ &\geq \frac{1-\varepsilon}{\varepsilon} u_i(\{1\}), \end{aligned}$$

so $1 \geq \frac{u_i(\{1\})}{\varepsilon}$, so $\varepsilon \geq u_i(\{1\})$. If $\pi_j\restriction_t = \emptyset$, then $n \geq 2$, so $u_0 \in \mathcal{U}_1 \cap \mathcal{U}_\mathrm{M}$ and thus u_0 assigns positive utility to each date; thus by definition of the FUDENBERG-MASKIN procedure, we have that $\pi_{j'}\restriction_t$ includes at most one time slot, so

$$\begin{aligned} u_i(\pi_j\restriction_t) &\geq \varepsilon - \varepsilon \\ &\geq u_i(\{1\}) - \varepsilon \\ &\geq u_i(\pi_{j'}\restriction_t) - \varepsilon. \end{aligned}$$

If $\pi_j\restriction_t \neq \emptyset$, then there is $t' \in \pi_j\restriction_t$, so since $\pi\restriction_t$ is *EFX* we have

$$\begin{aligned} u_i(\pi_j\restriction_t) &= u_j(\pi_j\restriction_t) \\ &\geq u_j(\pi_{j'}\restriction_t \backslash \{t'\}) \\ &= u_i(\pi_{j'}\restriction_t) - u_i(\{t'\}) \\ &\geq u_i(\pi_{j'}\restriction_t) - u_i(\{1\}) \\ &\geq u_i(\pi_{j'}\restriction_t) - \varepsilon. \end{aligned}$$

Thus in both cases, $u_i(\pi_j\restriction_t) \geq u_i(\pi_{j'}\restriction_t) - \varepsilon$. Since $i, j, j' \in N$ were arbitrary, thus each agent measures any two partial schedules at time t to be within ε of one another, so $\pi\restriction_t$ is *ε-perfectly equitable*. Since $t \in T$ was arbitrary, thus π is *eternally ε-perfectly equitable*, as desired. ■

Appendix 4

In this appendix, we provide a detailed explanation of Example 2.

EXAMPLE 2: *The* ITERATIVE GREEDY ALGORITHM *does not guarantee proportionality.* Let $N = \{1, 2\}$, and consider the schedule

$$S^* = \{1\} \cup \{5x + 1, 5x + 2, 5x + 4\}_{x \in \{1,2,\ldots\}}.$$

Thus S^* is constructed by dividing T into groups of five time slots; taking only the first time slot from the first group; then taking the first, second, and fourth time slots from every group thereafter. An agent with discount factor δ measures S^* to be worth $\frac{1}{2}$ if and only if

$$(1 - \delta) + \left(\frac{\delta + \delta^2 + \delta^4}{\delta + \delta^2 + \delta^3 + \delta^4 + \delta^5} \right) \delta^5 = \frac{1}{2}.$$

This equation has two solutions in $(0, 1)$, approximately 0.536 and approximately 0.923.[15] Let δ^* be the smaller of these, and let u_1 be geometric discounting with δ^*.

We claim that $S^* = \mathcal{G}_1(\frac{1}{2}|T)$. By construction, $u_1(S^*) = \frac{1}{2}$; thus it suffices to show that for each $t \in T \backslash S^*$,

$$\begin{aligned} u_1(\{t' \in S^* | t' < t\}) + u_1(\{t\}) > \frac{1}{2} \\ = u_1(\{t' \in S^* | t' < t\}) + u_1(\{t' \in S^* | t' > t\}), \end{aligned}$$

[15] To simplify verification, for $\delta > 0$, this equation can be re-written $2\delta^8 + 2\delta^6 - \delta^4 - \delta^3 - \delta^2 - \delta + 1 = 0$. We remark that by the Rule of Signs (Descartes, 1637; see for example Wang, 2004 and Komornik, 2006 for short proofs), since the sequence of coefficients $(2, 2, -1, -1, -1, -1, 1)$ changes sign twice, this equation has at most two positive real-valued solutions, and thus the two solutions in $(0, 1)$ are in fact the only positive solutions.

or equivalently, $u_1(\{t\}) > u_1(\{t' \in S^*|t' > t\})$. If this inequality holds for $t = 5$, then it holds for $t \in \{2, 3, 4\}$; thus we need only prove it holds for $t \in \{5x\}_{x\in\{1,2,...\}}$ and $t \in \{5x + 3\}_{x\in\{1,2,...\}}$. A discount factor δ satisfies (i) the desired inequality for all $t \in \{5x\}_{x\in\{1,2,...\}}$, and (ii) the desired inequality for all $t \in \{5x + 3\}_{x\in\{1,2,...\}}$, if and only if, respectively,

$$\text{(i')}\ 1 - \delta > \left(\frac{\delta + \delta^2 + \delta^4}{\delta + \delta^2 + \delta^3 + \delta^4 + \delta^5}\right)\delta, \text{ and}$$
$$\text{(ii')}\ 1 - \delta > \left(\frac{\delta + \delta^3 + \delta^4}{\delta + \delta^2 + \delta^3 + \delta^4 + \delta^5}\right)\delta.$$

It is easy to see that if $\delta \in (0, 1)$, then (i') implies (ii'). Moreover, if $\delta \in (0, 1)$, then (i') is satisfied if and only if δ is less than approximately 0.552.[16] Since $\delta^* \approx 0.536$, thus $S^* = \mathcal{G}_1(\frac{1}{2}|T)$, as claimed.

It is easy to see that a sufficiently patient agent values S^* to be approximately $0.6 > \frac{1}{2}$. To be concrete, if agent 2 has discount factor 0.95, then $u_2(S^*) \approx 0.530 > \frac{1}{2}$.

Altogether, then, with discount factors $\delta_1 = \delta^*$ and $\delta_2 = 0.95$, if 1 uses the Greedy Algorithm to construct his own schedule, then the resulting allocation is not *proportional*. If instead 2 uses the Greedy Algorithm to construct his own schedule, then he will take the first two time slots, which 1 values more than $\frac{1}{2}$, so again the resulting allocation is not *proportional*.

Appendix 5

In this appendix, we prove Theorem 4.

Theorem 4 (Restated): Fix an environment. If $n = 2$, then for the domain $\mathcal{U}_1 \cap \mathcal{U}_{\text{G}}$, there is no *strategy-proof* mechanism that always selects *proportional* allocations.

Proof: Assume, by way of contradiction, that $\mathcal{M}$ is a *strategy-proof* mechanism that always selects *proportional* allocations. To ease notation, let the agents directly report their discount factors, so that they together report $(\delta_1, \delta_2) \in [0.5, 1)^N$ with associated utility functions $(u_{\delta_1}, u_{\delta_2})$.

First, if $(\delta_1, \delta_2) = (0.5, 0.5)$, then by *proportionality*, one agent receives $\{1\}$ and the other receives $\{2, 3, ...\}$. Assume, without loss of generality, that $\mathcal{M}(0.5, 0.5)$ assigns $\{1\}$ to 1 and $\{2, 3, ...\}$ to 2.

Second, define $\pi \equiv \mathcal{M}(0.5, 0.52)$. We cannot have $\pi_2 \supseteq \{1\}$; else $\pi_1 \subseteq \{2, 3, ...\}$ and by *proportionality* we have $u_{0.5}(\pi_1) \geq \frac{1}{2}$, so $\pi_1 = \{2, 3, ...\}$, so $u_{0.52}(\pi_2) = u_{0.52}(\{1\}) < \frac{1}{2}$, contradicting *proportionality*. Thus $\pi_1 \supseteq \{1\}$, so $\pi_2 \subseteq \{2, 3, ...\}$. We cannot have $\pi_2 \subsetneq \{2, 3, ...\}$; else

$$\begin{aligned} u_{0.52}(\mathcal{M}_2(0.5, 0.5)) &= u_{0.52}(\{2, 3, ...\}) \\ &> u_{0.52}(\pi_2) \\ &= u_{0.52}(\mathcal{M}_2(0.5, 0.52)), \end{aligned}$$

contradicting *strategy-proofness*. Altogether, then, $\pi_2 = \{2, 3, ...\}$ and $\pi_1 = \{1\}$.

[16] To simplify verification, for $\delta > 0$, (i') can be re-written $\delta^5 + \delta^4 + \delta^2 + \delta < 1$ and (ii') can be re-written $\delta^5 + \delta^4 + \delta^3 + \delta < 1$.

Finally, define $\pi' \equiv \mathcal{M}(0.51, 0.52)$. We cannot have $\pi_1' \supseteq \{1\}$; else by *proportionality* we have $\pi_1' \supsetneq \{1\}$, so by the previous paragraph we have

$$\begin{aligned} u_{0.5}(\mathcal{M}_1(0.51, 0.52)) &= u_{0.5}(\pi_1') \\ &> u_{0.5}(\{1\}) \\ &= u_{0.5}(\mathcal{M}_1(0.5, 0.52)), \end{aligned}$$

contradicting *strategy-proofness*. Then $\pi_2' \supseteq \{1\}$ and $\pi_1' \subseteq \{2, 3, ...\}$. We cannot have $\pi_2' \cap \{2, 3, 4, 5, 6\} \neq \emptyset$; else

$$\begin{aligned} u_{0.51}(\pi_2') &\geq u_{0.51}(\{1, 6\}) \\ &\approx 0.5069, \end{aligned}$$

so $u_{0.51}(\pi_1') < \frac{1}{2}$, contradicting *proportionality*. But then $\pi_1' \supseteq \{2, 3, 4, 5, 6\}$, so

$$\begin{aligned} u_{0.52}(\pi_1') &\geq u_{0.52}(\{2, 3, 4, 5, 6\}) \\ &\approx 0.5002, \end{aligned}$$

so $u_{0.52}(\pi_2') < \frac{1}{2}$, contradicting *proportionality*. ■

Appendix 6

In this appendix, we prove Theorem 5.

THEOREM 5 (RESTATED): Fix an environment. If $n = 2$, then for the domain $\mathcal{U}_1 \cap \mathcal{U}_{\mathrm{M}}$, there is no *myopic* mechanism that always selects *proportional* allocations.

PROOF: Assume, by way of contradiction, that $\mathcal{M}$ is a *myopic* mechanism that always selects *proportional* allocations. Because $\mathcal{M}$ is *myopic*, it always assigns date 1 to the same agent whenever both agents assign utility 0.39 to this date; assume, without loss of generality, that in this case $\mathcal{M}$ assigns date 1 to agent 1.

Let $u_1 \in \mathcal{U}_1 \cap \mathcal{U}_{\mathrm{M}}$ be such that $(0.39, 0.3, 0.11, 0.1)$ gives the utilities of the first four dates and thereafter u_1 iteratively assigns one-half of the remaining utility to the next date. Similarly, let $u_2 \in \mathcal{U}_1 \cap \mathcal{U}_{\mathrm{M}}$ be such that $(0.39, 0.13, 0.13, 0.13)$ gives the utilities of the first four dates and thereafter u_2 iteratively assigns one-half of the remaining utility to the next date. Finally, define $\pi \equiv \mathcal{M}(u_1, u_2)$. By assumption, $\pi_1 \supseteq \{1\}$. We cannot have $\pi_1 \cap \{2, 3, 4\} \neq \emptyset$; else

$$\begin{aligned} u_2(\pi_1) &\geq u_2(\{1, 4\}) \\ &= 0.52, \end{aligned}$$

so $u_2(\pi_2) < \frac{1}{2}$, contradicting *proportionality*. But then $\pi_1 \cap \{2, 3, 4\} = \emptyset$, so

$$\begin{aligned} u_1(\pi_1) &\leq u_1(\{1\} \cup \{5, 6, ...\}) \\ &= 0.49, \end{aligned}$$

so $u_1(\pi_1) < \frac{1}{2}$, contradicting *proportionality*. ■

Appendix 7

In this appendix, we prove Theorem 6. Before proceeding, we remark that formally, the Iterative Apportionment procedure involves some arbitrary tie-breaking, and we prove that the resulting allocation is *proportional* regardless of how this is done.

Theorem 6 (Restated): Fix an economy. If (i) $n = 3$, and (ii) for each $i \in N$, $u_i \in \mathcal{U}_2$, then the Iterative Apportionment procedure constructs a *proportional* allocation.

Proof: Assume the hypotheses. In the Iterative Apportionment procedure, the three agents fill a first basket and assign its contents S_1 to an agent i_1, then the remaining two agents fill a second basket and assign its contents S_2 to an agent i_2, and finally the last agent receives the rest of the time slots.

Step 1: Construct $S_1 \in \mathcal{S}$ and choose $i_1 \in N$ such that (i) $u_{i_1}(S_1) \geq \frac{1}{3}$, and (ii) for each $j \in N \backslash \{i_1\}$, $u_j(S_1) \leq \frac{1}{3}$.

The agents fill the first basket by considering the time slots 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 $\frac{1}{3}$. If there are no flags, then t is added to the basket and the agents move to the next time slot. If there is one flag, then let i_1 be the agent who placed the flag, let S_1 denote the contents of the basket with t, and assign S_1 to i_1; in this case we are clearly done with Step 1. If there are two or three flags, then t is skipped and the agents move to the next time slot.

If there is no time slot with one flag, then let S_1 denote the first basket's limit schedule, or the union of its schedules across all time periods. In this case, we claim there is $i_1 \in N$ such that $u_{i_1}(S_1) = \frac{1}{3}$. Indeed, by construction, for each $j \in N$, $u_j(S_1) \leq \frac{1}{3} < 1 = u_j(T)$, so there is at least one time slot for which an agent places a flag. Assume, by way of contradiction, that there is a maximum t with a flag, and let i be an agent who places a flag for t. Since $u_i \in \mathcal{U}_1$, thus

$$\begin{aligned} u_i(S_1) &= u_i(\{t' \in S_1 | t' < t\}) + u_i(\{t+1, t+2, ...\}) \\ &\geq u_i(\{t' \in S_1 | t' < t\}) + u_i(\{t\}) \\ &> \frac{1}{3}, \end{aligned}$$

contradicting $u_i(S_1) \leq \frac{1}{3}$. Thus there is an infinite collection of flags, so there is at least one agent who places an infinite collection of flags; let i_1 be any such agent, and let T_1 denote the infinite collection of time slots for which i_1 places a flag. For each $t \in T_1$,

$$\begin{aligned} u_{i_1}(S_1) + u_{i_1}(\{t\}) &\geq u_{i_1}(\{t' \in S_1 | t' < t\}) + u_{i_1}(\{t\}) \\ &> \frac{1}{3}, \end{aligned}$$

so $u_{i_1}(\{t\}) > \frac{1}{3} - u_{i_1}(S_1)$. Since $\lim_{t \in T_1} u_{i_1}(\{t\}) = 0$, thus $u_{i_1}(S_1) \geq \frac{1}{3}$, so $u_{i_1}(S_1) = \frac{1}{3}$, as desired.

Step 2: Construct $S_2 \in \mathcal{S}$ and choose $i_2 \in N$ such that (i) $u_{i_2}(S_2) \geq \frac{1}{3}$, and (ii) for the unique $j \in N \backslash \{i_1, i_2\}$, $u_j(S_2) \leq \frac{1}{3}$.

Define $N' \equiv N\backslash\{i_1\}$. The two agents in N' fill a second basket to construct S_2 using the same procedure that was used to construct S_1; we provide the details to avoid ambiguity. The two remaining agents fill the second basket by considering the time slots in $T\backslash S_1$ 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 $\frac{1}{3}$. If there are no flags, then t is added to the basket and the agents move to the next time slot. If there is one flag, then let i_2 be the agent who placed the flag, let S_2 denote the contents of the basket with t, and assign S_2 to i_2; in this case we are clearly done with Step 2. If there are two flags, then t is skipped and the agents move to the next time slot.

If there is no time slot with one flag, then let S_2 denote the second basket's limit schedule, or the union of its schedules across all time periods. In this case, we claim there is $i_2 \in N'$ such that $u_{i_2}(S_2) = \frac{1}{3}$. Indeed, by construction, for each $j \in N'$, $u_j(S_2) \leq \frac{1}{3} < \frac{2}{3} \leq u_j(T\backslash S_1)$, so there is at least one time slot for which an agent places a flag. Assume, by way of contradiction, that there is a maximum t with a flag. Let $N_1 \subseteq N$ denote the agents who placed a flag for t during the construction of S_1, and let $N_2 \subseteq N' \subseteq N$ denote the agents who placed a flag for t during the construction of S_2. Since $t \in T\backslash S_1$, thus $|N_1| \geq 2$, and by assumption, $|N_2| = 2$; since $|N| = 3$, thus there is $i \in N_1 \cap N_2$. Since $i \in N_2$, thus $i \neq i_1$, so by Step 1 we have $u_i(S_1) \leq \frac{1}{3}$, so since $i \in N_1$ we have

$$\begin{aligned} u_i(\{t' \in S_1 | t' < t\}) + u_i(\{t\}) &> \frac{1}{3} \\ &\geq u_i(S_1) \\ &= u_i(\{t' \in S_1 | t' < t\}) + u_i(\{t' \in S_1 | t' > t\}), \end{aligned}$$

and thus $u_i(\{t\}) > u_i(\{t' \in S_1 | t' > t\})$. Moreover, we have $u_i \in \mathcal{U}_2$, so altogether

$$\begin{aligned} u_i(\{t+1, t+2, ...\}\backslash S_1) &= u_i(\{t+1, t+2, ...\}) - u_i(\{t' \in S_1 | t' > t\}) \\ &> 2u_i(\{t\}) - u_i(\{t\}) \\ &= u_i(\{t\}). \end{aligned}$$

But then

$$\begin{aligned} u_i(S_2) &= u_i(\{t' \in S_2 | t' < t\}) + u_i(\{t+1, t+2, ...\}\backslash S_1) \\ &> u_i(\{t' \in S_2 | t' < t\}) + u_i(\{t\}) \\ &> \frac{1}{3}, \end{aligned}$$

contradicting $u_i(S_2) \leq \frac{1}{3}$. Thus both agents in N' place an infinite collection of flags; let i_2 be either of these agents, and let T_2 denote the infinite collection of time slots for which i_2 places a flag. For each $t \in T_2$,

$$\begin{aligned} u_{i_2}(S_2) + u_{i_2}(\{t\}) &\geq u_{i_2}(\{t' \in S_2 | t' < t\}) + u_{i_2}(\{t\}) \\ &> \frac{1}{3}, \end{aligned}$$

so $u_{i_2}(\{t\}) > \frac{1}{3} - u_{i_2}(S_2)$. Since $\lim_{t \in T_2} u_{i_2}(\{t\}) = 0$, thus $u_{i_2}(S_2) \geq \frac{1}{3}$, so $u_{i_2}(S_2) = \frac{1}{3}$, as desired.

Step 3: Conclude.

Let i_3 denote the unique member of $N\backslash\{i_1, i_2\}$, and let π be the allocation such that $\pi_{i_1} = S_1$, $\pi_{i_2} = S_2$, and $\pi_{i_3} = T\backslash(\pi_{i_1} \cup \pi_{i_2})$. By Step 1, $u_{i_1}(\pi_{i_1}) \geq \frac{1}{3}$ and $u_{i_3}(\pi_{i_1}) \leq \frac{1}{3}$. By Step 2, $u_{i_2}(\pi_{i_2}) \geq \frac{1}{3}$ and $u_{i_3}(\pi_{i_2}) \leq \frac{1}{3}$. Thus $u_{i_3}(\pi_{i_3}) \geq \frac{1}{3}$, so π is *proportional*, as desired. ■

Appendix 8

In this appendix, we prove Theorem 7. Before proceeding, we remark that formally, the Simultaneous Apportionment procedure involves some arbitrary tie-breaking, and we prove that the resulting allocation is *proportional* regardless of how this is done.

Theorem 7 (Restated): Fix an economy. If (i) $n = 3$, and (ii) for each $i \in N$, $u_i \in \mathcal{U}_2$, then the Simultaneous Apportionment procedure constructs a *proportional* allocation.

Proof: Assume the hypotheses. In the Simultaneous Apportionment procedure, the three agents begin to fill three baskets together, and there are three cases:

- The agents exhaust the time slots and the baskets are assigned arbitrarily.
- At some point one of the agents takes a basket that nobody else values more than $\frac{1}{3}$, then the remaining agents continue with the remaining baskets.
- At some point one of the agents takes a basket that someone else values more than $\frac{1}{3}$, and the rest of the allocation is immediately determined with a procedure involving Divide and Choose.

Step 1: The three agents fill the three baskets, possibly stopping after time slot t_1 when basket S_1 is assigned to agent i_1, who values S_1 at least $\frac{1}{3}$.

To begin, there are three empty baskets. The time slots are assigned to baskets in sequence, and when assigning time slot t, agent i places a flag in basket j if and only if he believes the value of basket j would exceed $\frac{1}{3}$ should it receive t. Time slot t is then assigned to any basket with a minimal number of flags. If this basket has zero flags, then the agents continue to the next time slot. If this basket has at least one flag, then let i_1 be any agent who placed a flag in this basket, let S_1 denote the contents of this basket with t, and let t_1 denote this time slot; in this case we are clearly done with Step 1.

If every time slot is placed in a basket with zero flags, then by construction, all agents agree that each basket is worth no more than $\frac{1}{3}$, so they agree that each basket is worth $\frac{1}{3}$, so the baskets may be arbitrarily assigned to the agents to form a *proportional* allocation.

Step 2: If the three agents stopped filling the three baskets at some t_1, then complete the allocation.

By construction, the basket that was assigned to i_1 had at least one flag at t_1. We first claim that no agent placed three flags at t_1. Indeed, assume by way of contradiction that some agent i placed three flags at t_1, and let A, B, and C denote the contents of the three baskets before t_1 is assigned. Then

$$\begin{aligned}
u_i(\{1,2,...,t_1-1\}) + 3u_i(\{t_1\}) &= \sum_{S\in\{A,B,C\}} u_i(S\cup\{t_1\}) \\
&> \frac{1}{3}+\frac{1}{3}+\frac{1}{3} \\
&= 1,
\end{aligned}$$

so $3u_i(\{t_1\}) > u_i(t_1) + u_i(\{t_1+1, t_1+2, ...\})$, contradicting $u_i \in \mathcal{U}_2$.

Since (i) each agent placed at most two flags at t_1 and (ii) the basket assigned to i_1 had the minimal number of flags at t_1, thus the basket assigned to i_1 either had one flag or two flags at t_1. Moreover, if this basket had two flags at t_1, then each agent placed two flags at t_1.

Case 1: The basket assigned to i_1 had one flag at t_1.

In this case, the two agents continue as before with the remaining two baskets; we provide the details to avoid ambiguity. The remaining time slots are assigned to the remaining baskets in sequence, and when assigning time slot t, agent i places a flag in basket j if and only if he believes the value of j would exceed $\frac{1}{3}$ should it receive t. Time slot t is then assigned to any basket with a minimal number of flags. If this basket has zero flags, then the agents continue to the next time slot. If this basket has at least one flag, then let i_2 be any agent who placed a flag in this basket, let S_2 denote the contents of this basket with t, and let t_2 denote this time slot.

If every time slot is placed in a basket with zero flags, then by construction, the remaining agents agree that each remaining basket is worth no more than $\frac{1}{3}$. Moreover, since the basket assigned to i_1 had one flag at t_1, thus the remaining agents agree that the remaining baskets are together worth $\frac{2}{3}$. Altogether, then, the remaining agents agree that each remaining basket is worth $\frac{1}{3}$, so the remaining baskets may be arbitrarily assigned to the remaining agents to form a *proportional* allocation.

If t_2 is placed in a basket with at least one flag, then let i_3 be the unique member of $N\backslash\{i_1, i_2\}$ and let π be the allocation such that $\pi_{i_1} = S_1$, $\pi_{i_2} = S_2$, and $\pi_{i_3} = T\backslash(\pi_{i_1}\cup\pi_{i_2})$. If the basket assigned to i_2 has one flag at t_2, then π is clearly *proportional*. If the basket assigned to i_2 has two flags at t_2, then both agents placed two flags at t_2. Let A denote the contents of the basket without t_2 right after t_2 is assigned; since $u_{i_3} \in \mathcal{U}_1$ and since i_3 placed a flag in both baskets at t_2, thus

$$\begin{aligned}
u_{i_3}(\pi_{i_3}) &= u_{i_3}(A) + u_{i_3}(\{t_2+1, t_2+2, ...\}) \\
&\geq u_{i_3}(A) + u_{i_3}(\{t_2\}) \\
&> \frac{1}{3},
\end{aligned}$$

so π is *proportional.*

Case 2: The basket assigned to i_1 had two flags at t_1, and each agent placed two flags at t_1.

In this case, each basket received two flags at t_1. Let i_2 be the agent other than i_1 who placed a flag in the assigned basket at t_1, and let S_2 denote the contents of the remaining basket that received a flag from i_2 at t_1. Let i_3 denote the unique member of $N \backslash \{i_1, i_2\}$, and let S_3 denote the contents of the remaining basket that did not receive a flag from i_2 at t_1. Necessarily, i_3 placed a flag in the basket of S_3 at t_1.

Let i_2 divide the remaining time slots into two parts H and H' he considers to have equal value using the GREEDY ALGORITHM; formally, define

$$H \equiv \mathcal{G}_{i_2}(\frac{1}{2}u_{i_2}(\{t_1+1, t_1+2, ...\})|\{t_1+1, t_1+2, ...\}), \text{ and}$$
$$H' \equiv \{t_1+1, t_1+2, ...\} \backslash H.$$

By Proposition 1, i_2 indeed considers H and H' to have equal value. Let H_3 denote a member of $\{H, H'\}$ that i_3 considers most valuable, and let H_2 denote the other member of $\{H, H'\}$.

Let π be the allocation such that $\pi_{i_1} = S_1$, $\pi_{i_2} = S_2 \cup H_2$, and $\pi_{i_3} = S_3 \cup H_3$. For each $j \in \{i_2, i_3\}$, since $u_j \in \mathcal{U}_2$ and since j placed a flag in the basket for S_j at t_1, thus

$$\begin{aligned} u_j(\pi_j) &= u_j(S_j) + u_j(H_j) \\ &\geq u_j(S_j) + \frac{1}{2}u_j(\{t_1+1, t_1+2, ...\}) \\ &\geq u_j(S_j) + u_j(\{t_1\}) \\ &> \frac{1}{3}, \end{aligned}$$

so π is *proportional*, as desired. ■

Appendix 9

In this appendix, we prove Theorem 8.

THEOREM 8 (RESTATED): Fix an economy and let $\varepsilon \in (0,1]$. If for each $i \in N$, $u_i \in \mathcal{U}_{\mathrm{M}}$, then the ROUND-ROBIN procedure constructs an allocation that is *EF1* and *eternally EF1*. If, moreover, for each $i \in N$ we have $u_i \in \mathcal{U}_{\frac{1-\varepsilon}{\varepsilon}}$, then this allocation is moreover *ε-perfectly equitable* and *eternally ε-perfectly equitable*.

PROOF: Assume the hypotheses and let π be the ROUND-ROBIN allocation. Moreover, let $\varepsilon \in (0,1]$, let $i, j, j' \in N$, and let $t \in T$.

First, we claim that (i) $u_i(\pi_j) \geq u_i(\pi_{j'} \backslash \{j'\})$, and (ii) $u_i(\pi_j \restriction_t) \geq u_i(\pi_{j'} \restriction_t \backslash \{j'\})$. Indeed, by construction, (i) both π_j and $\pi_{j'}$ are infinite, and moreover (ii) for each $k \in \{1, 2, ...\}$, the kth-earliest member of π_j is earlier than the kth-earliest member of $\pi_{j'} \backslash \{j'\}$. Since $u_i \in \mathcal{U}_{\mathrm{M}}$, thus we have our claim, as desired.

Second, we claim that if for each $i \in N$ we have $u_i \in \mathcal{U}_{\frac{1-\varepsilon}{\varepsilon}}$, then (i) $u_i(\pi_j) \geq u_i(\pi_{j'}) - \varepsilon$, and (ii) $u_i(\pi_j \restriction_t) \geq u_i(\pi_{j'} \restriction_t) - \varepsilon$. Indeed, since $u_i \in \mathcal{U}_{\frac{1-\varepsilon}{\varepsilon}}$, thus as argued in the proof of Theorem 2 we have $\varepsilon \geq u_i(\{1\})$, so since $u_i \in \mathcal{U}_{\mathrm{M}}$ we have $\varepsilon \geq u_i(\{1\}) \geq u_i(\{j'\})$. The desired claim follows immediately from the first claim.

To conclude, since $i, j, j' \in N$ and $t \in T$ were arbitrary, thus (i) by the first claim, we have that π is *EF1* and *eternally EF1*, and (ii) by the second claim, we have that if for each $i \in N$ we have $u_i \in \mathcal{U}_{\frac{1-\varepsilon}{\varepsilon}}$, then π is *ε-perfectly equitable* and *eternally ε-perfectly equitable*. Since $\varepsilon \in (0,1]$ was arbitrary, we are done. ■

Appendix 10

In this appendix, we prove Theorem 10. Before proceeding, we remark that formally, the Envy Graph procedure involves some arbitrary tie-breaking, and we prove that the resulting allocation is *EFX* regardless of how this is done.

Theorem 10 (Restated): Fix an economy. If for each $i \in N$, $u_i \in \mathcal{U}_{\text{M}}$, then the Envy Graph procedure constructs an allocation that is *EFX* (and thus *EF1*).

Proof: Let π be the Envy Graph allocation. For each $t \in T$ and each $i \in N$, let π_i^t denote the schedule in the basket held by i immediately after t is assigned; by construction, $u_i(\pi_i^1) \leq u_i(\pi_i^2) \leq ... \leq u_i(\pi_i)$. Define $T' \equiv \{t \in T | \text{for each } i \in N, \pi_i^t \subseteq \pi_i\}$; by construction, T' is infinite.

Let $i, j \in N$. Since each agent's utility function belongs to $\mathcal{U}_{\text{M}}$, thus the agents share a common ranking of the time slots, so as argued by Plaut and Roughgarden (2020) in the proof of their Theorem 6.2, for each $t \in T$ we have $u_i(\pi_i^t) \geq u_i(\pi_j^t \backslash \{\max \pi_j^t\})$. If π_j is finite, then there is $t^* \geq \max \pi_j$ such that $\pi_j^{t^*} = \pi_j$, and thus

$$\begin{aligned} u_i(\pi_i) &\geq u_i(\pi_i^{t^*}) \\ &\geq u_i(\pi_j^{t^*} \backslash \{\max \pi_j^{t^*}\}) \\ &= u_i(\pi_j \backslash \{\max \pi_j\}), \end{aligned}$$

from which the desired conclusion directly follows. If π_j is infinite, then

$$\begin{aligned} u_i(\pi_i) &= \lim_{t \in T'} u_i(\pi_i^t) \\ &\geq \lim_{t \in T'} u_i(\pi_j^t \backslash \{\max \pi_j^t\}) \\ &= \lim_{t \in T'} u_i(\pi_j^t) - \lim_{t \in T'} u_i(\{\max \pi_j^t\}) \\ &= u_i(\pi_j) - 0, \end{aligned}$$

from which the desired conclusion directly follows. ■

References

Amanatidis, Georgios, Georgios Birmpas, George Christodoulou, and Evangelos Markakis (2017). "Truthful Allocation Mechanisms Without Payments: Characterization and Implications on Fairness." In: *EC '17: Proceedings of the 2017 ACM Conference on Economics and Computation*, 545-562.

Anbarci, Nejat, Ching-Jun Sun, and M. Utku Ünver (2021). "Designing practical and fair sequential team contests: The case of penalty shootouts." *Games and Economic Behavior* 130, 25-43.

Arrow, Kenneth and Debreu, Gerard (1954). "Existence of an equilibrium for a competitive economy." *Econometrica* 22, 265-290.

Austin, A. Keith (1982). "Sharing a cake." *The Mathematical Gazette* 66, 212-215.

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.

Aziz, Haris and Chun Ye (2014). "Cake Cutting Algorithms for Piecewise Constant and Piecewise Uniform Valuations." In: *Web and Internet Economics: 10th International Conference (WINE 2014)*, 1-14.

Babaioff, Moshe, Tomer Ezra, and Uriel Feige (2021). "Fair and Truthful Mechanisms for Dichotomous Valuations." In: *Proceedings of the Thirty-Fifth AAAI Conference on Artificial Intelligence (AAAI-21)*, 5119-5126.

Barman, Siddharth and Sanath Krishnamurthy (2019). "On the Proximity of Markets with Integral Equilibria." In: *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence (AAAI-19)*, 1748-1755.

Barman, Siddharth, Sanath Krishnamurthy, and Rohit Vaish (2018). "Finding Fair and Efficient Allocations." In: *EC '18: Proceedings of the 2018 ACM Conference on Economics and Computation*, 557-574.

Barman, Siddharth and Paritosh Verma (2022). "Truthful and Fair Mechanisms for Matroid-Rank Valuations." In: *Proceedings of the Thirty-Sixth AAAI Conference on Artificial Intelligence (AAAI-22)*, 4801-4808.

Bei, Xiaohui, Ning Chen, Guangda Huzhang, Biaoshuai Tao, and Jiajun Wu (2017). "Cake Cutting: Envy and Truth." In: *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence (IJCAI-17)*, 3625-3631.

Bernstein, Sergei (1917). "Attempt at an axiomatic foundation of probability theory [in Russian]." *Communications of the Kharkov Mathematical Society* [in Russian] 15, 209-274. Translation In: Sheynin, O. (2005). *Probability and Statistics: Russian Papers of the Soviet Period.* Berlin, Germany: NG Verlag.

Bezáková, Ivona and Varsha Dani (2005). "Allocating Indivisible Goods." *ACM SIGecom Exchanges* 5, 11-18.

Bilò, Vittorio, Ioannis Caragiannis, Michele Flammini, Ayumi Igarashi, Gianpiero Monaco, Dominik Peters, Cosimo Vinci, and William Zwicker (2022). "Almost envy-free allocations with connected bundles." *Games and Economic Behavior* 131, 197-221.

Brams, Steven and Mehmet Ismail (2018). "Making the Rules of Sports Fairer." *SIAM Review* 60, 181-202.

Brams, Steven, Mehmet Ismail, and D. Marc Kilgour (2023). "Fairer Shootouts in Soccer: The m-n Rule." Working paper.

Brams, Steven, Mehmet Ismail, D. Marc Kilgour, and Walter Stromquist (2018). "Catch-Up: A Rule That Makes Service Sports More Competitive." *The American Mathematical Monthly* 125, 771-796.

Brams, Steven and Alan Taylor (1995). "An Envy-Free Cake Division Protocol." *The American Mathematical Monthly* 102, 9-18.

Brams, Steven and Alan Taylor (2000). *The Win-Win Solution: Guaranteeing Fair Shares to Everybody.* New York, New York: W.W. Norton.

Brams, Steven, Alan Taylor, and William Zwicker (1995). "Old and new moving-knife schemes." *The Mathematical Intelligencer* 7, 30-35.

Bu, Xiaolin, Jiaxin Song, and Biaoshuai Tao (2023). "On existence of truthful fair cake cutting mechanisms." *Artificial Intelligence* 319, Article 103904.

Budish, Eric (2011). "The Combinatorial Assignment Problem: Approximate Competitive Equilibrium form Equal Incomes." *Journal of Political Economy* 119, 1061-1103.

Caragiannis, Ioannis, Christos Kaklamanis, Panagiotis Kanellopoulos, and Maria Kyropoulou (2009). "On Low-Envy Truthful Allocations." *Algorithmic Decision Theory: First International Conference (ADT 2009)*, 111-119.

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.

Chateauneuf, Alain (1985). "On the Existence of a Probability Measure Compatible with a Total Preorder on a Boolean Algebra." *Journal of Mathematical Economics* 14, 43-52.

Chen, Yiling, John Lai, David Parkes, and Ariel Procaccia (2013). "Truth, justice, and cake cutting." *Games and Economic Behavior* 77, 284-297.

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

Cho, Wonki and William Thomson (2023). "Strategy-Proofness in Private Good Economies with Linear Preferences: An Impossibility Result." Working paper, SSRN.

Chun, Youngsub (2016). *Fair Queueing.* Cham, Switzerland: Springer International Publishing.

Conitzer, Vincent, Rupert Freeman, and Nisarg Shah (2017). "Fair Public Decision Making." In: *EC '17: Proceedings of the 2017 ACM Conference on Economics and Computation*, 629-646.

de Finetti, Bruno (1937). "La prévision: ses lois logiques, ses sources subjectives [in French]." *Annales d'Institut Henri Poincaré* 7, 1-68.

Descartes, René (1637). *Discours de la Méthode pour bien conduire sa raison, et chercher la vérité dans les sciences. Plus La Dioptrique. Les Météores. Et La Géométrie* [in French]. Leiden, the Netherlands: Jan Maire. Translation of *La Géométrie*: Descartes, René. (1954). *The Geometry of René Descartes with a Facsimile of the First Edition.* Translators: Smith, David and Marcia Latham. New York, New York: Dover Publications, Inc.

Dolan, Robert (1978). "Incentive Mechanisms for Priority Queuing Problems." *The Bell Journal of Economics* 9, 421-436.

Dubins, Lester and Edwin Spanier (1961). "Hot to Cut A Cake Fairly." *The American Mathematical Monthly* 68, 1-17.

Eisenberg, Edmund and David Gale (1959). "Consensus of subjective probabilities: the pari-mutuel method." *Annals of Mathematical Statistics* 30, 165-168.

Erdős, Pál, István Joó, and Vilmos Komornik (1990). "Characterization of the unique expansions $1 = \sum_{i=1}^{\infty} q^{-n_i}$ and related problems." *Bulletin de la Société Mathématique de France* 118, 377-390.

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.

Gale, David and Shapley, Lloyd (1962). "College admissions and the stability of marriage." *American Mathematical Monthly* 69, 9-15.

Gourvès, Laurent, Jérôme Monnot, and Lydia Tlilane (2014). "Near Fairness in Matroids." In: *Proceedings of the Twenty-first European Conference on Artificial Intelligence (ECAI'14)*, 393-398.

Halpern, Daniel, Ariel Procaccia, Alexandros Psomas, and Nisarg Shah (2020). "Fair Division with Binary Valuations: One Rule to Rule Them All." In: *Web and Internet Economics: 16th International Conference (WINE 2020)*, 370-383.

Hill, Theodore (1987). "Partitioning General Probability Measures." *The Annals of Probability* 15, 804-813.

Jackson, Matthew and Leeat Yariv (2015). "Collective Dynamic Choice: The Necessity of Time Inconsistency." *American Economic Journal: Microeconomics* 7, 150-178.

Jones, Rafe (2011). "Achievement Sets of Sequences." *The American Mathematical Monthly* 118, 508-521.

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.

Kettering, Jeremy and Kochov, Asen (2021). "Consumption Smoothing and Discounting in Infinite-Horizon, Discrete-Choice Problems." *Mathematics of Operations Research* 47, 1957-1969.

Klaus, Bettina and Eiichi Miyagawa (2001). "Strategy-proofness, solidarity, and consistency for multiple assignment problems." *International Journal of Game Theory* 30, 421-435.

Kochov, Asen and Yangwei Song (2023). "Intertemporal Hedging and Trade in Repeated Games with Recursive Utility." Working paper, SSRN.

Kolm, Serge-Christophe (1971). *Justice et Equité* [in French]. Paris, France: Centre d'Etudes Prospectives et d'Economie Mathématique Appliquées à la Planification. Translation: Kolm, Serge-Christophe (2002). *Justice and Equity.* Translator: Harold F. See. Cambridge, Massachusetts: The MIT Press.

Komornik, Vilmos (2006). "Another Short Proof of Descartes's Rule of Signs." *American Mathematical Monthly* 113, 829-830.

Koopman, Bernard (1940). "The axioms and algebra of intuitive probability." *Annals of Mathematics* 41, 269-292.

Koopmans, Tjalling (1960). "Stationary Ordinal Utility and Impatience." *Econometrica* 28, 287-309.

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

Leo, Greg (2017). "Taking turns." *Games and Economic Behavior* 102, 525-547.

Levmore, Saul and Elizabeth Cook (1981). *Super Strategies for Puzzles and Games.* Garden City, New York: Doubleday and Company.

Liapounoff, Aleksandr (1940). "Sur les fonctions-vecteurs complètement additives [in Russian]." *Izvestiya Akademii Nauk SSSR, Seriya Matematicheskaya* 4, 465-478.

Lipton, Richard, Evangelos Markakis, Elchanan Mossel, and Amin Saberi (2004). "On Approximately Fair Allocations of Indivisible Goods." In: *EC '04: Proceedings of the 5th ACM Conference on Electronic Commerce*, 125-131.

Lowry, S. Todd (1987). *The Archaeology of Economic Ideas.* Durham, North Carolina: Duke University Press.

Mackenzie, Andrew (2019). "A foundation for probabilistic beliefs with or without atoms." *Theoretical Economics* 14, 709-778.

McKenzie, Lionel (1954). "On equilibrium in Graham's model of world trade and other competitive systems." *Econometrica* 22, 147-161.

Mossel, Elchanan and Omer Tamuz (2010). "Truthful Fair Division." In: *Algorithmic Decision Theory: Third International Symposium (SAGT 2010)*, 288-299.

Nash, John (1950). "The bargaining problem." *Econometrica* 18, 155-162.

Pazner, Elisha and David Schmeidler (1978). "Egalitarian equivalent allocations: A new concept of economic equity." *The Quarterly Journal of Economics* 92, 671-687.

Plaut, Benjamin and Tim Roughgarden (2020). "Almost envy-freeness with general valuations." *SIAM Journal on Discrete Mathematics* 34, 1039-1068.

Pólya, George and Gabor Szegö (1925). *Aufgaben und Lehrsätze aus der Analysis* [in German]. Berlin, Germany: Springer-Verlag. Translation of first volume: Pólya, George and Gabor Szegö (1978). *Problems and Theorems in Analysis I.* Translator: Dorothee Aeppli. Berlin, Germany: Springer-Verlag.

Rachmilevitz, Shiran (2013). "Endogenous bid rotation in repeated auctions." *Journal of Economic Theory* 148, 1714-1725.

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

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.

Savage, Leonard (1954). *The Foundations of Statistics.* Wiley, NY: Dover Publications.

Schummer, James (1997). "Strategy-proofness versus efficiency on restricted domains of exchange economies." *Social Choice and Welfare* 14, 47-56.

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.

Steinhaus, Hugo (1949). "Sur la division pragmatique." *Econometrica* 17 (Supplement), 315-319.

Stromquist, Walter (1980). "How to cut a cake fairly." *The American Mathematical Monthly* 87, 640-644. Addendum (1981), volume 88, 613-614.

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

Suksumpong, Warut (2019). "Fairly allocating contiguous blocks of indivisible items." *Discrete Applied Mathematics* 260, 227-236.

Thomson, William (2007). "Children crying at birthday parties. Why?" *Economic Theory* 31, 501-521.

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

Varian, Hal (1974). "Equity, envy and efficiency." *Journal of Economic Theory* 29, 217-244.

Villegas, Cesareo (1964). "On Quantitative Probability σ-Algebras." *Annals of Mathematical Statistics* 35, 1787-1796.

Wang, Xiaoshen (2004). "A Simple Proof of Descartes's Rule of Signs." *American Mathematical Monthly* 111, 525-526.

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