Efficient Egalitarian Equivalent Allocations over a Single Good*

Marco LiCalzi
Dept. of Applied Mathematics
University of Venice

Antonio Nicolò
Dept. of Economics
University of Padua

(October 2005)

Abstract. This paper studies efficient and egalitarian allocations over a single heterogeneous and infinitely divisible good. We prove the existence of such allocations using only measure-theoretic arguments. Under the additional assumption of complete information, we identify a sufficient condition on agents' preferences that makes it possible to apply the Pazner-Schmeidler rule for uniquely selecting an efficient egalitarian equivalent allocation. Finally, we exhibit a simple procedure that implements the Pazner-Schmeidler selection in a subgame-perfect equilibrium.

Keywords: egalitarism, equity, fair division, implementation, no-envy.

JEL Classification Numbers: D630, C700, D300.
MathSci Subject Classification: 91B32, 91B54, 28B05.

Correspondence to:

Antonio Nicolò	Department of Economics, University of Padua
	Via del Santo 33
	37123 Padua, Italy
Phone:	[++39] 049-827-4285
Fax:	[++39] 049-827-4211
E-mail:	antonio.nicolo@unipd.it

* We thank the seminar audience at Boston University for their comments. The research on this paper greatly benefited from the warm hospitality of Boston University. Financial support from MIUR is acknowledged.

1 Introduction

A fair division problem arises when two or more agents are called to divide a good over which they claim equal rights. The oldest known examples include Abraham and Lot arguing over land division (Genesis 13), and Prometheus and Zeus disputing a pile of meat (Hesiod's Theogony). A recent dramatic example is the carving of Bosnia and Herzegovina as an independent entity within the Dayton Accords that put an end to a 3-year civil war over the spoils of the former Republic of Yugoslavia.

There are several situations where the solution of a fair division problem cannot call on instruments like prices, monetary compensations, or auctions. This may be due to liquidity constraints; or to the psychological difficulty to bring a dispute down to monetary evaluations; or to political constraints, as in the case of Bosnia and Herzegovina; or to the presence of judicially enforceable rights — such as under U.S. law — "to seek partition in kind, or physical division, of jointly owned land"; see Miceli and Sirmans (2000).

This paper studies the problem of fair division when the dispute must be resolved using division in kind. We are interested in devising a procedure that can help the parties to reach an outcome that is both fair and efficient. We assume that the disputed object is a single infinitely divisible good over which agents have heterogeneous preferences and that there are no consumption externalities. The canonical example is the division of a cake, when agents have different (additive) preferences over different slices; see Steinhaus (1948). A less obvious example is the case of a finite (or countable) number of homogeneous infinitely divisible goods, where the aggregate endowment is viewed as the single heterogeneous good at stake; see Chambers (2005, Section 5).

There are two main ordinal concepts in the fair division literature. The first is the envy-free principle which states that each party should (weakly) prefers its share to anyone else's. This was proposed by Gamow and Stern (1958, pp. 117–119), but became widely known after Foley (1967). Any efficient envy-free allocation is ex post stable because no one desires to exchange what he received with anyone else's share. However, this solution concept suffers from a multiplicity problem that makes it less satisfactory from an ex-ante, or procedural, point of view. There are in general many efficient envy-free allocations, and each of them provides different payoffs to the agents. Therefore, they are likely to disagree on how to select one among these allocations. The divide-and-choose mechanism under complete information, for instance, selects among all the efficient envy-free allocations the division that maximizes the payoff to the divider — so conflict is likely to shift over how the divider is chosen.

An alternative normative concept is the egalitarian equivalent criterion which states that each party should be indifferent between getting his share and some reference bundle, identical for all agents. This was introduced by

Pazner and Schmeidler (1978) to overcome the problem that efficient no-envy allocation may not exist at all for economies with non-convex preferences or with production. As different reference bundles lead to different shares, the multiplicity problem over efficient and egalitarian equivalent allocations resurfaces in the choice of the reference bundle. Pazner and Schmeidler (1978) suggests to circumvent the difficulty by focusing only on those reference bundles that are proportional to the total endowment. (Assuming efficiency, this leads to a unique selection.) Sprumont and Zhou (1999) axiomatizes this "Pazner-Schmeidler" rule for exchange economies with convex preferences where the endowment is a finite number of homogeneous infinitely divisible goods.

It is not immediately obvious how to extend the "Pazner–Schmeidler" rule when the endowment is a single heterogeneous good. Consider the division of a contested cake among a group of people who have equal claims on it. The agents may evaluate the value of a piece of the cake along different attributes: its crust, its filling, its weight, the number of strawberries on it, and so on. The challenge is how to make sure that all relevant attributes are proportionally represented in the reference bundle. Moreover, if the parties themselves agree that a criterion should be represented whenever an agent cares about it, is there a way to elicit this strategic information from each party?

We answer these questions under the assumption that each agent can partition the disputed cake into a finite (or countable) number of parcels that he (but not necessarily the other parties) view as homogeneous. The intuition is the following. Each agent divides the cakes in as many parcels as he likes. Equally sized morsels from the same parcel carry the same utility to the agent, so that each parcel is a homogeneous good for the agent. Note that equally sized morsels from two different parcels may carry different utility to him; and, similarly, equally sized morsels from an agent's parcel may give different utilities to another agent. Consider now the common refinement of all the agents' partitions. Each parcel in this new and finer partition is a homogeneous good for each party. This brings us back to the standard setting for the Pazner-Schmeidler rule. Hence, we choose the reference bundle among those that are proportional to this common refinement. Under efficiency, the selection of the reference bundle to define the egalitarian equivalent allocation is again unique.

Clearly, in the search for a procedure to implement the efficient egalitarian equivalent allocation with respect to this special reference bundle, we also need to overcome the difficulty to devise a game in which each agent must announce his own (truthful) partition of the cake. Lying over one's partition may lead to a different reference bundle and hence to a better share for the liar. We provide a simple procedure which implements the desired outcome as a subgame-perfect equilibrium, under the assumption that agents have complete information about their preferences. The procedure is simple in the sense of Thomson (2005). It generalizes a mechanism suggested in Crawford (1979) as ameliorated in Demange (1984). Their mechanism derives an efficient egalitarian equivalent allocation for a finite collection of homogenous goods. Our procedure simultaneously "discovers" the right way to partition the heterogeneous good.

The paper is organized as follows. Section 2 describes our model, which is a standard version of the classical setup for cake division problems. Section 3 proves the existence of efficient egalitarian equivalent allocations for a single heterogeneous good using only measure-theoretic assumptions; the only other existence result we are aware of is more general in scope but requires additional topological assumptions; see Berliant et alii (1992). Section 4 describes the assumptions that define the economic environment over which our procedure can be applied. Section 5 states the implementation result. Long proofs are relegated in the appendix.

2 The model

Our model is an abstraction of the classical problem where a cake (or a piece of land) must be allocated among several agents. There is a measurable space $(\Omega, \mathcal{F})$, where Ω is the object to be divided among the n agents and $\mathcal{F}$ is a σ-algebra over Ω. We say that an element of $\mathcal{F}$ is a *parcel* and that an $\mathcal{F}$-measurable subset of a parcel is a *morsel*, which is a nicer term than "subparcel". Any subset of Ω mentioned in the following is an element of $\mathcal{F}$, and hence a parcel.

For $n \geq 2$, let $N = \{1, 2, \ldots, n\}$ be the (finite) set of agents. Agents have preferences over parcels of Ω. Each agent i is endowed with a utility function $u_i : \mathcal{F} \to \mathbb{R}^+$ that is a nonatomic probability measure on $\mathcal{F}$. (Since preferences are invariant up to a positive rescaling of the utility function, $u_i(\Omega) = 1$ is only a normalization.) A measure u_i is *nonatomic* if, for each parcel A and each x in $(0, u(A))$, there exists another parcel $B \subseteq A$ such that $u_i(B) = x$. Hence, the range of each u_i is the (convex) interval $[0, 1]$.

A utility function u over parcels is *absolutely continuous* with respect to another measure μ over $\mathcal{F}$ if $\mu(A) = 0$ implies $u(A) = 0$ for any parcel A. Clearly, any utility function u_i is absolutely continuous with respect to the measure $\mu = \sum_{i=1}^{n} u_i$. We make the assumption that the utility functions are mutually absolutely continuous; that is, if $u_i(A) = 0$ for some parcel A, then $u_j(A) = 0$ for any agent j. Since agents agree on the null parcels, we say that a parcel has zero (or positive) measure without specifying a measure.

An (ordered) partition $P = (p_1, \ldots, p_k)$ of Ω constituted only by parcels is called a *parceled k-partition.* An *allocation* $X = (x_1, \ldots, x_n)$ is a parceled n-partition, where x_i is the parcel assigned to agent i in N. An allocation X is *efficient* (or *weakly efficient*, respectively) if there exists no other allocation

$Y = (y_1, \ldots, y_n)$ such that $u_i(y_i) \geq u_i(x_i)$ for all i, with the strict inequality holding for some i (or $u_i(y_i) > u_i(x_i)$ for all i). Any efficient allocation is also weakly efficient. The converse is true under our assumption that agents have preferences that are mutually absolutely continuous; see Akin (1995, Lemma 9).

There are several criteria to evaluate the fairness of an allocation. For instance, an allocation X is *proportional* if $u_i(x_i) \geq (1/n)$ for all i; and it is *equitable* if $u_i(x_i) = u_j(x_j)$ for all i and j. These two notions of fairness hinge on the demanding assumption that interpersonal preferences are comparable. The main fairness criteria based on ordinal preferences are two. An allocation X is *envy-free* if $u_i(x_i) \geq u_i(x_j)$ for all i and j, and it is *egalitarian equivalent* (for short, EE) if there exists a *reference parcel* A such that $u_i(x_i) = u_i(A)$ for all i. Any envy-free allocation is proportional, but the converse is true only if $n = 2$.

Under our setup, the following existence results are known. Dubins and Spanier (1961) proves the existence of efficient and proportional allocations for preferences which may not be mutually absolutely continuous. It notes that adding this assumption ensures that all efficient allocations are equitable. Maccheroni and Marinacci (2003) gives sufficient conditions to extend the existence result for proportional allocations when the utility functions are subadditive. Weller (1985) proves the existence of weakly efficient and envy-free allocations; efficiency follows immediately under mutual absolute continuity.

More existence results are known under related setups, which additionally assume that Ω is a subset of $\mathbb{R}^k$. For instance, Stromquist (1980) proves the existence of envy-free allocations for a planar cake using a larger class of preferences, but restricting the set of admissible partitions. Berliant et alii (1992) has several results. It gives a stronger version of Weller's (1985) result assuming that the utility functions are absolutely continuous with respect to the Lebesgue measure. And it proves the existence of efficient and egalitarian equivalent allocations for a general class of preferences that must however be continuous in a complicated topology described in Berliant and Dunz (2004).

3 Existence of efficient EE partitions

This section proves the existence of efficient and egalitarian equivalent allocations in our setup. Contrary to Berliant et alii (1992), we make no topological assumptions so that the proof does not rely on the structure of $\mathbb{R}^k$.

We need a few definitions. Let $u = (u_1, \ldots, u_n)$ be the vector of the n agents' utility functions on the measurable space $(\Omega, \mathcal{F})$. The set $R(u) = \{u(A) : A \in \mathcal{F}\}$ in $\mathbb{R}^n$ is the *range* of u. The range of u spans the vector of utilities that the agents can achieve if they are all given the same parcel. By assumption, each u_i is a nonatomic probability measure. Then, by Lyapunov's convexity theorem, $R(u)$ is a compact and convex subset of $\mathbb{R}^n$.

Let Π be the set of all parceled n-partitions of Ω. The set $RP(u) = \{(u_1(x_1), \ldots, \; u_n(x_n)) : X \in \Pi\}$ in $\mathbb{R}^n$ is the *partition range* of u. The partition range is sometimes called the Individual Pieces Set; see Barbanel (2005). The partition range of u spans the vector of utilities that the agents can achieve by dividing up the cake according to some allocation. Dvoretzky et alii (1951) derives from Lyapunov's convexity theorem a more general result, which implies that the partition range is also a compact and convex subset of $\mathbb{R}^n$.

We also call on the following three lemmata, which assume that u is a (finite) vector of nonatomic probability measures. The first two results correspond to Corollary 1.1 and Lemma 5.3 in Dubins and Spanier (1961), respectively. They do not require preferences that are mutually absolutely continuous.

Lemma 1 *Given an integer k and positive weights $\alpha_1, \ldots, \alpha_k$ with $\sum_j \alpha_j = 1$, there exists a parceled k-partition $X = (x_1, \ldots, x_k)$ such that $u_i(x_j) = \alpha_j$ for all $i = 1, \ldots, n$ and $j = 1, \ldots, k$.*

For $k = n$ and $\alpha_j = 1/n$ for all j, this implies that the partition range of u always contains the point $(1/n, \ldots, 1/n)$. The next lemma, instead, concerns the range of u and implies that it always contains the whole *diagonal*.

Lemma 2 *For any t in $[0, 1]$ there exists a parcel A_t such that $u_i(A_t) = t$ for each i.*

Our last lemma characterize efficient allocations when preference are mutually absolutely continuous; see Barbanel and Zwicker (1997, Theorem 1). Section 7C in Barbanel (2005) discusses the case where mutual absolute continuity does not hold.

Lemma 3 *If preferences are mutually absolutely continuous, an allocation is efficient if and only if it maximizes a convex combination of the utility functions.*

The following is the main result in this section.

Theorem 1 *There exists at least one allocation which is efficient and egalitarian equivalent.*

Proof: Consider the set $RP(u)$. If the vector $\mathbf{1} = (1, \ldots, 1)$ belongs to $RP(u)$, there exists an allocation such that every agent has utility $1 = u_i(\Omega)$.

Clearly, this allocation is efficient and egalitarian equivalent with respect to the reference parcel Ω.

So, assume that the vector $\mathbf{1}$ is not in $RP(u)$. For any t in $[0,1]$, let $\mathbf{t} = (t, \ldots, t)$. By Lemma 1, $RP(u)$ contains $\mathbf{1/n}$ and so it is not empty. By compactness of $RP(u)$, the continuous function$f(\mathbf{t}) = t$ attains a maximum $t^* < 1$. The vector $\mathbf{t}^*$ is in $RP(u)$, so there exists some allocation X^* such that $u_i(x_i^*) = t^*$ for each i. By Lemma 2, there exists some parcel A such that $u_i(A) = t^*$ and thus X^* is an egalitarian equivalent allocation with respect to the reference parcel A.

Moreover, since $\mathbf{t}^*$ is also a boundary point of the convex set $RP(u)$, there is a supporting hyperplane for $RP(u)$ going through $\mathbf{t}^*$. Thus, for some convex combination of weights $(\alpha_1, \ldots, \alpha_n)$, $\sum_i u_i(x_i^*) = \sum_i \alpha_i t_i \geq \sum_i \alpha_i u_i$ for any u in $RP(u)$ and hence by Lemma 3 X^* is also efficient. $\square$

4 An economic environment for EE allocations

The existence result in Theorem 1 is not constructive. In other words, it does not tell us how to find X^*. It does not even try and ask agents if they know what X^* should be. This leads naturally to the question of designing a procedure that generates an efficient and egalitarian equivalent allocation.

Crawford (1979) provides a solution to this problem when there is a finite number of perfectly divisible homogeneous goods, under the assumption that there is complete information about agents' (continuous and strongly monotonic) preferences. Given a *numeraire bundle* x that is desirable for all agents, each agent bids a price for the right to propose the allocation. If everybody accepts the winner's proposal, this is carried out. If an agent refuses the proposal, the final allocation is derived from the equal share rule in which everybody gets $(1/n)$ of the original endowment as follows: the divider gives up a fraction p (equal to his bid) of x that is equally shared among the other agents.

As it turns out, there is a unique price p^* which makes every agent indifferent between the roles of divider and chooser. The procedure generates a final allocation that is efficient and egalitarian equivalent with respect to the reference bundle formed by the union of a fraction $(1/n)$ of the original endowment and a fraction p^* of the numeraire bundle. Demange (1984) improves on this scheme by proposing a version that avoids the infeasible off-equilibrium allocations present in Crawford's (1979) procedure.

When viewed as an implementation result for their economic environment, these results exhibit a limitation. Although the choice of the numeraire bundle affects the final allocation, the procedure assumes that the numeraire bundle is given exogenously, circumventing the problem of how agents come to agree on it. Crawford (1979) points out that either plain money or a bundle proportional to the total endowment are likely to be focal choices for the numeraire bundle. This latter choice defines the Pazner-Schmeidler rule for the selection of the numeraire bundle, recently axiomatized in Sprumont and Zhou (1999).

In our setup with just one heterogenous good, there is no money and it is not clear how to define a proportional bundle. Our contribution in this section is to identify an economic environment where the definition of a proportional bundle should be uncontroversial, making the use of the Pazner-Schmeidler rule intuitively natural. We then extend the Crawford-Demange procedure accordingly and provide a general method to achieve efficient and egalitarian equivalent allocations for a single heterogeneous good.[1]

We enrich the setup in Section 2 with three assumptions, that define the economic environment investigated in this section.

(A1) There is a measure μ on the measurable space $(\Omega, \mathcal{F})$.

(A2) For each i, the utility u_i is absolutely continuous with respect to μ.

(A3) For each i, there exists a finite parceled partition $P^i = (p^i_1, \ldots, p^i_{m_i})$ such that, if two parcels A, B satisfy $A \cup B \subseteq p^i_j$ then

$$\mu(A) = \mu(B) \Rightarrow u_i(A) = u_i(B)$$

Assumption (A1) requires that there is a common "objective" measure. If Ω is a subset of $\mathbb{R}^k$, this might be the Lebesgue measure. Assumption (A2) requires that an agent attaches no utility to sets that have size zero, where the *size* of a parcel is simply its μ-measure. Finally, Assumption (A3) states that there is a partition P^i which divides Ω into m_i parcels, each of which can be considered a homogeneous good for agent i. Assumption (A3) can be considerably relaxed by allowing P^i to be constituted by a countably infinite number of parcels. After few obvious modifications, all of our results would still hold. So it is only in the interest of simplicity that we assume that P^i is finite.

An example may help to assess the import of these three assumptions. Suppose that the object to be allocated is a chocolate chip cookie, which we may think of as a subset of $\mathbb{R}^3$. Presumably, agents care only about the cookie dough or the chocolate chips, although possibly in different guises. Then (A1) is satisfied by taking μ to be the Lebesgue measure on $\mathbb{R}^3$ and (A2) holds for instance if agents attach no utility to morsels with empty interior. Finally, (A3) is satisfied by assuming that each agent partitions the cookie into two sets: the dough and the chips.

While (A1) and (A3) may look restrictive when presented in an abstract setting, we have not been able to find applications in which it is not natural

[1] In this respect, Demange (1984) suggests a lottery-based extension of the procedure, but its demanding assumptions impose that agents be expected utility maximizers and that the whole good be randomly assigned to a single agent off the equilibrium path.

to assume that there is an objective way to measure Ω and that an agent cares differently about the parcels he gets in more than a finite (or countable) number of ways. Even the standard example of an exchange economy with a finite (or countable) number of homogeneous goods implicitly assumes both (A1) and (A3); see Chambers (2005, Section 5). Assumption (A2) is a more technical condition which we need only to ensure the efficiency of a specific family of allocations defined below; see footnote 2 below.

Any parceled finite refinement of P^i defines another partition that satisfies (A3). Ordering such finite partitions by inclusion defines a lattice on the set of parceled partitions of Ω. We denote by $P_1 \vee P_2$ and $P_1 \wedge P_2$ the common coarsening and the common refinement of two partitions P_1 and P_2, respectively.

For definiteness, we associate to each agent i the *canonical partition* P_i^c corresponding to the common coarsening of all finite parceled partitions that satisfy (A3). This is the maximal element in the lattice. Clearly, the canonical partition P_i^c is the only partition among those satisfying (A3) for which two parcels $A \subseteq p_j^i$ and $B \subseteq p_l^i$ $(j \neq l)$ such that $\mu(A) = \mu(B)$ are associated with distinct utilities $u_i(A) \neq u_i(B)$.

Consider now the common refinement $P^c = \bigwedge_{i \in N} P_i^c$ of the canonical partitions for each agent, which we call the *natural partition* for Ω. This finer (parceled) partition extends (A3) to all agents, in the sense that if two parcels A, B satisfy $A \cup B \subseteq p_j$ in P^c then

$$\mu(A) = \mu(B) \Rightarrow u_i(A) = u_i(B), \text{for all } i \text{ in } N. \tag{1}$$

Therefore, the natural partition P^c of Ω divides the original heterogeneous good in a finite set of parcels which each player views as homogeneous. Note that taking the common refinement of arbitrary parceled partitions that satisfy (A3) for all i leads to a common partition finer than P^c for which (2) still holds. The natural partition P^c is simply the smallest one among those that satisfy (2): focusing on it entails no loss of generality.

Let $P = (p_1, \ldots, p_m)$ be the natural partition (or a refinement of it). Then each agent is indifferent among morsels of equal size from a single parcel p_j and we can construct a bundle by adjoining proportional morsels from each parcel of the natural partition. Formally speaking, a *proportional bundle* of size λ obtains if we choose λ in $[0, 1]$ and then pick a morsel e_j from each parcel p_j such that $\mu(e_j) = \lambda\mu(p_j)$ for all $j = 1, \ldots, m$. The proportional bundle is $A = \bigcup_{j=1}^m e_j$. We do not need to specify how a morsel e_j is chosen from the parcel p_j because, as far as $\mu(e_j) = \lambda\mu(p_j)$, each agent is indifferent over the actual morsel chosen.

On the contrary, if we attempt the same construction using an arbitrary partition P, morsels of equal size from the same parcel may carry different utility for an agent; therefore, an exact specification of how morsels are chosen is necessary before an agent can evaluate the outcome of this process.

In other words, in general we need to assume that P is (or refines) the natural partition to ensure that the value of a proportional bundle of size λ to each agent is independent of specific details in the choice of morsels. When this is the case, we say that the utility of the bundle is *process invariant.* In particular, under process invariance, the identity of the *divider* who actually gets to choose the morsels from a parcel is irrelevant.

Suppose again that $P = (p_1, \ldots, p_m)$ is the natural partition (or a refinement) and choose λ in $[0, 1]$. We can repeat the procedure n times and construct n proportional bundles of size λ by choosing (or instructing a divider to choose) morsels of size $\lambda\mu(p_j)$ from each parcel p_j in P. If $\lambda = 1/n$, the divider can pick mutually exclusive morsels from each parcel that jointly exhaust Ω (up to a μ-null set[2]). Doing so generates an allocation that splits Ω into n proportional bundles of equal size. By process invariance, each agent is indifferent among any of these n bundles of size $1/n$. Hence, this allocation is proportional and envy-free; however, in general it is not efficient.

If $\lambda > (1/n)$, the divider can still iterate the procedure and construct n (overlapping) bundles of equal size λ. Since these are not mutually exclusive, the resulting collection of equally sized parcels is not an allocation. More generally, let P be the natural partition (or a refinement) and suppose that we are given weights $\lambda_1, \ldots, \lambda_n$ in $[0, 1]$. We can repeat the procedure n times, instructing a divider to choose at each stage k morsels of size $\lambda_k\mu(p_j)$ from each parcel p_j in P. Again, the outcome is process invariant but, unless $\sum_i \lambda_i = 1$, does not generate an allocation.

The outcome of this construction is an (ordered) collection of n parcels which have sizes proportional to a vector $\boldsymbol{\lambda} = (\lambda_1, \ldots, \lambda_n)$ of positive weights. Even when it is not an allocation, this collection associates each agent i with a bundle of size λ_i. We want to use the utility an agent receives from his bundle as a benchmark for the utility he gets from a given allocation. Since it is based on bundling proportionally-sized morsels from each parcel, we call this collection of parcels a $\boldsymbol{\lambda}$-*proportional benchmark* over the partition P. When $\sum_i \lambda_i = 1$, we assume without loss of generality that the parcels in the collection are mutually exclusive and we speak of a $\boldsymbol{\lambda}$-*proportional benchmark allocation.*

Whereas a proportional allocation must be feasible and make the *utility* of each agent at least $(1/n)$, a proportional benchmark requires only that the *size* of the parcel associated with agent i is proportional to λ_i. In general, the agents' utility for a $\boldsymbol{\lambda}$-proportional benchmark depends both on the partition P used and on the details of the process. When P is the natural partition (or a refinement), however, the utility of the $\boldsymbol{\lambda}$-proportional benchmark is process invariant.

The next section shows how to implement egalitarian equivalent alloca-

[2] By (A2), a residual set of size zero carries no utility for any agent. From now on, we assume without loss of generality that no residual set of size zero is left over.

tions with respect to a reference parcel which is a proportional bundle of size $\lambda^* \geq 1/n$. Equivalently, for each agent the final allocation generates the same utility as a $\lambda^*\mathbf{1}$-proportional benchmark, where $\lambda^*\mathbf{1}$ denotes a vector with all the components equal to λ^*. This is the case of interest when all agents have equal claims on the good to be divided. The name of the next section is a pun on similar titles in the literature, as well as a reminder that in this environment there may be a three-dimensional cookie (made of chocolate chips and dough) at the center of the table.

5 How to cut a cookie fairly

We exhibit a procedure that implements an efficient and egalitarian equivalent allocation when the object to be divided is a heterogeneous good that is infinitely divisible, in the setup of Section 2. The solution concept is subgame perfection. We show that there are unique equilibrium payoffs, with a final allocation that is efficient, proportional and egalitarian equivalent. We assume complete information about preferences, as well as (A1-A3). Any partition of Ω mentioned in the following is formed by parcels; hence, for simplicity, we drop the "parceled" adjective.

Our procedure is inspired by the mechanism described in Demange (1984) to improve the version proposed in Crawford (1979). We have found a few different game forms that do the job. The one presented here is especially expedient because on the equilibrium path the only messages announced are n bids (one for each agent), the final allocation as chosen by an agent (called divider), and $n-1$ "yes" from the remaining agents (called choosers) who accept the divider's proposal; off the equilibrium path, the worst case has a chooser saying "no", after which each agent announces a partition and the final allocation is a suitable proportional benchmark picked by the lowest-bid chooser. The procedure develops in two stages.

Stage 1. Each agent i in N simultaneously announces a bid b_i between $1/n$ and 1. Agents are ordered by decreasing bids using if necessary an arbitrary tie-breaking rule, so that $b_1 \geq b_2 \geq \ldots \geq b_n$. Agent 1 is called the *divider*, while any other agent is a *chooser*.

Stage 2. The last stage consists of several consecutive moves, one for each agent.

Move 0. The divider proposes a feasible allocation X.
Move 1. Chooser n accepts or refuses. If he refuses, each agent i announces[3] a partition P^i and chooser n picks[4] a proportional benchmark allocation X^n

[3] The timing of these announcements is irrelevant for the equilibrium outcome.
[4] The identity of the agent who selects the allocation after a refusal is immaterial for the equilibrium outcome, but nominating chooser n slightly simplifies the proof.

using the partition $\bigwedge_{i=1}^{n} P^i$ and the vector of weights $\boldsymbol{\lambda}^n$ defined below; then the game stops and X^n is the final allocation. If he accepts, move 2 is played.

$\vdots$

Move $i+1$. Chooser $n-i$ accepts or refuses. If he refuses, each agent i announces a partition P^i and chooser n picks a $\boldsymbol{\lambda}^{n-i}$-proportional benchmark allocation X^{n-i} over the partition $\bigwedge_{i=1}^{n} P^i$; then the game stops and X^{n-i} is the final allocation. If he accepts, move $i+2$ is played.

$\vdots$

Move $n-1$. Chooser 2 accepts or refuses. If he refuses, each agent i announces a partition P^i and chooser n picks a $\boldsymbol{\lambda}^2$-proportional benchmark allocation X^2 over the partition $\bigwedge_{i=1}^{n} P^i$; then the game stops and X^2 is the final allocation. If he accepts, the game stops and the final allocation X is the divider's proposal.

If agent i is the first to refuse the divider's proposal, this affects only the choice of the vector $\boldsymbol{\lambda}^i$ of convex weights used by agent n to pick a proportional benchmark allocation. These vectors of convex weights are defined hereafter.

Proportionality weights. To construct the vectors $\boldsymbol{\lambda}^i$ (for $i = 2, \ldots, n$) of proportionality coefficients, we define a pecking order for agents. Suppose that players sit around a circle arranged in standard (clockwise) increasing order. When i refuses, the circle is walked counterclockwise starting from i and the pecking order is $i < i-1 < \ldots < 2 < n < n-1 < \ldots < i+1$. Let $\sigma^k(i)$ denote the k-th agent in this pecking order. We are ready to define the vector $\boldsymbol{\lambda}^i = (\lambda_1^i, \ldots, \lambda_n^i)$ used if i refuses the proposed allocation. The refusing agent $i = \sigma^1(i)$ is assigned the proportionality coefficient $\lambda_i^i = b_i$. Following the pecking order, a later chooser $\sigma^i(k)$ is assigned the coefficient

$$\lambda^i_{\sigma^i(k)} = \begin{cases} b_{\sigma^i(k)} & \text{if } \sum_{s=1}^{k} b_{\sigma^i(s)} \leq 1 \\ \max\left(0, 1 - \sum_{s=1}^{k-1} b_{\sigma^i(s)}\right) & \text{otherwise.} \end{cases}$$

In other words, later choosers are assigned a proportionality coefficient equal to their bid when this is still a feasible convex weight, or otherwise its truncation to 1. Whenever there is a refusal, the divider is assigned the (same) coefficient $\lambda_1^i = \max(0, 1 - \sum_{s=2}^{n} b_s)$, regardless of the refuser i's identity; for short, we denote this by β_1. For each i, the positive weights in $\boldsymbol{\lambda}^i$ add up to 1, so that (up to a set of size zero, irrelevant by (A2)) the $\boldsymbol{\lambda}^i$-proportional benchmark can also be made into an allocation exhausting the whole good.

This concludes the description of the procedure. We prove two results in this section. The first one is that the equilibrium allocations are proportional, in the sense that each agent receives a parcel which carries utility of at least $1/n$ to him. The second and main result is that this procedure has unique equilibrium payoffs and yields an allocation X that is efficient and egalitarian equivalent.

We need a few definitions. Let $\mathcal{A}$ be the set of all possible allocations. Given a (parceled) partition P and a vector $\boldsymbol{\lambda}$ of convex weights, let $\Pi(\boldsymbol{\lambda}|P)$ denote the set of $\boldsymbol{\lambda}$–proportional benchmark allocations over P. Clearly, this set contains (uncountably) many possible allocations. However, when $P = P^c$ is the natural partition (or a refinement), by process invariance each agent is indifferent among the allocations in $\Pi(\boldsymbol{\lambda}|P^c)$; hence, the utility that each agent i obtains by a $\boldsymbol{\lambda}$-proportional benchmark allocation over P is well-defined.

Theorem 2 *Any equilibrium allocation for the procedure described above is proportional.*

Proof: We prove the stronger statement that each agent has a strategy that ensures that he receives the utility level $\bar{u} = (1/n)$ associated with the proportional benchmark allocation $\Pi(\frac{1}{n}\mathbf{1}|P^c)$.

The strategy is the following. The agent bids $(1/n)$ in the first stage. In the second stage, if he is the divider, he chooses the $\frac{1}{n}\mathbf{1}$-proportional benchmark allocation over P^c; if he is a chooser, he refuses any proposed allocation; whenever asked for, he announces the partition P^c.

If the agent ends up being the divider and his proposal is accepted, the final allocation is in $\Pi(\frac{1}{n}\mathbf{1}|P^c)$. If it is refused by an agent j, announcing $P^1 = P^c$ makes sure that $\bigwedge_{i=1}^n P^i = P^c$ (or a refinement of it). Hence, the final allocation is in $\Pi(\boldsymbol{\lambda}^j|P^c)$. But $\boldsymbol{\lambda}^j = \frac{1}{n}\mathbf{1}$, because if the divider's bid is $(1/n)$ all other bids are equal to $(1/n)$. Thus, the final allocation is still in $\Pi(\frac{1}{n}\mathbf{1}|P^c)$. The divider gets utility $\bar{u} = (1/n)$.

If instead the agent is chooser i, every agent $n, n-1, \ldots, i+1$ before him in the pecking order has bid $b_j = (1/n)$. So, if a chooser $j = i, \ldots, n$ refuses, the vector $\boldsymbol{\lambda}^j$ has $\lambda^j_k = (1/n)$ for $k = i, \ldots, n$. Hence, if a refusal occurs up to (and including) i's move, the final allocation is in $\Pi(\boldsymbol{\lambda}^j|P^c)$ which by a similar reasoning to the above yields utility $\bar{u} = (1/n)$ to agent i (as well as to $j = i+1, \ldots, n$). □

The following is the main result in this section. Its proof, including a few lemmata, can be found in the appendix. We assume that a chooser who is indifferent always prefers the move that keeps his play simpler and, subordinately, the move that ends the game sooner; for instance, a tie between accepting or refusing a proposal is broken by accepting because this avoids him having to announce a partition. The statement uses the following piece of notation. For $\Pi(\boldsymbol{\lambda}|P^c)$ being the set of $\boldsymbol{\lambda}$–proportional benchmark allocations over the natural partition P^c (or a refinement), let $\pi_i(\boldsymbol{\lambda}|P^c)$ denote the set of parcels that these allocations may assign to agent i. By process invariance, $u_i[\pi_i(\boldsymbol{\lambda}|P^c)]$ is well-defined. Finally, the equilibrium concept used is subgame perfection; for short, we speak simply of "equilibrium".

Theorem 3 *The procedure described above has unique equilibrium payoffs, with final allocations that are efficient and egalitarian equivalent. In every equilibrium allocation, each agent i is indifferent between the parcel he receives and getting $\pi_i(\lambda^*\mathbf{1}|P^c)$, where $\lambda^* = \max\{\lambda :$ there exists $Y \in \mathcal{A}$ with $u_i(y_i) \geq u_i[\pi_i(\lambda\mathbf{1}|P^c)]$ for each $i\}$.*

A Appendix

We recall and extend a piece of our notation to make the proofs in this appendix easier to read. Let $\Pi(\boldsymbol{\lambda}|P)$ denote the set of $\boldsymbol{\lambda}$–proportional benchmark allocations over P and let $\pi_i(\boldsymbol{\lambda}|P)$ be the set of parcels that these allocations may assign to agent i. Since $\pi_i(\boldsymbol{\lambda}|P)$ depends only on the i-th component of $\boldsymbol{\lambda}$, we abuse notation and write simply $\pi_i(\lambda_i|P)$ when we need to highlight i's proportionality coefficient. If $P = P_i^c$ (or a refinement, and in particular the natural partition P^c), the agent is indifferent over any element in this set and we denote his utility by $v_i(\lambda_i|P_i^c) = u_i[\pi_i(\lambda_i|P_i^c)]$. Clearly, $v_i(\lambda_i|P_i^c) = v_i(\lambda_i|P)$ whenever P is a refinement of P_i^c; for simplicity, in the following we use the common refinement $P^c = \bigwedge_{i\in N} P_i^c$ so that we have $v_i(\lambda_i|P^c) = v_i(\lambda_i|P_i^c)$ for each i.

By (A2) and the nonatomicity of u_i, $v_i(\lambda_i|P^c)$ is a strictly increasing continuous function of λ_i. We extend the function $v_i(\lambda_i|\cdot)$ to an arbitrary partition P by letting $v_i(\lambda_i|P) = \max u_i[\pi_i(\lambda_i|P)]$ be the maximum utility that i can obtain by a parcel in $\pi_i(\lambda_i|P)$. (This maximum is well defined because the partition range of u is compact.) In other words, $v_i(\lambda_i|P)$ is the maximum utility that an agent can get by any parcel associated with a $\boldsymbol{\lambda}$–proportional benchmark allocation over a partition P, when his proportionality coefficient if λ_i.

We assume that a chooser who is indifferent always prefers the move that keeps his play simpler and, subordinately, the move that ends the game sooner. Hence, when indifferent, he chooses acceptance over refusing the divider's proposal. When he is indifferent between refusing now and having some later chooser refuse, he prefers to refuse now.

We begin by proving that there exists a (subgame perfect) equilibrium in which the players' announcements after any refusal induce the natural partition (or a refinement of it); that is, $\bigwedge_{i\in N} P^i = P^c$. For instance, this may occur if each player truthfully announces his canonical partition P_i^c; then $\bigwedge_{i\in N} P_i^c = P^c$. Note that each player *alone* can enforce (a refinement of) his own canonical partition by making the announcement $P_i = P_i^c$. Therefore each player i holds a form of *veto power* on the partition used after a refusal:

exercising this, he can secure a payoff $v_i(\lambda_i^k|P_i^c) = v_i(\lambda_i^k|P^c)$ after a refusal from agent k.

In general, there may be several different equilibrium strategy profiles supporting the use of the natural partition after a refusal. To see why, suppose that there exists an equilibrium in which at least two agents announce the canonical partition P^c if some chooser refuses the divider's proposal. Then $\bigwedge_{i=1}^n P^i = P^c$ (or a refinement) and thus, by process invariance, announcing a different partition is payoff-equivalent for the other players. Since we are interested only in the uniqueness of equilibrium payoffs, we can afford not to chase all the equilibrium strategy profiles. It is convenient to focus on the case where the common refinement of the partitions announced after a refusal is P^c; for short, we say that the common partition after a refusal is P^c.

Lemma 4 *Suppose that the common partition after a refusal is P^c. Then a proposed allocation X is accepted if and only if $u_i(x_i) \geq v_i(b_i|P^c)$ for each $i \geq 2$.*

Proof: The proof is by backwards induction. Suppose that the game has reached agent 2. Then he can accept x_2 or choose a $\boldsymbol{\lambda}^2$-proportional benchmark allocation X^2 over the partition P^c. By our tie-breaking rule, he accepts if and only if $u_2(x_2) \geq v_2(b_2|P^c)$. Anticipating this, if 2 is going to accept, then agent 3 accepts x_3 if and only if $u_3(x_3) \geq v_3(b_3|P^c)$. The result follows by induction. □

Lemma 5 *Suppose that the common partition after a refusal is P^c. If the proposed allocation X is refused, the final allocation is a $\boldsymbol{\lambda}^n$-proportional benchmark X^n chosen by player n (who moves second).*

Proof: Suppose that X is not accepted and let i be the smaller index in $\{2, \ldots, n\}$ such that $u_i(x_i) < v_i(b_i|P^c)$. By Lemma 4, i refuses X if the game reaches him. Anticipating this, agent $i+1$ (if any) has a choice between getting a proportional bundle of size λ_{i+1}^i over the partition P^c (if he accepts and lets the game reach i) or a proportional bundle of size b_{i+1} over the same partition P^c (if he refuses). But $\lambda_{i+1}^i \leq b_{i+1}$, so the first option can never lead to a higher utility for $i+1$. Therefore, by our tie-breaking rule, he prefers to stop the game. The result follows by induction. □

The next two lemmata describe the optimal strategy for the divider, given a vector $\mathbf{b} = (b_1, \ldots, b_n) = (b_1, \mathbf{b}_{-1})$ of bids from the first stage and conditional on the common partition after a refusal being P^c. Whenever there is a refusal, the proportionality coefficient for the divider is $\beta_1 = \max(0, 1 - \sum_{s=2}^n b_s)$, regardless of the refuser i's identity. Suppose that, if the proposed allocation X is refused by a chooser, the common partition is P^c. Then the divider's utility is $v_1(\beta_1|P^c)$ by process invariance. On the other hand, by Lemma 4 and process invariance again, the allocation X is accepted if and only if $u_i(x_i) \geq v_i(b_i|P^c)$ for every $i \geq 2$. Let

$$\mathcal{A}(\mathbf{b}_{-1}) = \{X \in \mathcal{A} : \ u_i(x_i) \geq v_i(b_i|P^c) \text{ for each } i \geq 2\}$$

be the set of allocations which are accepted.

Lemma 6 *Suppose that the common partition after a refusal is P^c. For any $\mathbf{b}_{-1}$ such that $\mathcal{A}(\mathbf{b}_{-1}) \neq \emptyset$ there exists an allocation which maximizes $u_1(x_1)$. A maximizing allocation $X^*(\mathbf{b}_{-1})$ satisfies $u_i(x_i^*) = v_i(b_i|P^c)$ for each $i \geq 2$.*

Proof: Let $S(\mathbf{b}_{-1}) = [0,1] \times [v_2(b_2|P^c), 1] \times [v_n(b_n|P^c), 1]$ denote the cartesian product of n intervals. This is a compact and convex subset of $\mathbb{R}^n$. Similarly, the partition range $RP(u)$, which spans the vector of utilities that the agents can achieve under a feasible allocation, is a nonempty, compact and convex subset of $\mathbb{R}^n$.

When $\mathcal{A}(\mathbf{b}_{-1}) \neq \emptyset$, there exists at least one allocation which maps to a vector of utilities in $S(\mathbf{b}_{-1})$. Hence, the intersection of $RP(u)$ and $S(\mathbf{b}_{-1})$ is not empty. As this intersection is also compact (and convex), there exists (at least) an allocation $X^*(\mathbf{b}_{-1})$ which maximizes $u_1(x_1)$.

Now, suppose that at X^* there is some $i \geq 2$ such that $u_i(x_i^*) > v_i(b_i|P^c)$. By the nonatomicity of u_i, we can always cut away a morsel from x_i^* and reduce the utility of i down to $v_i(b_i|P^c)$, transferring the morsel to agent 1's parcel. By mutual absolute continuity of preferences, this strictly increases the utility of agent 1 and therefore X^* cannot be optimal. Therefore, at an optimal allocation X^*, the equality $u_i(x_i^*) = v_i(b_i|P^c)$ must hold for each $i \geq 2$. □

Given $\mathbf{b}_{-1}$ and conditional on the common partition after a refusal being P^c, the divider faces the choice of selecting an allocation which is accepted by everybody or another allocation which is eventually refused. In this second case, his utility is $v_1(\beta_1|P^c)$. Clearly, if $\mathcal{A}(\mathbf{b}_{-1}) = \emptyset$, there exists no acceptable allocation X^* so the divider ends up with $v_1(\beta_1|P^c)$. Instead, if $\mathcal{A}(\mathbf{b}_{-1}) \neq \emptyset$, he proposes $X^*(\mathbf{b}_{-1})$ if and only if $u_1(x_1^*) \geq v_1(\beta_1|P^c)$. The next lemma summarizes this. Let $f(\mathbf{b}_{-1}) = \max\{u_1(x_1) : \ X \in \mathcal{A}(\mathbf{b}_{-1})\}$, with the usual clause that $f(\mathbf{b}_{-1}) = -\infty$ if $\mathcal{A}(\mathbf{b}_{-1}) = \emptyset$. By Lemma 6, this is well-defined and moreover, on $\mathcal{A}(\mathbf{b}_{-1}) \neq \emptyset$,

$$f(\mathbf{b}_{-1}) = \max\{u_1(x_1) : \ X \in \mathcal{A} \text{ and } u_i(x_i) = v_i(b_i|P^c) \text{ for each } i \geq 2\}. \quad (2)$$

Lemma 7 *Suppose that the common partition after a refusal is P^c. For any $\mathbf{b}_{-1}$, agent 1 proposes an acceptable allocation $X^*(\mathbf{b}_{-1})$ and the vector of final equilibrium payoffs is*

$$(f(\mathbf{b}_{-1}), v_2(b_2|P^c), \ldots, v_n(b_n|P^c))$$

if and only if $f(\mathbf{b}_{-1}) \geq v_1(\beta_1|P^c)$. *Otherwise, agent 1's proposal is refused by agent* n *and the vector of final equilibrium payoffs is*

$$(v_1(\beta_1|P^c), v_2(\lambda_2^n|P^c), \ldots, v_n(\lambda_n^n|P^c)),$$

with $\lambda_i^n \leq b_i$ *for each* $i \geq 2$.

The next lemma notes two useful properties for $f(\mathbf{b}_{-1})$.

Lemma 8 *Let* D *be the interior set of* $\{\mathbf{b}_{-1} : \mathcal{A}(\mathbf{b}_{-1}) \neq \emptyset\}$. *The function* $f(\mathbf{b}_{-1})$ *is (component-wise) strictly decreasing and continuous on* D.

Proof: By definition, $f(\mathbf{b}_{-1}) \geq 0$ on D. Recall that $v_i(b_i|P^c)$ is a strictly increasing and continuous function of b_i for each i. By strict monotonicity, $b_i < b'_i$ for some $i \geq 2$ implies $v_i(b_i|P^c) < v_i(b'_i|P^c)$. If $(b_2, \ldots, b_i \ldots, b_n)$ and $(b_2, \ldots, b'_i \ldots, b_n)$ are in D, by (2) and mutual absolute continuity of preferences, $f(b_2, \ldots, b_i \ldots, b_n) > f(b_2, \ldots, b'_i \ldots, b_n)$.

To prove continuity, let $\mathbf{u} = (u_1, \ldots, u_n)$ denote a vector of utilities for each player and observe that

$$f(\mathbf{b}_{-1}) = \max\{u_1 : \mathbf{u} \in RP(u) \text{ and } u_i = v_i(b_i|P^c) \text{ for each } i \geq 2\}.$$

For $\mathbf{b}_{-1}$ in D, let $\Gamma(\mathbf{b}_{-1}) = \{\mathbf{u} \in RP(u) : u_i = v_i(b_i|P^c) \text{ for each } i \geq 2\}$ be the correspondence mapping $\mathbf{b}_{-1}$ into $RP(u)$. Since $RP(u)$ is nonempty, compact and convex, the continuity of each v_i implies that $\Gamma(\mathbf{b}_{-1})$ is compact-valued and continuous. Hence, by the Maximum theorem, the function $f(\mathbf{b}_{-1}) = \max\{u_1 : \mathbf{u} \in \Gamma(\mathbf{b}_{-1})\}$ is continuous. □

The next lemma pins down equilibrium behavior, under the assumption that the common partition after a refusal is P^c.

Lemma 9 *Suppose that the common partition after a refusal is* P^c. *The only possible equilibrium move in the first stage is that everyone makes the same bid* $b^* = \max\{\lambda : X \in \mathcal{A} \text{ and } u_i(x_i) \geq v_i(\lambda|P^c) \text{ for each } i\}$.

Proof: First, we show that b^* is well-defined. Let $\times_{i=1}^n [v_i(\lambda|P^c), 1]$ be the cartesian product of the n intervals $[v_i(\lambda|P^c), 1]$, for $i = 1, \ldots, n$. For any λ in $[1/n, 1]$, let $C(\lambda) = \times_{i=1}^n [v_i(\lambda|P^c), 1] \bigcap RP(u)$. By the proof of Theorem 2, $C(1/n) \neq \emptyset$. Moreover, $\lambda_1 > \lambda_2$ implies $C(\lambda_1) \subset C(\lambda_2)$. Since $C(\lambda)$ is a decreasing collection of nested compact subsets, the finite intersection property implies that, for $\lambda^* = \sup\{\lambda : C(\lambda) \neq \emptyset\}$, $C(\lambda^*) \neq \emptyset$. Clearly, λ^* is equal to b^*.

Second, we show that if everybody bids b^* then the divider proposes an acceptable allocation. By Lemma 7, the divider does so if and only if $f(\mathbf{b}_{-1}) \geq v_1(\beta_1|P^c)$. By Theorem 2, $f(\mathbf{b}_{-1}) \geq v_1(1/n|P^c)$. Moreover, $\beta_1 \leq 1/n$ and $v_1(\cdot|P^c)$ is an increasing function, so $f(\mathbf{b}_{-1}) \geq v_1(1/n|P^c) \geq v_1(\beta_1|P^c)$.

Third, we show that an equilibrium move in the first stage cannot have $b_i < b_1$ for some i. Since $b_1 \geq b_i$ for each i by construction, this implies that the equilibrium bids must all be equal. Indeed, suppose $b_i < b_1$ for some i. There are two cases. If $\beta_1 > 0$, then $1 - \beta_1 = \sum_{i=2}^{n} b_i < 1$. By choosing a bid b'_i such that $b_i < b'_i < b_1$ and $b'_i + \sum_{j \neq 1,i} b_j < 1$, i remains chooser and strictly increases his payoff to $v_i(b'_i|P^c)$. If $\beta_1 = 0$, the inequality in the paragraph above strengthens to $f(\mathbf{b}_{-1}) \geq v_1(1/n|P^c) > v_1(\beta_1|P^c)$ by the strict monotonicity of v_1. Then $\mathbf{b}_{-1}$ is in D and by Lemma 8 agent i can find a bid b'_i with $b_i < b'_i < b_1$ and $f(b_2, \ldots, b'_i \ldots, b_n) > v_1(\beta_1|P^c)$. The divider is still agent 1 who chooses an acceptable allocation, which strictly increases i's payoff to $v_i(b'_i|P^c)$. In either case, b_i is not a best reply.

The last paragraph implies that in equilibrium all bids are equal to a common value b. We now prove that $b = b^*$. Let $w_i(b) = \max\{u_i(x_i) : X \in \mathcal{A}$ and $u_j(x_j) \geq v_j(b|P^c)$ for $j \neq i\}$. In other words, let $w_i(b)$ be the payoff to i if he is the divider and everybody has made the same bid b. Clearly, $w_i(b) < v_i(b|P^c)$ if and only if $b > b^*$. Suppose that the common equilibrium bid is $b < b^*$. By the continuity of w_i, chooser i can slightly raise his bid to b' and become the divider, which gets him a payoff $w_i(b') > v_i(b|P^c)$ so that b is not a best reply. Now, suppose that the common equilibrium bid is $b > b^*$. Then the divider's payoff is $w_1(b) < v_1(b|P^c)$. Again, by continuity of w_1, he can slightly lower his bid to b' and become a chooser, which gets him a payoff $v_1(b') > w_1(b|P^c)$ so that b is not a best reply. Hence, the only possible equilibrium must have everybody bidding b^*. Using Lemma 6 and the fact that $w_i(b) < v_i(b|P^c)$ if and only if $b > b^*$, it follows that when everybody bids b^* each agent i (regardless of his role as divider or chooser) receives a payoff $v_i(b^*|P^c)$ and thus he is indifferent between the two roles. Therefore, everybody bidding b^* is the (only) equilibrium. □

Lemma 10 *Suppose that the common partition after a refusal is P^c. Then the only possible equilibrium outcome is efficient.*

Proof: The proof of Lemma 9 shows that the divider proposes an acceptable allocation that satisfies (2). □

The next lemma establishes the existence of an equilibrium in which the common partition after a refusal is P^c. Although there exist other (perhaps more intuitive) equilibria, this particular one provides the stepping stone for the proof of Theorem 3.

Lemma 11 *There exists an equilibrium in which the common partition after a refusal is P^c.*

Proof: Suppose that the common partition after a refusal is P^c. (We do not claim yet that this is part of an equilibrium.) Let $\bar{u}_i$ be the payoff obtained by agent i. Each agent i has a limited form of veto power, because he can enforce a (refinement) of his canonical partition P_i^c. Moreover, $v_i(\lambda|P_i^c) = v_i(\lambda|P^c)$ for any λ. Hence, if a refusal occurs, a player can indifferently announce P_i^c or P^c (or any other refinement of P_i^c) and secure the payoff $\bar{u}_i$. We assume that he announces P^c.

By Lemma 10, the only possible equilibrium outcome is efficient and produces utility $\bar{u}_i$ for each agent i. If this were not an equilibrium, there would be some agent k who could profitably deviate and obtain a utility $u_k > \bar{u}_k$. By efficiency, however, this implies that some other agent i should then obtain a utility $u_i < \bar{u}_i$, contradicting the fact that he can secure the payoff $\bar{u}_i$ by announcing P^c after a refusal. This establishes the claim. □

Proof of Theorem 3. The same argument of Lemma 11 shows that there may no exist equilibrium payoffs different from those achieved when the common partition after a refusal is P^c. This establishes uniqueness. Lemma 10 and Lemma 9 prove that the equilibrium outcome is efficient and egalitarian equivalent, respectively. The proof of Lemma 9 also shows that each agent i is indifferent between the parcel received and the $\lambda^*\mathbf{1}$-proportional benchmark allocation. □

References

[1] E. Akin (1995), "Vilfredo Pareto cuts the cake", *Journal of Mathematical Economics* **24**, 23–44.

[2] J. Barbanel (2005), *The Geometry of Efficient Fair Division*, Cambridge University Press.

[3] J. Barbanel and W.S. Zwicker (1997), "Two applications of a theorem of Dvoretzky, Wald, and Wolfowitz", *Theory and Decision* **43**, 203–207.

[4] M. Berliant, W. Thomson and K. Dunz (1992), "On the fair division of a heterogeneous commodity", *Journal of Mathematical Economics* **21**, 201–206.

[5] M. Berliant and K. Dunz (2004), "A foundation of location theory: Existence of equilibrium, the welfare theorems, and core", *Journal of Mathematical Economics* **40**, 593–618.

[6] C.P. Chambers (2005), "Allocation rules for land division", *Journal of Economic Theory* **121**, 236–258.

[7] V.P. Crawford (1979), "A procedure for generating Pareto-efficient egalitarian-equivalent allocations", *Econometrica* **47**, 49–60.

[8] G. Demange (1984), "Implementing efficient egalitarian equivalent allocations", *Econometrica* **52**, 1167–1177.

[9] A. Dvoretzky, A. Wald, and A. Wolfowitz (1951), "Relations among certain ranges of vector measures", *Pacific Journal of Mathematics* **1**, 59–74.

[10] L.E. Dubins and E.H. Spanier (1961), "How to cut a cake fairly", *American Mathematical Monthly* **68**, 1–17.

[11] D. Foley (1967), "Resource allocation and the public sector", *Yale Economic Essays* **7**, 45–98.

[12] G. Gamow and M. Stern (1958), *Puzzle–Math*, New York: The Viking Press.

[13] F. Maccheroni and M. Marinacci (2003), "How to cut a pizza fairly: Fair division with decreasing marginal evaluations", *Social Choice and Welfare* **20**, 457–65.

[14] T.J. Miceli and C.F. Sirmans (2000), "Partition of real estate; or breaking up is (not) hard to do", *Journal of Legal Studies* **29**, 783–796.

[15] E.A. Pazner and D. Schmeidler (1978), "Egalitarian equivalent allocations: A new concept of economic equity", *Quarterly Journal of Economics* **92**, 671–687.

[16] Y. Sprumont and L. Zhou (1999), "Pazner-Schmeidler rules in large societies", *Journal of Mathematical Economics* **31**, 321–339.

[17] H. Steinhaus (1948), "The problem of fair division", *Econometrica* **16**, 101–104.

[18] W. Stromquist (1980), "How to cut a cake fairly", *American Mathematical Monthly* **87**, 640–644.

[19] W. Thomson (2005), "Divide-and-permute", *Games and Economic Behavior* **52**, 186–200.

[20] D. Weller (1985), "Fair division of a measurable space", *Journal of Mathematical Economics* **14**, 5–17.