To Give or Not to Give: Fair Division for Single Minded Valuations[*†]

Simina Brânzei
Hebrew University
of Jerusalem
Jerusalem, Israel

Yuezhou Lv
Tsinghua University
Beijing, China

Ruta Mehta
University of Illinois
Urbana-Champaign
Champaign, USA

Abstract

Single minded agents have strict preferences, in which a bundle is acceptable only if it meets a certain demand. Such preferences arise naturally in scenarios such as allocating computational resources among users, where the goal is to fairly serve as many requests as possible. In this paper we study the fair division problem for such agents, which is complex due to discontinuity and complementarities of preferences.

Our solution concept—the *competitive allocation from equal incomes* (CAEI)—is inspired from market equilibria and implements fair outcomes through a pricing mechanism. We study existence and computation of CAEI for multiple divisible goods, discrete goods, and cake cutting. Our solution is useful more generally, when the players have a target set of goods, and very small positive values for any bundle other than their target set.

1 Introduction

The question of dividing scarce resources among multiple participants in a way that is fair has remained a pressing question that intrigued humans for a long time, be it for dividing land among citizens[1] or allocating organizational resources to its members. The formal study of *fair division* as we know it started during World War II with works of Steinhaus, Knaster and Banach [Steinhaus, 1951]. Much of the earlier work focused on one heterogeneous divisible good, i.e. the cake cutting problem [Brams and Taylor, 1996; Robertson and Webb, 1998], while other models and notions of fairness were studied later in a growing body of literature [Moulin, 2004]. Fair division has been recently studied in the computer science community, as problems in resource allocation and fair division in particular are arguably relevant for the design of multiagent systems, including manufacturing and scheduling, airport traffic, and industrial procurement [Chevaleyre *et al.*, 2006; Procaccia, 2013].

Prominent examples of fair division models that surfaced recently in real scenarios include problem of allocating computational resources (e.g. CPU, memory, bandwith) among the users of a system [Ghodsi *et al.*, 2011], or the problem of allocating university courses to students in a way that is fair and efficient. The latter motivated the introduction of a notion of fairness known as A-CEEI [Budish, 2011], which approximates the well known ideal notion of fairness from economics, the *competitive equilibrium from equal incomes* [Foley, 1967; Varian, 1974]. Othman *et al.* [2014] studied the complexity of A-CEEI with pessimistic conclusions for general preferences. Nevertheless, the A-CEEI solution is used to allocate courses to students at the Wharton Business School at the University of Pennsylvania.

The largest body of the literature on fair division models for multiple divisible and indivisible goods, including cake cutting, focuses on additive valuations, which capture *perfect substitutes*, i.e. goods that can replace each other in consumption, such as Pepsi and Coca-Cola. However, many real allocation problems have complementarities in the preferences to various degrees. For example, if the user of a computer system wants to run a specific task, they may need 2 units of CPU and 5 units of RAM. It is not useful if the user gets only 1 unit of CPU but 6 or more units of RAM – they simply cannot run the task because CPU is a bottleneck resource.

An important scenario with complements is that of *single minded* valuations, where the agent values a particular bundle and nothing less, and anything extra does not add to the value. Such valuations arise naturally, for example assembling a bike requires a set of parts, or an agent's computational task can finish if and only if it is allocated a required bundle of resource. The latter example represents basis for more complex models

[*]The full version of the paper can be found on the archive at http://arxiv.org/abs/1602.09088.

[†]We are grateful to Erel Segal-Halevi for useful feedback. The authors aknowledge support from the Danish National Research Foundation and The National Science Foundation of China (under the grant 61361136003) for the Sino-Danish Center for the Theory of Interactive Computation and from the Center for Research in Foundations of Electronic Markets (CFEM), supported by the Danish Strategic Research Council. A part of this work was done while Simina Brânzei and Ruta Mehta were visiting the Simons Institute for the Theory of Computing. Simina Brânzei was also supported by ISF grant 1435/14 administered by the Israeli Academy of Sciences and Israel-USA Bi-national Science Foundation (BSF) grant 2014389, as well as the I-CORE Program of the Planning and Budgeting Committee and The Israel Science Foundation.

[1]Divide-and-choose is mentioned already in the Bible, in the Book of Genesis (chapter 13)

that take into account the dynamics over time. Due to their immense applicability, there is an extensive body of literature on such single minded valuations in areas such as auctions and exchange markets (see, e.g., [Briest *et al.*, 2005; Nisan *et al.*, 2007; Ledyard, 2007; Robu *et al.*, 2013; Feldman and Lucier, 2014; Feige *et al.*, 2015]), however work on fair division with such valuations is largely missing.

1.1 Our Contribution

We study fair division among single minded agents for three main scenarios, namely multiple divisible goods, cake cutting, and discrete goods. Our main solution concept—the "competitive allocation from equal incomes" (CAEI)—is inspired from market equilibria and implements fair outcomes through a pricing mechanism. The CAEI solution can be seen as a relaxation of the standard competitive equilibrium from equal incomes and is a more intuitive notion for discrete resources than the latter. However many of our results, most notably for multiple divisible goods, carry over to the competitive equilibrium from equal incomes solution. The CAEI solution concept is sandwhiched between the competitive equilibrium from equal incomes and envy-freeness, and it prunes (via the prices), the least desirable envy-free allocations. Nevertheless, CAEI exists if and only if envy-free allocations exist in all the fair division models we consider.

For multiple divisible goods and cake cutting we show existence of CAEI, while for discrete goods we give a succinct characterization of instances that admit this solution; then we give an efficient algorithm to find one in all cases. Maximizing social welfare turns out to be NP-hard in general, however we obtain efficient algorithms for a number of settings, (i) divisible and discrete goods when the number of different *types* of agents is a constant, (ii) cake cutting with contiguous demands and (iii) cake cutting with constant number of agents.

Our results also carry over to valuation models where the agents have a desired target set (e.g. the set of parts of a bike), while having very small positive value (ϵ) for any bundle that does not contain their desired set. For such scenarios (which are compatible with experimental evidence that people like to accumulate "stuff", even if unneeded), the solution computed is an ϵ-CAEI.

2 Background

We begin by presenting the model and solution concept for the most general setting. Let $N = \{1, \ldots, n\}$ be a set of agents and $\mathcal{R}$ a set of resources, where $\mathcal{R}$ is a compact subset of a Euclidean space. Each agent i is equipped with a valuation function V_i over the resources, such that for each subset $S \subseteq \mathcal{R}$, $V_i(S)$ represents the valuation of agent i for bundle S. The goal is to allocate the resources among the agents in a way that is fair.

Agent i is said to be *single minded* if there exists a bundle $D_i \subseteq \mathcal{R}$ that i cannot be happy without; that is, for any other bundle $S \subseteq \mathcal{R}$ we have that $V_i(S) = 1$ if $D_i \subset S$ and $V_i(S) = 0$ otherwise.

An *allocation* $\mathbf{x} = (\mathbf{x}_1, \ldots, \mathbf{x}_n)$ is a partition of the set $\mathcal{R}$ of resources such that $\mathbf{x}_i \subseteq \mathcal{R}$ is the bundle received by agent i, the bundles are not intersecting and add up to the whole space.

Our solution concept—the "competitive allocation from equal incomes" (CAEI)—is inspired from market equilibria and implements fair outcomes through a pricing mechanism. More formally, each agent is given an artificial unit of currency by the center. It can be used to acquire goods, and has no intrinsic value. The agent wants to spend its budget to acquire a bundle of items that maximizes its utility. A CAEI outcome is defined as a tuple $\langle p, \mathbf{x}\rangle$, where $\mathbf{x}$ is an allocation and $p : \mathcal{R} \to \mathbb{R}_+$ is an integrable, non-negative price density function, such that for each subset $S \subseteq \mathcal{R}$, $p(S) = \int_{x\in S} p(x)dx$.

Formally, $\langle p, \mathbf{x}\rangle$ is a CAEI solution if and only if:

- For all $i \in N$, $\mathbf{x}_i$ maximizes agent i's utility given prices p and its unit budget.
- All the resources are allocated: $\bigcup_{i=1}^{n} \mathbf{x}_i = \mathcal{R}$.
- No agent overspends: $p(\mathbf{x}_i) \leq 1$ for all $i \in N$.

The difference between the CAEI solution and the competitive equilibrium from equal incomes (CEEI), is that under CEEI, the agents are additionally required to spend all of their budget. Since we are in a fair division setting, neither the center nor the agents have value for the units of currency, thus the requirement of spending all the budgets can be relaxed naturally. More importantly, in several of the scenarios we study, the CEEI solution is not as tractable due to discontinuity of the valuations and/or discrete goods.

Note that every CAEI allocation is envy-free, since if an agent envied another agent's bundle it could just buy it instead, as the endowments are equal; the prices provide a mechanism for indicating the interest level in the different resources. However, since the valuations are discontinuous, CAEI outcomes are not necessarily efficient. Thus our aim will be to compute CAEI allocations with improved welfare guarantees.

The *social welfare* of an allocation $\mathbf{x}$ is the sum of utilities of all the agents: $SW(\mathbf{x}) = \sum_{i=1}^{n} V_i(\mathbf{x}_i)$.

3 Multiple Divisible Goods

We now formalize the model with multiple divisible goods and single minded agents. There is a set $N = \{1, \ldots, n\}$ of agents and a set $M = \{1, \ldots, m\}$ of goods. W.l.o.g., each good comes in one unit that is infinitely divisible. Each agent i demands a bundle $D_i = \langle v_{i,1}, \ldots, v_{i,m}\rangle$, where $v_{i,j} \in [0, 1]$ denotes the fraction required by agent i from good j. The utility of i for a bundle $\mathbf{z} \in [0, 1]^m$ is $V_i(\mathbf{z}) = 1$ if $z_j \geq v_{i,j}$ for all $j \in M$, and $V_i(\mathbf{z}) = 0$ otherwise. Since every part of a good is equally valued by agents, it suffices to specify its per unit price, *i.e.*, $p_j \in \mathbb{R}_+$ for good j.

We first show that the CAEI solution is guaranteed to exist for multiple divisible goods with single minded valuations, and can moreover be computed in polynomial time. All the results in this section apply to the CEEI solution as well, which differs only by requiring that agents exhaust their unit of currency.

Theorem 1. *Given a fair division problem with multiple divisible goods and single minded agents, a CAEI solution is guaranteed to exist and can be computed in polynomial time.*

Proof. To prove the statement we leverage results on the existence and computation of market equilibria in Fisher markets with Leontief utilities (for background see, e.g., chapter 5 in [Nisan *et al.*, 2007]). Given fair division problem

$(N, M, \mathcal{D})$, we consider a Fisher market where the set of buyers is N and the set of items is M. Each buyer i has budget $B_i = 1$ and *Leontief utility* described by the vector $D_i = \langle v_{i,1}, \ldots, v_{i,m} \rangle$, i.e. for any bundle z, $u_i(\mathbf{z}) = \min_{j \in M : v_{i,j} > 0} \left\{ \frac{z_j}{v_{i,j}} \right\}$. Fisher markets with Leontief utilities always have exact market equilibria when each item is desired by at least one buyer and each buyer desires at least one item [Maxfield, 1997].

Let $(\mathbf{x}, \mathbf{p})$ be any such equilibrium, where $x_{i,j}$ is the fraction received by agent i from good j and p_j is the price of good j. Denote the Leontief utility of buyer i in the market by $u_i^{\mathcal{L}}(\mathbf{x}_i, \mathbf{p}) = \min_{j \in M : v_{i,j} > 0} \left\{ \frac{x_{i,j}}{v_{i,j}} \right\}$. We argue the market equilibrium $(\mathbf{x}, \mathbf{p})$ is a CAEI solution for the fair division problem with single minded agents. To this end, we must show that all the items are allocated, each agent spends no more than a unit, and gets in exchange an optimal bundle at those prices. Clearly the first two requirements are met since $(\mathbf{x}, \mathbf{p})$ is a market equilibrium in the Fisher market with identical budgets.

We additionally show that each agent gets an optimal bundle in the fair division problem. For every agent i, if the allocation $\mathbf{x}$ satisfies the property that $x_{i,j} \geq v_{i,j}$ for all $j \in M$, then $V_i(\mathbf{x}_i) = 1$ and i gets its demand at these prices. Otherwise, there is an item k with $v_{i,k} > 0$ but $x_{i,k} < v_{i,k}$; then agent i does not get its demand, thus $V_i(\mathbf{x}_i) = 0$. Since $(\mathbf{x}, \mathbf{p})$ is a market equilibrium with respect to the Leontief utilities given by $\mathbf{v}$, we have $u_i^{\mathcal{L}}(\mathbf{x}_i, \mathbf{p}) \leq \frac{x_{i,k}}{v_{i,k}} < 1$.

Assume by contradiction that in the fair division problem agent i could afford its demand set at prices $\mathbf{p}$, i.e. $\exists \mathbf{y} \in [0,1]^m$ with $y_j \geq v_{i,j}\ \forall j \in M$, and $\mathbf{p}(\mathbf{y}) \leq 1$. Then in the Fisher market with Leontief utilities, buyer i could also purchase bundle $\mathbf{y}$ and get: $u_i^{\mathcal{L}}(\mathbf{y}, \mathbf{p}) = \min_{j : v_{i,j} > 0} \left\{ \frac{y_j}{v_{i,j}} \right\} \geq \min_{j : v_{i,j} > 0} \left\{ \frac{v_{i,j}}{v_{i,j}} \right\} = 1 > u_i^{\mathcal{L}}(\mathbf{x}_i, \mathbf{p})$, which is a strict improvement over $\mathbf{x}_i$, contradicting that $(\mathbf{x}, \mathbf{p})$ is a market equilibrium in the Fisher market. Thus the assumption was false and $(\mathbf{x}, \mathbf{p})$ is a CAEI.

Finally, Codenotti and Varadarajan (2004) showed that a Fisher market equilibrium can be computed in polynomial time via a convex program formulation, which can be used to compute a CAEI solution for single minded agents. □

As the next example illustrates, the solution computed in Theorem 1 is not necessarily optimal.

Example 1. Consider 2 agents and 2 goods, where the demand of agent 1 is $D_1 = \langle 0.5, 0.4 \rangle$ and of agent 2 is $D_2 = \langle 0, 0.6 \rangle$. The CAEI solution from Theorem 1 prices good 1 at $p_1 = 0$ and good 2 at $p_2 = 2$, and the allocation splits good 2 equally between the two agents. This way agent 2 gets zero utility, while agent 1 gets utility 1 (and some extra good that is of no use). Instead, another possible CAEI solution is to set prices $p' = \langle 1/3, 5/3 \rangle$, where agent 2 can now afford a quantity of 0.6 from good 2, and agent 1 still gets its demand. Thus the CAEI solution computed via the Leontief market equilibrium is dominated by another CAEI solution with better welfare. □

We show that while maximizing social welfare is in general NP-hard, the problem can be solved in polynomial time when the number of *types* of agents is constant. This allows handling possibly large numbers of agents when their demands fit some standard templates for resource requests.

Algorithm 1: COMPUTE-MAX-CAEI

Data: Agents $N = \{1, \ldots, n\}$, items $M = \{1, \ldots, m\}$, valuations $\mathbf{v}$, types τ
Result: Social welfare maximizing CAEI

$(\text{OPT}, \mathbf{x}^*, \mathbf{p}^*) \leftarrow (-\infty, \text{NULL}, \text{NULL})$
foreach $S \subseteq \mathcal{T}$ **do**
  $(\mathbf{m}, \mathbf{p}) \leftarrow \text{SOLVE-OPTIMAL-SUBSET}(S, \mathcal{M})$
  $x_{i,j} = \frac{m_{i,j}}{p_j},\ \forall i \in N,\ \forall j \in M$
  if $(\text{SW}(\mathbf{x}) > \text{OPT})$ **then**
    $\text{OPT} \leftarrow \text{SW}(\mathbf{x})$
    $(\mathbf{x}^*, \mathbf{p}^*) \leftarrow (\mathbf{x}, \mathbf{p})$
  end
end
return $(\mathbf{x}^*, \mathbf{p}^*)$

Theorem 2. *Computing a CEEI solution that maximizes social welfare for single minded valuations with divisible goods is NP-hard.*

The reduction from the NP-complete problem SET PACKING is deferred to the full version.

Theorem 3. *Given a fair division problem with single minded agents and divisible goods, a welfare maximizing CAEI solution can be found in polynomial time for a constant number of agent* types.

Proof. We are given a set $N = \{1, \ldots, n\}$ of agents, M of items, valuations $\mathcal{D}$, and $\mathcal{T}$ a set of types of agents, such that for each $i \in N$, $\tau(i) \in \mathcal{T}$ represents the type of agent i, and agents with same type have identical valuation ($\tau(i) = \tau(i') \Rightarrow \mathbf{v_i} = \mathbf{v_{i'}}$). Algorithm 1 solves this problem and runs in polynomial time for fixed $|\mathcal{T}|$.

$$\begin{aligned}
\text{maximize} \quad & \epsilon \\
\text{subject to} \quad & \sum_{j=1}^{m} p_j \cdot v_{i,j} \leq 1, && \forall i \in N : \tau(i) \in S \\
& \sum_{j=1}^{m} p_j \cdot v_{i,j} \geq 1 + \epsilon, && \forall i \in N : \tau(i) \in \bar{S} \\
& m_{i,j} \geq p_j \cdot v_{i,j}, && \forall j \in M, \forall i : \tau(i) \in S \\
& \sum_{i=1}^{m} m_{i,j} = p_j, && \forall j \in M \\
& \sum_{j=1}^{m} m_{i,j} \leq 1, && \forall i \in N \\
& \mathbf{p} \geq 0; \quad \mathbf{m} \geq 0; \quad \epsilon \geq 0
\end{aligned}$$

Procedure SOLVE-OPTIMAL-SUBSET is given as a linear program, where p_j is the price of good j, $m_{i,j}$ is the amount of money spent by agent i on good j, and ϵ is a variable that should be strictly positive in the optimal solution if the set S of agent types can be made simultaneously happy.

The observation underlying the algorithm is that agents of the same type must have identical utilities (but not necessarily identical bundles). The COMPUTE-MAX-CAEI procedure tries to compute a CAEI solution for every set S of types so that agents of type $\tau \in S$ are satisfied, while agents outside S cannot afford their demand. The LP constraints ensure these conditions are met and a positive solution is found if and only if there exists a CAEI solution for the set S. At the end, COMPUTE-MAX-CAEI selects the maximum over all S. □

4 Cake Cutting

Next we investigate the cake cutting problem with single minded agents. Cake cutting models the problem of allocating a heterogeneous divisible resource, such as land, time, mineral deposits or computer memory, among agents with different preferences. The "cake" is represented mathematically as the interval $[0, 1]$ and the agents have different preferences over the interval. A "piece" of cake A is a union of intervals: $A = (I_1, \ldots, I_m)$. The literature on cake cutting has seen very interesting algorithmic developments in recent years (see, e.g., [Kurokawa *et al.*, 2013; Segal-Halevi *et al.*, 2015c; Aziz and Mackenzie, 2016]), all of which are concerned with additive valuations in the one-dimensional model. Among the few exceptions we mention: [Caragiannis *et al.*, 2011] studied fair division with additive valuations constrained by a minimum length requirement (PUML), [Branzei *et al.*, 2013] studied externalities in cake cutting, while [Segal-Halevi *et al.*, 2015a; 2015b], studied cake cutting in two dimensions, where the agents also care about the shape of the pieces that they receive. In earlier work, a different multidimensional model of fair division has been explored under the name of pie-cutting by Brams et al. [Brams *et al.*, 2008], while a very general model of fair resource division has been explored in the works of DallAglio and Maccheroni and of Husseinov and Sagara , who prove the existence of fair and efficient allocations under mild conditions assuming that the valuations are continuous. However, single minded valuations are discontinuous, and so their results do not directly imply ours.

In this paper we focus on single minded agents, and so each agent i will demand a set D_i consisting of possibly several disjoint intervals: $D_i = (D_{i,1}, \ldots, D_{i,m_i})$. The utility of agent i for a piece of cake $A \subseteq [0, 1]$ is: $V_i(A) = 1$ if $D_i \subseteq A$ and $V_i(A) = 0$ otherwise.

Examples of single minded agents in cake cutting include cases where the agent requires land that has buildings on it (with some particular functionality) or wants to build a house in a particular location. The case where each single minded agent requires a contiguous piece (*i.e.,* $m_i = 1$) can in fact be mapped to a standard problem known as "interval scheduling" in operations research.

We first show that the CAEI solution always exists and can be computed efficiently in this model.

Theorem 4. *A CAEI solution is guaranteed to exist and can be computed in polynomial time in the cake cutting problem with single minded agents.*

Proof. For each $i \in N$, let $D_i = (D_{i,1}, \ldots, D_{i,m_i})$ be the disjoint intervals in its demand set, where $D_{i,j} = [l_{i,j}, r_{i,j}] \subseteq [0, 1]$. Divide each interval $D_{i,j}$ into two equal pieces, and let $F_{i,j} = \left\{l_{i,j}, \frac{l_{i,j}+r_{i,j}}{2}, r_{i,j}\right\}$ be the set of points resulting from this partition of $D_{i,j}$. Consider the set of all the distinct points thus obtained: $\mathcal{P} = (\bigcup_{i=1}^{n} \bigcup_{j=1}^{m_i} F_{i,j}) \cup \{0, 1\} = \{q_0, q_1, \ldots, q_m\}$, where $0 = q_0 < q_1 < \ldots < q_m = 1$. Then we can view every segment $I_k = [q_k, q_{k+1}]$ delimitated by two consecutive points in $\mathcal{P}$ as an indivisible good; moreover, good I_k is in the demand set of each agent i for which $[q_k, q_{k+1}] \subset D_i$. This gives a related allocation problem with m indivisible goods and n agents, where no agent has a demand set consisting of a single good. The latter property follows from the fact that we partitioned each interval $D_{i,j}$ into two pieces, which are viewed as distinct goods. From [Branzei *et al.*, 2015], a competitive equilibrium from equal incomes is guaranteed to exist and can be computed in polynomial time for every fair division problem with indivisible goods and single minded agents, where no agent has a demand set with only one good; let $(\mathbf{x}, \mathbf{p})$ be such an allocation and prices.

Then we can compute a CAEI in cake cutting by allocating the intervals in $\mathcal{P}$ in the same order (from left to right) as the indivisible goods are allocated under $\mathbf{x}$ and setting the price curve of each interval $[q_k, q_{k+1}]$ uniformly such that it sums up to the price of the corresponding indivisible good. □

When the demands are contiguous we can compute a welfare maximizing CAEI in polynomial time. The proof requires a more delicate construction, the details of which are in the full version of the paper (Branzei, Lv, Mehta [2016]).

Theorem 5. *A welfare-maximizing CAEI can be computed in polynomial time in cake cutting with single minded agents and contiguous demands. Moreover, when there are no identical agents, the welfare maximizing CAEI coincides with the solution to the pure optimization problem of maximizing welfare.*

Proof. The idea is to leverage a connection with the interval scheduling problem [Kleinberg and Tardos, 2005] to first decide the agents that get allocated, and then to construct an *asymmetric* pricing scheme to implement a price equilibrium.

The interval scheduling problem is as follows. There are n jobs to be run on a supercomputer, where each job i runs from time s_i to time f_i. There are multiple such requests arriving simultaneously, and the goal is to process as many of them as possible, but the computer can only run one job at a time. The question is how to schedule the jobs so that the maximum number of requests is served. The optimal allocation—that maximizes the number of jobs scheduled—is given by a greedy algorithm: Schedule first the job with the earliest finishing time, then remove the jobs intersecting with it, and repeat among the remaining jobs.

This problem can be mapped to cake cutting with single minded agents by considering an agent for each job, i.e. $N = \{1, \ldots, n\}$, and setting the demand of each agent i to the schedule of job i, that is $D_i = \{[s_i, f_i]\}$, where s_i, f_i are normalized in $[0, 1]$.

While our aim is to indeed maximize the number of completed jobs, we want to achieve this in a way that is fair. That is, we will compute an allocation of the cake $\mathbf{x}$ and a price curve $\mathbf{p}$ such that $(\mathbf{x}, \mathbf{p})$ represent a CAEI solution, where $p(x)$ is the value of the price density function at point x. This

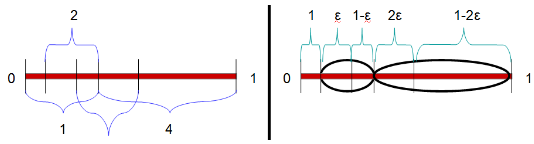

Figure 1: Four agents with contiguous single minded demands, indicated by the blue brackets the intervals (left). On the right, the two agents selected in the optimal schedule are circled and the prices marked on top of each differently priced interval (the price listed indicates the price of the whole interval for the corresponding bracket).

can be accomplished by allocating the agents in the same name order that the greedy algorithm allocates the corresponding jobs, with the caveat that if multiple agents have the same finishing time, then the one with the *latest* starting time must be selected. For each k-th agent that gets selected to receive their full demand, we select two very small intervals at the left and right endpoints of their demand, respectively, and price them (uniformly) to sum up to $k \cdot \epsilon$ and $1 - k \cdot \epsilon$, respectively. When there are no identical demands, the agents that don't get served but for which there is an unallocated piece of their demand, can receive a small such piece (at a total price of 1); otherwise, they get nothing (but their demand now has the property that it costs more than 1). This allocation can be done so that the whole cake is divided.

Then to handle identical agents, the algorithm can be modified such that whenever selecting the earliest finishing time and finding multiple agents with the same finish *and* start time, then take a very small interval at the leftmost point of the unallocated cake and divide it equally among the identified identical agents, setting the price uniformly to sum up to their budgets. Then iterate on the remaining cake. The complete proof is deferred to the full version, while an example with 4 agents is given in Figure 1. □

On the other hand, if the demands are not contiguous, then the problem of computing a welfare maximizing CAEI is NP-hard. The reduction from SET PACKING is in the full version.

Theorem 6. *Given a cake cutting problem with single minded valuations, computing a welfare-maximizing CAEI is NP-hard.*

We can also maximize social welfare for discontinuous demands in polynomial time when the number of agents is constant. The idea is to try out every possible set $S \subseteq N$ of *happy* agents, who will receive their required land at a price of one, while the remaining agents split the remaining cake at prices so they cannot afford their demand. Then select the outcome that maximizes welfare among all choices of S.

Theorem 7. *Given a cake cutting problem with single minded agents, a welfare maximizing CAEI can be computed in polynomial time when the number of agents is constant.*

5 Multiple Discrete Goods

Our setting for discrete goods is as follows. There is a set $N = \{1, \ldots, n\}$ of agents and $M = \{1, \ldots, n\}$ of items, where each item j comes in Q_j indivisible copies. Each agent i has a demand set $D_i \subseteq M$, and wants a copy of each item $j \in D_i$; D_i is not a multi-set. Without loss of generality, we assume that each type of item is required by some agent. We illustrate the model with an example: there are agents $N = \{1, 2\}$, items $M = \{1, 2, 3\}$, quantities $Q_1 = Q_3 = 1$, $Q_2 = 2$, and demands $D_1 = \{1, 2\}$, $D_2 = \{2, 3\}$.

The goal is to find a set of prices for each item—such that every copy from the same item has an identical price—together with an allocation $\mathbf{x}$ such that $(\mathbf{x}, \mathbf{p})$ is a CAEI solution. In our example, a CAEI solution is attained at: $p_1 = 0.5, p_2 = 0.5, p_3 = 0.5$ and $\mathbf{x}_1 = \{1, 2\}$, $\mathbf{x}_2 = \{2, 3\}$.

Next we provide a succinct characterization of the instances that admit a CAEI solution. To this end, a useful notion will be that of *over-demand*. We say that an item j is over-demanded among a set of agents S if the aggregate demand of the agents in S exceeds the available supply from item j, namely Q_j. We illustrate this in the next example.

Example 1. Consider a market $\mathcal{M}$ with single minded agents and discrete goods, where $M = \{1, 2\}$, with quantities $Q_1 = 2$ and $Q_2 = 4$, and agents $N = \{1, \ldots, 4\}$, with demands: $D_1 = D_2 = D_3 = \{1\}$, $D_4 = \{1, 2\}$. Then the aggregate demand of the set of agents $S = \{1, 2, 3\}$ with singleton demands, consists of 3 copies of item 1, which exceeds Q_1, the available supply from this item. Thus item 1 is over-demanded among S.

The next theorem shows that these are essentially the only instances that don't have a fair outcome. We say that demand of agent i is singleton if $|D_i| = 1$.

Theorem 8. *Given a fair division problem with single minded agents and discrete goods, a CAEI solution exists if and only if there is no set of agents with singleton demands that give rise to an over-demanded item among that set. Moreover, a solution can be computed in polynomial time if it exists.*

Proof. Consider Algorithm 2. Clearly, if there is a set of agents S with singleton demands, such that $D_i = \{j\}$ for some $j \in M$ and all $i \in S$, where $|S| > Q_j$, then no CAEI can exist. This is because $p_j \leq 1$ causes over-demand, and $p_j > 1$ makes it un-affordable and thereby un-sold.

We claim all other instances admit a CAEI. Algorithm 2 forms a sequence S of active agents, initially containing all the agents sorted in increasing order of $|D_i|$, and then iterates over the item types, at each step searching for an item that is desired by more agents among the active ones (i.e. in S) than there are copies available. If such an item j is found, then j is allocated at a price of $p_j = 1$, one copy to each of the first Q_j agents in S that want it. All the active agents that got item j are removed from S, and the algorithm searches for another such over-demanded *and* unallocated item among the updated set S. Otherwise, when no such item is found, all the goods left are in abundant supply for the agents in S and can be given for a small price (ϵ) among them. Note that these last agents, which are allocated bundles at a price of ϵ per good, receive not only items that they desire, but in fact all the remaining items.

We argue that the pricing rule ensures that a CAEI is computed. In particular, all the goods are allocated—at latest, the player (or set of players), which receive items in the last round of the algorithm, get all the remaining items. Moreover, the

Algorithm 2: COMPUTE-DISCRETE-CAEI

Input: Agents $N = \{1, \ldots, n\}$, items $M = \{1, \ldots, m\}$, demand sets D_i, quantities Q_j
Output: CAEI $(\mathbf{x}, \mathbf{p})$, or NULL if none exists
1 **for each** $(j \in M)$ **do**
2   **if** $(Q_j < card\left(\{i \in N \mid D_i = \{j\}\}\right))$ **then**
3   **return** NULL *// No solution*
4   **end**
5 **end**
6 $A \leftarrow \emptyset$ *// item types allocated so far*
7 $S \leftarrow$ SORT-BY-DEMAND-SET-SIZE(N) *// Sort the players in increasing order of demand set size*
8 **while** $(\exists j \in M \setminus A :$ OVER-DEMAND$(j, S) =$ TRUE$)$ **do**
9   $p_j \leftarrow 1$; ALLOCATE-OVERDEMAND(j, S); $A \leftarrow A \cup \{j\}$
10 **end**
11 $\epsilon \leftarrow 1/\left(1 + \sum_{j=1}^{m} Q_j\right)$ *// very small value*
12 **for each** $(j \in M \setminus A)$ **do**
13   $p_j \leftarrow \epsilon$; ALLOCATE-REMAINDER(j, S); $A \leftarrow A \cup \{j\}$
14 **end**
15 **Procedure** ALLOCATE-OVERDEMAND *(j, S)*
16   $Q \leftarrow Q_j$; $S_j \leftarrow \{k \in S \mid j \in D_k\}$ *// active agents that want item j*
17   *// Give one copy of j to every active agent that wants it until running out of copies*
18   **for each** $(i \in S_j)$ **do**
19   **if** $(Q > 0)$ **then**
20   append$(\mathbf{x}_i, j)$; $Q \leftarrow Q - 1$ *// Agent i gets a copy of j; decrease the # of copies left*
21       $S \leftarrow S \setminus \{i\}$ *// mark i as inactive*
22   **end**
23   **end**
24 **Procedure** ALLOCATE-REMAINDER *(j, S)*
25   $S_j \leftarrow \{k \in S \mid j \in D_k\}$ *// active agents that want item j*
26   **for each** $(i \in S_j)$ **do**
27   append$(\mathbf{x}_i, j)$
28   **end**
29   $K \leftarrow Q_j - |S_j|$ *// # of copies of j left unallocated*
30   **if** $(K > 0)$ **then**
31     $\ell \leftarrow S_j.last()$ *// The last player in S_j gets all the additional copies of j*
32   append$(\mathbf{x}_\ell, \{j\} \times K)$
33   **end**

pricing rule is such that each agent i is allocated all of its bundle at once, regardless of whether $\mathbf{x}_i$ contains i's required demand set or not. The agents are divided in two categories, namely those who get exactly one item at a price of 1 (Line 9) and those who get multiple items at a price of ϵ (Line 13). In both cases, the pricing rule ensures that the allocation $\mathbf{x}_i$ costs at most 1. By a case analysis, it can be seen that each agent gets an optimal bundle at the given prices; due to space constraints, this is delegated to the full version. □

We illustrate Algorithm 2 on an example.

Example 2. Let $N = \{1,2,3,4,5\}$, $M = \{1,2,3,4,5\}$, quantities $Q_1 = 2$, $Q_2 = 4$, $Q_3 = 2$, $Q_4 = 3$, $Q_5 = 2$, demands: $D_1 = \{1\}$, $D_2 = \{1,2\}$, $D_3 = \{1,3\}$, $D_4 = \{2,3,4\}$, $D_4 = \{2,3,4,5\}$. Only agent 1 has singleton demand $\{1\}$ and $|Q_1| \geq 1$, so there is a CAEI. Algorithm 2 sorts the agents in increasing order by demand set size and initializes $S = (1, 2, 3, 4, 5)$. The search for over-demanded items begins. Item 1 is wanted by 3 agents in S, namely $\{1, 2, 3\}$, but $Q_1 = 2 < 3$; it is over-demanded. Set $p_1 = 1$, $\mathbf{x}_1 = \mathbf{x}_2 = \{1\}$. Update $S \leftarrow S \setminus \{1, 2\} = \{3, 4, 5\}$. Next item 3 is over-demanded (wanted by agents $3, 4, 5$ but $Q_3 = 2$). Set $p_3 = 1$ and $\mathbf{x}_3 = \mathbf{x}_4 = \{3\}$. Update $S = \{5\}$. All the items are in sufficient quantities now, give all of them to agent 5 at price $\epsilon = 1/14$ each.

It is NP-hard to compute a welfare-maximizing CAEI (see [Branzei *et al.*, 2015]), but if we no longer insist on all the items being allocated, then we get a polynomial time algorithm by solving the problem as if the goods were divisible, and then rounding the solution.

Theorem 9. *Consider a fair division problem with agents N, m items in quantities $Q = (Q_1, \ldots, Q_m)$, and demand $D_i \subseteq \{1, \ldots, m\}$ for agent $i \in N$. Let $\mathcal{T}$ be the types of the agents, where same type agents have the same demand. Then a welfare maximizing CAEI, where all items need not be sold, can be computed in polynomial time if $|\mathcal{T}|$ is a constant.*

6 Discussion

We studied the computation and complexity of the *competitive allocation from equal incomes*, a method for fair division, for single minded agents for multiple divisible and discrete goods, and cake-cutting. Although a solution can be computed efficiently, social welfare maximizing solutions are hard to compute in general. However, we solved the latter efficiently for several interesting special scenarios. However, cases with constantly many goods, or altogether characterizing easy instances remain unresolved. will be interesting to settle these. We note that our results also work valuation classes where the agents have a desired target set (e.g. the set of parts of a bike), and very small positive value (ϵ) for any other bundle of goods. For such scenarios, the solution computed is an ϵ-CAEI.

In the full version of the paper, we also explain how the CAEI solution can rule out (via the prices) the least desirable envy-free allocations, while remaining envy-free on exactly the same instances where an envy-free solution is guaranteed to exist. On the other hand, the standard CEEI solution is not the right notion for cake cutting or discrete goods; for instance, with discrete goods, even when there exist goods in abundant quantities for everyone, it can be the case that the number of units is "wrong" and cannot result in the exact clearing prescribed by CEEI.

An immediate generalization is that of Leontief valuations, where even existence of succinct characterization of instances that admit CEEI or CAEI in case of discrete goods is not clear, while the bigger question of understanding fair division with complementarities remains a mystery at large.

References

[Aziz and Mackenzie, 2016] H. Aziz and S. Mackenzie. A discrete and bounded envy-free cake cutting protocol for four agents. In *Proceedings of STOC*, 2016.

[Brams and Taylor, 1996] S. J. Brams and A. D. Taylor. *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press, 1996.

[Brams *et al.*, 2008] S. Brams, M. Jones, and C. Klamler. *International Journal of Game Theory*, 36(3):353–367, 2008.

[Branzei *et al.*, 2013] S. Branzei, A. Procaccia, and J. Zhang. Externalities in cake cutting. In *IJCAI*, 2013.

[Branzei *et al.*, 2015] S. Branzei, H. Hosseini, and P. B. Miltersen. Characterization and computation of equilibria for indivisible goods. In *SAGT*, pages 244–255, 2015.

[Branzei *et al.*, 2016] S. Branzei, Y. Lv, and R. Mehta. To give or not to give: Fair division for single minded valuations, 2016. http://arxiv.org/abs/1602.09088.

[Briest *et al.*, 2005] P. Briest, P. Krysta, and B. V ocking. Approximation techniques for utilitarian mechanism design. In *STOC*, 2005.

[Budish, 2011] E. Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 6(119):1061–1103, 2011.

[Caragiannis *et al.*, 2011] I. Caragiannis, John K. Lai, and A. D. Procaccia. Towards more expressive cake cutting. In *IJCAI*, pages 127–132, 2011.

[Chevaleyre *et al.*, 2006] Y. Chevaleyre, P. E. Dunne, U. Endriss, J. Lang, M. Lemaître, N. Maudet, J. Padget, S. Phelps, J. A. Rodríguez-Aguilar, and P. Sousa. Issues in multiagent resource allocation. *Informatica*, 30:3–31, 2006.

[Codenotti and Varadarajan, 2004] B. Codenotti and K. R. Varadarajan. Efficient Computation of Equilibrium Prices for Markets with Leontief Utilities. In *ICALP*, pages 371–382, 2004.

[DallAglio and Maccheroni,] M. DallAglio and F. Maccheroni. Disputed lands. *Games and Economic Behavior*, 66(1):57–77.

[Feige *et al.*, 2015] U. Feige, M. Feldman, N. Immorlica, R. Izsak, B. Lucier, and V. Syrgkanis. A unifying hierarchy of valuations with complements and substitutes. In *AAAI*, pages 872–878, 2015.

[Feldman and Lucier, 2014] Michal Feldman and Brendan Lucier. Clearing markets via bundles. In *SAGT*, pages 158–169, 2014.

[Foley, 1967] D. Foley. Resource allocation and the public sector. *Yale Econ. Essays*, 7:45–98, 1967.

[Ghodsi *et al.*, 2011] A. Ghodsi, M. Zaharia, B. Hindman, A. Konwinski, S. Shenker, and I. Stoica. Dominant resource fairness: Fair allocation of multiple resource types. In *NSDI*, pages 323–336, 2011.

[Husseinov and Sagara,] F. Husseinov and N. Sagara. Existence of efficient envy-free allocations of a heterogeneous divisible commodity with nonadditive utilities. *Social Choice and Welfare*, 41(4), pages = 923-940).

[Kleinberg and Tardos, 2005] J. Kleinberg and E. Tardos. *Algorithm Design*. Pearson, 2005.

[Kurokawa *et al.*, 2013] D. Kurokawa, J. K. Lai, and A D. Procaccia. How to cut a cake before the party ends. In *AAAI*, 2013.

[Ledyard, 2007] J. O. Ledyard. Optimal combinatoric auctions with single-minded bidders. In *EC*, pages 237–242, 2007.

[Maxfield, 1997] R. Maxfield. General equilibrium and the theory of directed graphs. *Journal of Mathematical Economics*, 27(1):23–51, 1997.

[Moulin, 2004] H. Moulin. *Fair Division and Collective Welfare*. MIT Press, 2004.

[Nisan *et al.*, 2007] N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani (Eds.). *Algorithmic Game Theory*. Cambridge University Press, 2007.

[Othman *et al.*, 2014] A. Othman, C. H. Papadimitriou, and A. Rubinstein. The complexity of fairness through equilibrium. In *ACM-EC*, pages 209–226, 2014.

[Procaccia, 2013] A. D. Procaccia. Cake cutting: Not just child's play. *CACM*, 56(7):78–87, 2013.

[Robertson and Webb, 1998] J. M. Robertson and W. A. Webb. *Cake Cutting Algorithms: Be Fair If You Can*. A. K. Peters, 1998.

[Robu *et al.*, 2013] V. Robu, D. Parkes, T. Ito, and N. R. Jennings. Efficient interdependent value combinatorial auctions with single minded bidders. In *IJCAI*, 2013.

[Segal-Halevi *et al.*, 2015a] E. Segal-Halevi, A. Hassidim, and Y. Aumann. Envy-free cake-cutting in two dimensions. In *AAAI*, pages 1021–1028, 2015.

[Segal-Halevi *et al.*, 2015b] E. Segal-Halevi, S. Nitzan, A. Hassidim, and Y. Aumann. Fair and square: Cake-cutting in two dimensions. *CoRR*, abs/1510.03170, 2015.

[Segal-Halevi *et al.*, 2015c] Erel Segal-Halevi, Avinatan Hassidim, and Yonatan Aumann. Waste makes haste: Bounded time protocols for envy-free cake cutting with free disposal. In *Proceedings of AAMAS*, pages 901–908, 2015.

[Steinhaus, 1951] H. Steinhaus. *Mathematical Snapshots*. Oxford, 1951. second edition.

[Varian, 1974] H. Varian. Equity, envy and efficiency. *Journal of Econ. Theory*, 9:63–91, 1974.