

House Allocation With Fractional Endowments ^{*}

Stergios Athanassoglou [†]

Jay Sethuraman[‡]

July 2009; revised January 2010

Abstract

We study a generalization of the well known house allocation problem in which agents may own fractions of different houses summing to an arbitrary quantity, but have use for only the equivalent of one unit of a house. We also depart from the classical model by assuming that arbitrary quantities of each house may be available to the market. Justified envy considerations arise when two agents have the same initial endowment, or when an agent is in some sense disproportionately rewarded in comparison to her peers. For this general model, we design an algorithm to find a fractional allocation of houses to agents that satisfies ordinal efficiency, individual rationality, and no justified envy. Our results extend to the full preference domain. We show that individual rationality, ordinal efficiency, and no justified envy conflict with weak strategyproofness. We also show that individual rationality, ordinal efficiency and strategyproofness are incompatible. Finally, we prove that two reasonable notions of envy-freeness, no justified envy and equal-endowment no envy, conflict in the presence of ordinal efficiency and individual rationality. All our impossibility results hold in the strict preference domain.

Keywords: house allocation, fractional endowments, fairness, individual rationality

JEL Classifications: C72, C78, D71

^{*}We are grateful to two anonymous referees for the detailed comments on an earlier version of this paper.

[†]The Earth Institute at Columbia University, New York, NY; email: sa2164@columbia.edu

[‡]IEOR Department, Columbia University, New York, NY; email: jay@ieor.columbia.edu

1 Introduction

In this paper, we consider the problem of allocating a number of objects (say houses) to agents in an *efficient* and *fair* manner. Agents have complete and transitive preferences over the houses, and each agent wishes to be allocated the equivalent of *at most* one house. The distinguishing feature of our model is that agents may be endowed with *fractional* amounts of various houses. This model is a common generalization of several well-studied models that have received a lot of attention in the literature. If each agent is endowed with a distinct house, we recover the *housing market* model, first considered by Shapley and Scarf [9]. For this model, Shapley and Scarf proposed the Top-Trading Cycles mechanism (attributed to Gale) that finds the unique core allocation of the associated cooperative game. The TTC mechanism is Pareto efficient, strategyproof, and is individually rational (precise definitions are given later). Ma [7] later proved that the TTC mechanism is characterized by these properties. At the other extreme, if agents have no endowments, we recover the *random assignment* problem considered by, among others, Abdulkadiroglu and Sonmez [1], and Bogomolnaia and Moulin [4]. Abdulkadiroglu and Sonmez [1] study the *random priority* (RP) mechanism: agents are ordered randomly, they choose houses in this order, and each agent picks her most preferred house among the set of houses still available. This mechanism is *strategyproof*, *ex-post* Pareto efficient, and satisfies *equal treatment of equals*. An alternative mechanism—*probabilistic serial* (PS)—for the same problem was proposed by Bogomolnaia and Moulin [4]: at each point in time, agents consume their best available houses at unit rate. The resulting assignment can be implemented as a lottery over efficient deterministic assignments. Bogomolnaia and Moulin [4] showed that the PS mechanism finds an allocation that is *envy-free* and is *ordinally efficient* (a stronger form of efficiency), but satisfies strategyproofness only in a weaker sense. Katta and Sethuraman [5] extended the PS mechanism to the full preference domain, and proved that envy-freeness and ordinal efficiency are incompatible with even the weaker form of strategyproofness. Finally, Abdulkadiroglu and Sonmez [2] and Yilmaz [10] have considered house allocation problems with *existing tenants*:¹ in this model, some agents have no endowments (“new tenants”) and others are endowed with a distinct house (“existing tenants”). In such models, in addition to fairness and efficiency, it is natural to require *individual rationality*: absent such a requirement, agents may not participate in the mechanism in environments where such participation is voluntary. Abdulkadiroglu and Sonmez [2] designed a natural mechanism for this problem that specializes to the TTC mechanism when there are no new tenants, and to the

¹A later version of Yilmaz’s work includes a brief discussion of how his algorithm can be adapted to the fractional endowment setting considered in this paper.

RP mechanism when there are no existing tenants. Recently, Yilmaz [10] proposed a mechanism that specializes to the PS mechanism when there are no existing tenants, but is different from the TTC mechanism when there are no new tenants.² Note that in models with endowments expecting envy-free assignments that are also individually rational may be unreasonable, as these two requirements may be in obvious conflict with each other. A key contribution of Yilmaz [10] is his definition of *justified* and *unjustified* envy, which shows how to interpret the equity requirement when agents have different endowments.

Contributions. In our model, agents are allowed to have *arbitrary* endowments over the houses. Thus, the model we propose—House allocation with *fractional endowments*—is a common generalization of most of the existing models in the context of house allocation. For the house allocation problem with fractional endowments, we design an algorithm to find an assignment that is individually rational, ordinally efficient, and satisfies the *no justified envy* criterion of Yilmaz [10]. This algorithm generalizes the work of Yilmaz [10] who designed an algorithm with these properties for the special case of 0-1 endowments. Similar to Yilmaz’s algorithm, our algorithm is computationally efficient (its running-time is polynomial in the size of the input), and works by solving a sequence of maximum flow problems. These algorithms are in the spirit of earlier work of Katta and Sethuraman [5], and can be viewed as a generalization of their result to this substantially more general model. We further show that individual rationality, ordinal efficiency, and no justified envy are incompatible with weak-strategyproofness, a very mild incentive compatibility requirement. Somewhat surprisingly, we also show that ordinal efficiency and individual rationality alone are incompatible with strategyproofness. This negative result holds even in the canonical instance of the model in which there are n agents and n houses and each agent owns $1/n$ of every house. In the context of random assignment, this finding suggests that endowing each of n agents with $1/n$ of every house and allowing them to trade their “endowments” cannot lead to a truthful allocation procedure. Property rights are, in this sense, not helpful in overcoming related impossibility results in the random assignment problem (Bogomolnaia and Moulin [4]). Finally, we prove that, under ordinal efficiency and individual rationality, no justified envy conflicts with the fairness requirement that no two agents with equal endowments envy each other. All of our impossibility results hold even in the strict preference domain, and apply even when we substitute no justified envy with weaker equity criteria.

²Sethuraman [8] proposed a solution for this problem that specializes to the PS mechanism when there are no existing tenants and to the TTC mechanism when there are no new tenants.

Applications. House allocation and random assignment models have significant applications—the allocation of scarce on-campus housing is one celebrated example. Our model addresses any situation in which agents collectively own a bundle of goods, allowing for separate individual ownership. Owning a fraction of a good could reflect the probability of an agent actually possessing the good in question, provided that endowments sum to less than 1. In the context of dorm room allocation it could also reflect the relative “right” an agent has over a certain room.

Another interpretation of fractional individual ownership rests on thinking of goods as divisible entities. (This is in contrast to the classical models where goods are assumed to be indivisible and a fractional allocation of goods to agents is viewed as a lottery assignment.) This approach makes particular sense in markets where *time-sharing* is an option. If a good may be consumed at different times during the course of a given time period, then owning a fraction of it simply reflects the amount of time an agent is entitled to consuming it.

Finally, one can view our model as a way of improving upon a given lottery. Imagine a situation in which the final assignment of objects to agents will be made based on a given fractional assignment matrix. Interpreting this fractional assignment matrix as the endowment of the agents, the mechanism proposed here computes an alternative assignment matrix in which each agent’s random allocation stochastically dominates her endowment, yielding a “superior” lottery for each agent. Ordinal efficiency of the proposed mechanism implies that this new lottery cannot be improved upon for all the agents simultaneously.

Organization of the Paper. We discuss the model more formally in §2. Section 3 contains a description of the algorithm to find a solution for any given instance of the house allocation problem with fractional endowments; it also proves that the algorithm finds an assignment that is individually rational, ordinally efficient, and has no justified envy. Section 4 contains a collection of impossibility results. We discuss several extensions in §5 and end with a brief discussion of future research in §6.

2 Model Description

2.1 Model

Consider a market with n agents $I = \{1, 2, \dots, n\}$ and n houses $H = \{h_1, h_2, \dots, h_n\}$. Suppose agent i is endowed with e_{ij} units of house h_j , with each $e_{ij} \in [0, 1]$. We assume that each agent owns at most the equivalent of a full house, and that at most one unit of any house is owned by

the agents. In other words, the endowment matrix can be represented by a doubly sub-stochastic matrix.³

	h_1	h_2	\dots	h_n
1	e_{11}	e_{12}	\dots	e_{1n}
2	e_{21}	e_{22}	\dots	e_{2n}
\vdots	\vdots	\vdots	\vdots	\vdots
n	e_{n1}	e_{n2}	\dots	e_{nn}

with the rows indexed by the agents and the columns by the houses. Each agent i has (ordinal) preferences over the set of houses expressed by the complete and transitive relation \succeq_i .⁴ If houses h_j and h_k are such that $h_j \succeq_i h_k$ and $h_k \succeq_i h_j$ then agent i is indifferent between houses h_j and h_k , denoted by $h_j \sim_i h_k$. If $h_j \succeq_i h_k$, but $h_k \not\succeq_i h_j$, then agent i strictly prefers h_j to h_k , denoted by $h_j \succ_i h_k$. It is clear that the relation \sim_i is symmetric and transitive, and that the relation \succ_i is antisymmetric and transitive.

An *allocation* for agent i is a vector $p_i = (p_{i1}, p_{i2}, \dots, p_{in})$ such that $\sum_j p_{ij} \leq 1$. The interpretation is that in allocation p_i , agent i consumes p_{ij} units of house h_j . As is clear from the definition of an allocation, we consider environments in which agents desire at most the equivalent of “one house.” The allocations for all the agents can be described by an *assignment matrix*, with the rows indexing the agents and columns indexing the houses; like the endowment matrix, the assignment matrix will be a doubly sub-stochastic matrix. In this work, we explore mechanisms for allocating the houses to the agents satisfying some desirable properties, described next.

2.2 Mechanisms and Properties

A *mechanism* is a function that determines an assignment matrix for every possible profile of preferences and endowments. The desirable properties of a mechanism typically stem from *efficiency*, *truthfulness*, and *equity* considerations.

Efficiency. For agent i , an allocation p_i *dominates* q_i , denoted $p_i \succeq_i q_i$,

$$p_i \succeq_i q_i \Leftrightarrow \sum_{k \succeq_i h} p_{ik} \geq \sum_{k \succeq_i h} q_{ik}, \text{ for all } h \in H.$$

³All of the results extend to the more general case in which an arbitrary amount (instead of 1) of each house is available in the market. See §5 for this and other generalizations.

⁴All of our results extend in a straightforward manner to the case in which \succeq_i is quasi-transitive, and also to the case in which \succeq_i is a partial order. We omit the details.

If at least one of the above inequalities is strict, then p_i *strictly dominates* q_i , and is denoted $p_i \succ_i q_i$. (The dominance relation described here is simply the first-order stochastic dominance.) Note that \succeq is a partial order and certain allocations are not comparable: for example, getting the second best object for sure cannot be compared to getting 1/2 unit each of the best and worst objects. The dominance relation defined on individual allocations extends to assignment matrices in a natural way: an assignment matrix P *dominates* an assignment matrix Q if $p_i \succeq_i q_i$ for every agent i ; P *strictly dominates* Q if P dominates Q , and if $p_i \succ_i q_i$ for some agent i . An assignment matrix P is said to be *ordinally efficient* (or simply *efficient*) if P is *not* strictly dominated by any assignment matrix Q . A mechanism is efficient if it determines an efficient assignment matrix for every profile of preferences and endowments.

Individual Rationality. Consider an environment in which participation is voluntary. If the mechanism finds an allocation p_i for agent i such that $p_i \succeq_i e_i$, then agent i will *always* participate. Otherwise i may choose not to participate, which may result in an inefficient assignment. So, an assignment P is said to be *individually rational* if $p_i \succeq_i e_i$ for each agent i . A mechanism is individually rational if it determines an individually rational allocation for every agent.

Truthfulness. In many application contexts, preferences of the agents are not observable, but should be elicited from them. A natural, but fairly strong, requirement then is a mechanism in which it is a (weakly) dominant strategy for agents to reveal their preferences truthfully. As not every pair of allocations can be compared, there are two versions of this property. A mechanism is said to be *strategyproof* if for every agent i , the allocation she obtains by reporting her true preferences (weakly) dominates the allocation she obtains by reporting any other preference, regardless of what the other agents do. A mechanism is said to be *weakly strategyproof* if for every agent i , the allocation she obtains by reporting her true preferences is *not dominated* by the allocation she obtains by reporting any other preference, regardless of what the other agents do.

Equity. A minimal requirement of fairness is the familiar property of *equal treatment of equals* (ETE), which states that two agents with identical endowments and preferences should receive identical allocations. Formally a mechanism satisfies ETE if it finds an allocation such that $p_i = p_{i'}$ whenever $e_i = e_{i'}$ and $\succeq_i = \succeq_{i'}$, for any pair of agents i and i' .

A stronger requirement is *envy-freeness*, which states that each agent's allocation (weakly) dominates every other agent's allocation. That is, for any agent i , $p_i \succeq_i p_{i'}$ for every agent i' . Indeed this property has been considered in many economic contexts, notably in the house

allocation problem with no endowments. For the model with endowments, however, envy-freeness is too strong a requirement as it is in obvious conflict with individual rationality. For instance, suppose agents i and i' both have house h_j as their most preferred choice, but i owns h_j . In this case individual rationality dictates that i be allocated h_j fully, but i' will necessarily envy i in this allocation. Thus, any *reasonable* definition of envy in this context should take into account the fact that a mechanism may be forced to treat agents differently because they have different endowments. We define two notions of envy-freeness in this context (neither one of which implies the other), which are described next.

First, we can require envy-freeness only amongst agents who have identical endowments: this is a natural property and a reasonable requirement because two agents with identical endowments bring exactly the same resources to the group, so any differences in their final assignment should be explained solely by their preferences. We say that a mechanism satisfies *equal-endowment no envy* (EENE) if agents with the same initial endowments do not envy each other. In other words, a mechanism satisfies EENE if it finds an allocation P such that $p_i \succeq_i p_{i'}$ whenever $e_i = e_{i'}$.

An alternative notion—*no justified envy*—has been proposed by Yilmaz [10] for the house-allocation model with existing tenants, a special case of our model in which the endowment matrix is 0-1 and sub-stochastic. He distinguishes between two kinds of envy: *justified* and *unjustified*. The difference is explained by the following two examples, both due to Yilmaz.

Example 1. Consider the following instance of the house allocation problem with three agents $\{1, 2, 3\}$ and three houses $\{a, b, c\}$. Agent 1 prefers a to b and b to c ; agents 2 and 3 prefer b to a and a to c . The initial endowments are specified in braces, next to the preference ordering. Here, agent 1 is endowed with house b , agent 2 with a , and agent 3 with c .

$$\begin{aligned} 1 : & \quad a \succ b \succ c \quad \{b\} \\ 2 : & \quad b \succ a \succ c \quad \{a\} \\ 3 : & \quad b \succ a \succ c \quad \{c\} \end{aligned}$$

It is clear that the only individually rational and efficient assignment is

	a	b	c
1	1	0	0
2	0	1	0
3	0	0	1

Clearly agent 3 will envy both agents 1 and 2. However, this envy is *not justified* because it is not possible for agents 1 and 2 to give up any portion of their endowments to agent 3, receive a positive share of house c and still maintain individual rationality. In contrast, consider the following example:

Example 2.

$$\begin{aligned} 1 : \quad & a \succ c \succ b \quad \{b\} \\ 2 : \quad & b \succ c \succ a \quad \{a\} \\ 3 : \quad & b \succ a \succ c \quad \{c\} \end{aligned}$$

The assignment discussed earlier—giving a to 1, b to 2, and c to 3—is still individually rational and efficient. However there are other individually rational and efficient allocations because agents 1 and 2 are willing to give up some of b and a respectively for any house in the sets $\{a, c\}$ and $\{b, c\}$ respectively. In this context, if all of c is allocated to agent 3, then this agent could justifiably envy agents 1 and 2. This is because instead of giving agents 1 and 2 their best houses, the mechanism could have found a different allocation in which agents 1 and 2 do a little worse, still maintain individual rationality, and agent 3 does a little better. In particular, the assignment

	a	b	c
1	$\frac{1}{2}$	0	$\frac{1}{2}$
2	0	$\frac{1}{2}$	$\frac{1}{2}$
3	$\frac{1}{2}$	$\frac{1}{2}$	0

is individually rational, efficient, and is envy-free.

Yilmaz [10] formalizes these observations into the following definition: an agent i *justifiably envies* an agent i' if i 's allocation does not dominate i' 's, and if i 's allocation is an individually rational allocation for agent i' . Formally, i *justifiably envies* i' if

$$p_i \not\succeq_i p_{i'} \text{ and } p_i \succeq_{i'} e_{i'}.$$

Equivalently, i *does not justifiably envy* i' if

$$p_i \succeq_i p_{i'} \text{ or } p_i \not\succeq_{i'} e_{i'}.$$

We say that a mechanism satisfies *no justified envy* if in the assignment it determines, no agent justifiably envies any other agent.⁵ This definition of NJE is motivated by the following consideration: if $p_i \not\succeq_i p_{i'}$, then agent i could potentially prefer the allocation $p_{i'}$ to p_i ; however, if NJE

⁵We could require further that the allocation of agent i' dominate i 's endowment. But this makes it more difficult for justified envy to exist, so no justified envy becomes easier to satisfy.

is satisfied, this implies that p_i is not even an individually rational allocation for agent i' . In a way, this definition of NJE traces any potential envy in the final allocation to the difference in the endowments the agents start with.

We conclude this section by pointing out that the model we consider generalizes some of the most prominent models studied in the house allocation literature. In particular:

- If the endowment matrix is a permutation matrix, we recover the classical house trading model of Shapley and Scarf [9] in which each agent owns a distinct house.
- If the endowment matrix is identically zero, we get the random assignment problem considered by Abdulkadiroglu and Sonmez [1], Bogomolnaia and Moulin [4], Katta and Sethuraman [5], and others.
- If the endowment matrix is $\{0,1\}$ with each column sum at most 1 and each row sum at most 1, we obtain the house allocation problem with existing tenants, considered by Abdulkadiroglu and Sonmez [2] and Yilmaz [10].

3 The Controlled-Consuming (CC) Algorithm

In this section we design an efficient algorithm to find an allocation satisfying individual rationality, ordinal efficiency, and no justified envy. To make the discussion transparent and to keep the notation short, we shall restrict attention to the case in which the agents have strict preferences and have doubly stochastic endowment matrices. In Section 5, we show how our algorithm can be adapted to deal with indifferences and more general endowment profiles.

The CC algorithm falls under the general class of simultaneous eating algorithms, first introduced by Bogomolnaia and Moulin [4]. In particular, it allows each agent to “eat” her most preferred available house at rate 1, *as long as there is some way to complete the assignment so that the individual rationality constraints are not violated*; this continues until some house is completely consumed, or some individual rationality constraint is in danger of being violated. In the latter case the agents, whose continued consumption of their best available houses would violate some individual rationality constraint, are forbidden from consuming their most preferred houses even if they are available, and they move on to their next best house.

3.1 Flows and Cuts

A *network* (V, A) consists of a set V (called nodes) and a set A (called arcs) of ordered pairs of distinct elements of V , along with some additional data associated with V and A such as

capacities and costs. A useful model to keep in mind is that of a transportation network: the nodes represent demand or supply points of a commodity, and the arcs indicate potential ways in which the commodity can be transported, with the costs and capacities having the obvious interpretation. The relevant network problem for the purposes of this paper is the *maximum flow* problem, which is defined as follows: given a network, a *source* node s , a *sink* node t , and capacities $u(\cdot)$ on the arcs, find the maximum amount of flow that can be sent from s to t . Related to this problem is the problem of finding a *minimum* capacity s - t cut: an s - t cut is any collection of nodes S that includes s and excludes t , and the capacity of any such cut is the sum of the capacities of the arcs (i, j) with $i \in S$ and $j \notin S$. It is obvious that the capacity of any s - t cut is an upper bound on the total flow that can be sent from s to t ; therefore the *minimum* capacity s - t cut is an upper bound on the *maximum* s - t flow. A fundamental result in network flow theory is that the maximum s - t flow is exactly the same as the minimum capacity s - t cut (for background on maximum flows, see Ahuja et al. [3]). A more general model involves a network in which certain arc capacities are a function of a parameter λ ; the maximum flow (equivalently, the minimum cut) is therefore a function of λ , and the problem of interest is to understand the dependence of these quantities on this parameter. This is called the *parametric* maximum flow problem. The CC algorithm, as we shall describe later, works by finding a maximum flow in a suitably defined parametric network. We first illustrate the algorithm on an example.

3.2 An Illustrative Example

Consider the following instance:

$$\begin{array}{lll} 1 : & a \succ c \succ b & \{.99b, .01c\} \\ 2 : & b \succ a \succ c & \{.99a, .01c\} \\ 3 : & b \succ a \succ c & \{.01a, .01b, .98c\} \end{array}$$

As described earlier, the algorithm finds the final allocations by solving a sequence of maximum-flow problems on specific networks associated with the given instance. The networks all have the same set of nodes and arcs, but data associated with the network such as arc-capacities and other auxiliary information maintained by the algorithm will change over time. The nodes of the network are as follows: for each agent i , we introduce 3 nodes $i_{(1)}, i_{(2)}, i_{(3)}$, one for each “preference level”; there is a node for each house; and finally, there is a source node s and a sink node t . The arcs of the initial network are as follows: the source is connected to each node $i_{(k)}$, with the capacity of the arc being e_{ih} if h is the k th most preferred house of agent i ; each “house” node

is connected to the sink with an arc of capacity 1; and finally, there are k infinite capacity arcs from node $i_{(k)}$ to the house nodes, one to each of agent i 's k most preferred houses.




Figure 1: Initial network

The initial network is shown in Figure 3.

We first make a few observations:

- Any flow from s to t determines an assignment in a natural way: the amount of house h allocated to agent i , denoted p_{ih} , is given by the total amount of flow in the arcs $(i_{(k)}, h)$, for $k = 1, 2, 3$.
- The maximum flow from s to t is 3 (use the endowments as flows).
- Any flow of 3 units from s to t determines an individually rational assignment: the only way to send 3 units of flow from s to t is for each arc from s to $i_{(k)}$ to carry a flow equal to its capacity, which is equal to i 's endowment of her k th best house; the only way for this flow to reach the sink is via one of the arcs leaving $i_{(k)}$, and each of these arcs is to a house that i (weakly) prefers to her k th best house. So the individual rationality constraints will be satisfied for every agent i .

Iteration 1. The flow given by the endowments (i.e., the assignment in which each agent is given her endowment) is individually rational, but may not be ordinally efficient. To find an

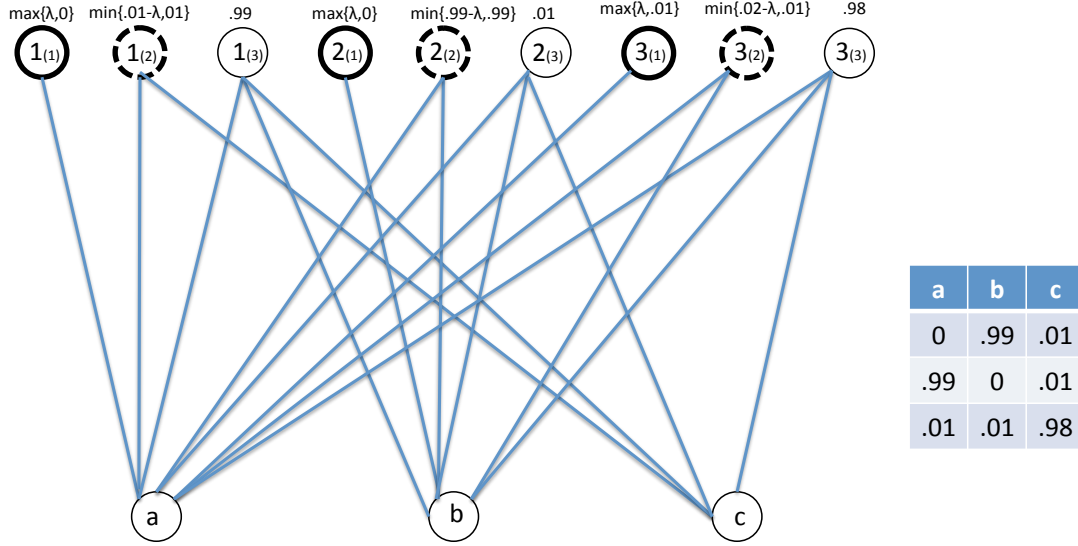


Figure 2: Initial network

ordinally efficient assignment, we employ a variation of the “simultaneous eating” algorithm due to Bogomolnaia and Moulin [4]: for each agent i , we consider increasing the capacity on the arc $(s, i_{(1)})$ at unit rate; because of this increase, we can decrease the capacity (at unit rate) on the *first positive capacity arc* in the sequence $(s, i_{(2)}), (s, i_{(3)}), \dots, (s, i_{(n)})$, and still maintain individual rationality. The algorithm we propose does exactly this with one important exception: if an agent i has a positive endowment e of her most preferred house, the capacity of the arc $(s, i_{(1)})$ is set to this endowment until time e so as to maintain individual rationality. This is seen, for instance, in Figure 4, which focuses on agent 3 and shows how the arc-capacities from the source to the nodes of agent 3 vary during the course of the algorithm. As agent 3 is endowed with 0.01 units of his most preferred house b , the capacity of arc $(s, 3_{(1)})$ will be 0.01 until λ

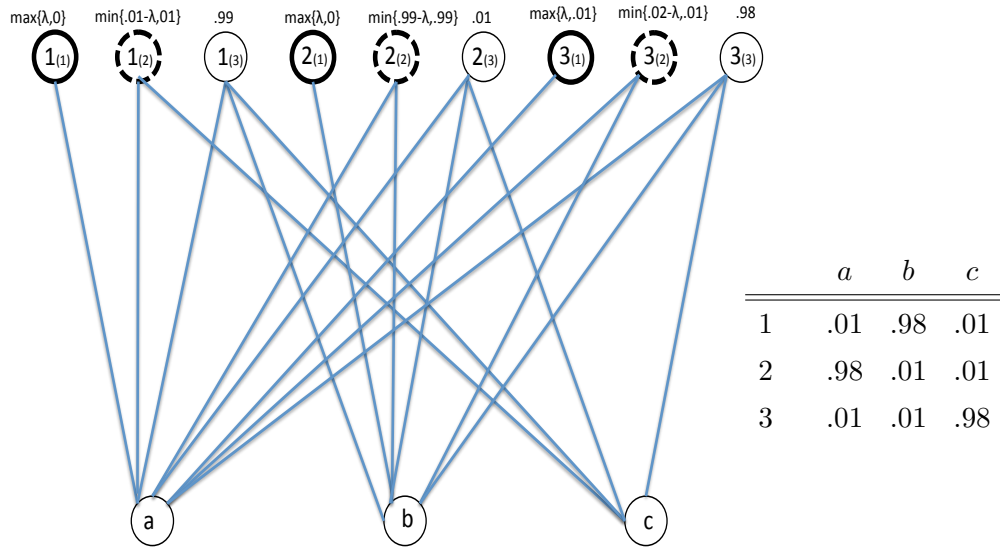


Figure 3: Initial network

reaches that value, after which it will simply be λ .

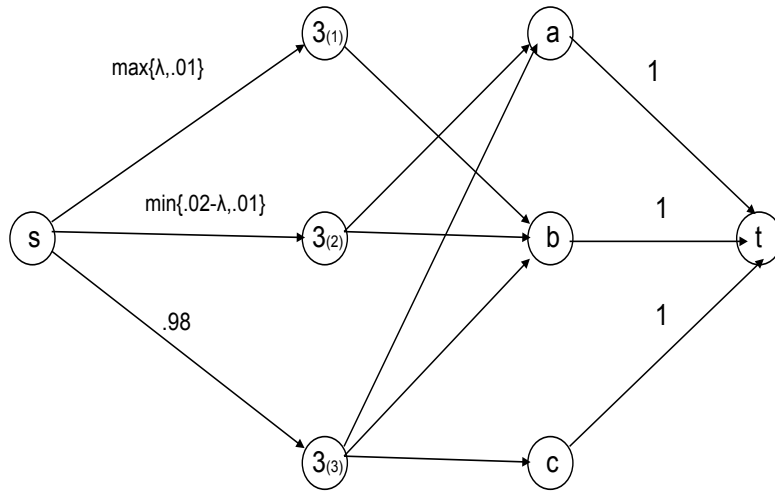


Figure 4: Changing capacities as λ increases.

We gradually increase λ and continue doing so as long as the maximum flow from s to t is still 3, and as long as none of the arc capacities drops below zero. The former is to ensure that we stay within the class of individually rational allocations, and the latter ensures that the trade-off of less- to more-preferred houses is feasible for each agent. For example, the capacity of arc $(s, 1_{(2)})$ drops to zero when $\lambda = 0.01$: the interpretation is that agent 1 consumes 0.01 of his most-preferred house (house a) by giving up the same amount of house c that he is endowed with; any additional consumption of a by this agent cannot be due to a demand of house c by the other agents, because agent 1 is not *entitled* to any part of house c any longer; so any further consumption of a by agent 1 must be due to his giving up an equal amount of house b . This is incorporated into the algorithm by setting the capacity of $(s, 1_{(2)})$ to zero and by setting the capacity of $(s, 1_{(3)})$ to $0.99 - \lambda$.⁶ The value of λ at the end of the first iteration is thus 0.01 and the allocation at this time is

	a	b	c
1	.01	.98	.01
2	.98	.01	.01
3	.01	.01	.98

Iteration 2. We continue increasing λ until it reaches .02, at which point the capacity of the arc $(s, 3_{(2)})$ becomes zero. At this point the network is updated so that the capacity of the arc $(s, 3_{(3)})$ is set to $0.98 - \lambda$. The value of λ at the end of this iteration is 0.02 and the allocation at this time is

	a	b	c
1	.02	.97	.01
2	.97	.02	.01
3	0	.02	.98

Iteration 3. We continue increasing λ until it reaches .5, at which point the maximum flow in the network is about to drop below 3. That is, the maximum s - t flow will be below 3 for any larger value of λ . This is so because the nodes $2_{(1)}$ and $3_{(1)}$ receive λ units of flow from the source and they can send this flow only to node b , which can send only one unit of flow to the sink. That house b is a bottleneck can also be seen by examining the min-cut: The relevant min-cut is $\{s, 2_{(1)}, 3_{(1)}, b\}$ with a capacity of $4 - 2\lambda$, which is below 3 for any $\lambda > 0.5$. Thus house b becomes

⁶The CC algorithm keeps track of these two entities—the best house for an agent and the next house for which he has a positive endowment—for each agent over time in two arrays called *best* and *next*.

unavailable to agents 2 and 3 ⁷, and their allocation of house b is set at .5. Their “best” house now becomes a , and their “next” houses are updated to c . The allocation at the time $\lambda = 0.5$ is given by

	a	b	c
1	.5	0	.5
2	.49	.5	.01
3	.01	.5	.49

Note that for agents 2 and 3, the capacity of the arcs from the source to their “best” houses will now be increased at unit rate. Notice, further, that at this stage all agents are “eating” house a .

Iteration 4. We continue increasing λ until it reaches .505, at which point the maximum flow in the network is about to drop below 3. The relevant minimum-cut in this case is $\{s, 1_{(1)}, 2_{(1)}, 2_{(2)}, 3_{(1)}, 3_{(2)}, a, b\}$ with a capacity of $3.505 - \lambda$. Thus, house a is declared unavailable to agents 1 and 3 and their allocation of house a is set at .505, and .005 respectively. In effect, the algorithm discovers that any additional allocation of house a to agents 1 or 3 will cause the individual rationality condition for agent 2 to be violated. Thus, even though house a is still available, it has to be made unavailable to agents 1 and 3, if the final allocation is to be individually rational for agent 2. Agent 1’s best house is updated to c and her next house to b . Agent 3’s best house is updated to c , while her next house is now undefined. The allocation at $\lambda = 0.505$ is given by

	a	b	c
1	.505	0	.495
2	.49	.5	.01
3	.005	.5	.495

Iteration 5. We continue increasing λ until it reaches .99, at which point the maximum flow in the network is about to drop below 3. The relevant minimum-cut is once again $\{s, 1_{(1)}, 2_{(1)}, 2_{(2)}, 3_{(1)}, 3_{(2)}, a, b\}$ with a capacity of $3.99 - \lambda$: thus, house a is declared unavailable to agent 2 and her allocation of house a is set at .49. Her best house is updated to c , while her next house is now

⁷The CC algorithm uses a set \mathcal{A} of ordered agent-house pairs to keep track of the set of houses available for each agent.

undefined, which we denote by U . The allocation when $\lambda = 0.99$ is thus

	a	b	c
1	.505	0	.495
2	.49	.5	.01
3	.005	.5	.495

Iteration 6. We continue increasing λ from 0.99. No other updates are made until the algorithm terminates at time $\lambda = 1$. Thus, the final allocation is given by

	a	b	c
1	.505	0	.495
2	.49	.5	.01
3	.005	.5	.495

The role of endowments and strategic behavior. We briefly comment on the CC algorithm's focus on individual endowments and vulnerability to strategic behavior. As the following example suggests, it can be profitable to claim that one's endowment is more valuable than it actually is. In particular, a valuable social endowment may give an agent added bargaining power over obtaining her most preferred houses.

Suppose that agent 1 misrepresents her preferences and submits $a \succ b \succ c$, instead of $a \succ c \succ b$. She thus overstates the value of her endowment to her. With this new preference profile, the CC algorithm will compute the following allocation

	a	b	c
1	.99	0	.01
2	.01	.98	.01
3	0	.02	.98

It is clear that agent 1's new allocation emphatically dominates her old one. By misstating that house b is her second most preferred, agent 1 can ensure that 2 trades her endowment of house a almost entirely with her. Notice that agent 2 is also quite happy with this trickery as her new allocation also clearly dominates her old one. The only agent who is left out in the cold is 3 who is in effect compelled to keep her endowment of house c .

The above example further highlights how agents with identical preferences (agents 2 and 3) may receive very different allocations depending on the attractiveness of their endowments for

other agents. In our example, agent 2's endowment is far more desirable to agent 1 than agent 3's endowment. This, however, does not give agent 2 an intrinsic edge over 3. The structure of the CC algorithm ensures that agent 3 is allowed to improve her allocation, so long as IR constraints are respected. In this regard, agent 1's IR constraints are key. When she declares her preferences truthfully, her coveted endowment of house b will be equally distributed to agents 2 and 3. When she lies, agent 2 ends up getting all but .01 of agent 1's endowment of b .

3.3 The Algorithm

We now formally present the algorithm. Recall that $I = \{1, 2, \dots, n\}$ denotes the set of agents and $H = \{h_1, h_2, \dots, h_n\}$, the set of houses. We assume strict preferences and a doubly stochastic endowment matrix. Let $h_{i_{(k)}}$ denote agent i 's k th most preferred house (thus $h_{i_{(1)}} \succeq_i h_{i_{(2)}} \succeq_i \dots \succeq_i h_{i_{(n)}}$). For convenience, denote $e_{i, h_{i_{(k)}}}$ by $e_{i_{(k)}}$. As mentioned earlier, the algorithm finds the final assignment by solving a sequence of (parametric) maximum-flow problems on specific networks associated with the given instance. The networks all have the same set of nodes and arcs, but some of the arc-capacities change during the course of the algorithm. The nodes of the network are as follows:

- **(agent nodes)** for each agent i , there are n nodes $i_{(1)}, i_{(2)}, \dots, i_{(n)}$, one for each “preference level”;
- **(house nodes)** for each house h_j , there is a node labelled h_j ; and
- a source node s and a sink node t .

The arcs of the network are as follows: the source is connected to each agent node $i_{(k)}$ with an arc whose capacity is denoted $u_{i_{(k)}}$; each house node h_j is connected to the sink with an arc of capacity 1; and finally, there are k infinite capacity arcs from node $i_{(k)}$ to the house nodes, one to each of agent i 's k most preferred houses. During the course of the algorithm the capacities of some of the arcs emanating from the source node will be varied; all other arc-capacities remain fixed. (To capture this, we sometimes use an additional superscript for $u_{i_{(k)}}$ to make this dependence explicit.) Finally, the algorithm maintains the following additional information that is critical to its operation: (i) a set \mathcal{A} of *available* agent-house pairs; and (ii) for each agent i , a *best* house index b_i and a *next* house index n_i .

Initial Network

If $(i, h) \in \mathcal{A}$, we say that house h is available for agent i . Initially, every house is available for each agent so that \mathcal{A} consists of all possible agent-house pairs. The initial capacities, $u_{i_{(k)}}^0$, of the arc connecting the source s to the node $i_{(k)}$ is set to the corresponding endowment $e_{i_{(k)}}$. The *best* house for agent i is her most preferred house among the houses available to her; if her best house is $h_{i_{(k)}}$, her *next* house, if any, is the smallest $j > k$ for which the arc $(s, i_{(j)})$ has positive capacity. Equivalently, agent i 's best house index, b_i , is k and her next house index, n_i , is j . (If there is no next house for agent i , n_i is undefined.) Note that the best house is defined with respect to the set \mathcal{A} whereas the next house is defined with respect to the arc-capacities of the associated network.

An iteration of the CC algorithm

The algorithm progresses by examining a sequence of networks at times $0 = \lambda^0 \leq \lambda^1 \leq \lambda^2 \leq \dots \leq \lambda^z = 1$. At each of these instants, the network is updated (some arc-capacities are changed), as is the additional information that it maintains (the set \mathcal{A} , and the best and next house indices for each agent). To complete the description of the algorithm we specify how, given all the data at time λ^t , the algorithm finds λ^{t+1} and updates the network as well as the auxiliary information.

Let λ be a parameter that will be gradually increased from its current value of λ^t . Consider the network at time λ^t and make the following changes for each agent i :

- Set the capacity of the arc $(s, i_{(b_i)})$ to

$$\max \left\{ \lambda - \sum_{\ell=1}^{b_i-1} u_{s, i_{(\ell)}}^t, \quad u_{s, i_{(b_i)}}^t \right\}. \quad (1)$$

- Set the capacity of the arc $(s, i_{(n_i)})$ to

$$\min \left\{ \sum_{\ell=1}^{b_i} u_{s, i_{(\ell)}}^t + u_{s, i_{(n_i)}}^t - \lambda, \quad u_{s, i_{(n_i)}}^t \right\}. \quad (2)$$

All other arc capacities are maintained at their values at time λ^t .

If the maximum in Expression (1) is achieved by the first term, we say that agent i is *consuming* her best house; if the maximum in Expression (1) is achieved by the second term, we say that agent i is *claiming* her best house. Clearly an agent can consume or claim a house only if it is available to her, and only when it is her best house. The interpretation of these two steps is very

straightforward: for each agent, we increase the capacity of the arc to the best available house at unit rate (the first term in Expression (1)), except when individual rationality requires a larger quantity of that house to be set aside for this agent (the second term in Expression (1)). In the former case, the increase is accompanied by a corresponding decrease in the guarantee of the next best house, which explains Expression (2).

We now solve a parametric maximum-flow problem on this updated network by gradually increasing λ from its current value of λ^t . Observe that for $\lambda = \lambda^t$, the maximum s - t flow is n . Define λ^{t+1} as the earliest time at which at least one of the following events occurs:

- (a) The capacity of some arc *becomes* zero;
- (b) Any further increase of λ will cause the maximum s - t flow to be strictly below n ;
- (c) The value of λ is 1.

We first obtain the new network by fixing the capacities of all the arcs to be their values at $\lambda = \lambda^{t+1}$.

Event (c) defines the termination condition for the algorithm: any maximum s - t flow (necessarily of value n) in the final network can be interpreted as an allocation of houses to the agents: the amount of house h allocated to agent i , denoted p_{ih} , is given by the total amount of flow in the arcs $(i_{(k)}, h)$, for $k = 1, 2, \dots, n$.

If Event (b) occurs, the algorithm identifies a minimum s - t cut whose capacity is strictly below n for any $\lambda > \lambda^{t+1}$. Such a cut will be of the form $s \cup X_{t+1} \cup Y_{t+1}$, where X_{t+1} is a subset of the agent nodes of the form $i_{(k)}$ and Y_{t+1} , a subset of the house nodes.⁸ Moreover, for any node $i_{(k)} \in X_{t+1}$, each of its neighboring house nodes must be in Y_{t+1} (otherwise the cut will have infinite capacity). Consider any agent i . Suppose her total consumption (including all her copies in X_{t+1}) of the bottleneck set Y_{t+1} increases with λ , then agent i 's best house, which is necessarily in Y_{t+1} , is declared *unavailable* to her. This occurs if and only if $i_{(b_i)} \in X_{t+1}$ and $i_{(n_i)} \notin X_{t+1}$; in this case the pair $(i, h_{i_{(b_i)}})$ is removed from \mathcal{A} , and her best house index is incremented (as her best house is no longer available to her). This is done for each agent i . The next house indices for all the agents are updated, and the algorithm continues.

If Event (a) occurs, the best house indices do not change, but the next house indices of at least one of the agents changes; we recalculate the next house indices of all the agents and the algorithm continues.

⁸When there are many such minimum cuts, we pick one with the maximum number of nodes on the source side. Such a cut is unique, see Lovasz and Plummer [6].

This completes one iteration of the algorithm. A formal description of the algorithm appears as Algorithm 1.

Relationship to Yilmaz’s algorithm. The CC algorithm generalizes the earlier algorithm of Yilmaz [10] that was designed for the special case of 0-1 endowments, which, in turn, generalized an earlier algorithm of Katta and Sethuraman [5] for the case of no endowments. Therefore the CC algorithm shares a number of features with these two algorithms. The one key difference is in the network construction: in the case of fractional endowments it may be necessary to make as many copies of each agent node as the number of distinct objects, whereas for the case of 0-1 endowments a single additional copy suffices. Given the need to work with n agent-nodes for each agent, it is critical that any additional consumption of the best object for each agent be compensated in a way that maximizes potential trading opportunities in the future: this is done by reducing the capacity of the arc from the source to that agent’s entitlement of the *next* best object; a different choice may not result in an efficient outcome! This is facilitated by the auxiliary information maintained by the algorithm. These issues do not arise in the case of 0-1 endowments or when there are no endowments: in these special cases, once an agent node becomes part of the bottleneck set, either an object is completely consumed, or the group of agents in the bottleneck set can be isolated and their final allocation can be determined by solving a subproblem in isolation. In other words, once the IR constraint becomes binding for a group of agents, that group of agents will compete for the rest of their endowments, and this determines the subproblem that the algorithm solves. In contrast, for the case of fractional endowments, there are potentially n different IR conditions for each agent, and it is not possible to view these in isolation. It is this feature that makes the CC algorithm somewhat more complicated to describe.

3.4 Properties

We show that the CC mechanism is individually rational, ordinally efficient, and satisfies no justified envy and no-envy for agents with equal endowments.

Proposition 1 *The CC mechanism is individually rational.*

Proof. Fix an agent i and any $k \in \{1, 2, \dots, n\}$. We show, by induction on t , that

$$\sum_{\ell=1}^k u_{i_{(\ell)}}^t \geq \sum_{\ell=1}^k e_{i_{(\ell)}} \quad (3)$$

Algorithm 1: The CC Algorithm

begin

$\mathcal{A}^0 = \{(i, h) \mid i \in I, h \in H\}$

for $i \in I$ **do**

$u_{i(k)}^0 = e_{i(k)}$, for $k = 1, 2, \dots, n$
 $b_i = 1$, $n_i = \min_{\ell > 1} \{\ell \mid u_{i(\ell)}^0 > 0\}$

$t = 0$, $\lambda^0 = 0$

while $\lambda^t < 1$ **do**

for $i \in I$ **do**

$u_{i(b_i)}(\lambda) = \max\{\lambda - \sum_{\ell=1}^{b_i-1} u_{i(\ell)}^t, u_{i_{b_i}}^t\}$
 $u_{i(n_i)}(\lambda) = \min\{\sum_{\ell=1}^{b_i} u_{i(\ell)}^t + u_{i(n_i)}^t - \lambda, u_{i_{n_i}}^t\}$

 Gradually increase λ from λ^t until (a) some arc capacity *becomes* zero; or (b) the maximum s - t flow in the network is below n for any larger value of λ ; or (c) $\lambda = 1$

 Set $\lambda^{t+1} = \lambda$, $\mathcal{A}^{t+1} = \mathcal{A}^t$

for $i \in I$ **do**

$u_{i(b_i)}^{t+1} = \max\{\lambda^{t+1} - \sum_{\ell=1}^{b_i-1} u_{i(\ell)}^t, u_{i_{b_i}}^t\}$
 $u_{i(n_i)}^{t+1} = \min\{\sum_{\ell=1}^{b_i} u_{i(\ell)}^t + u_{i(n_i)}^t - \lambda^{t+1}, u_{i_{n_i}}^t\}$

if *maximum s - t flow is below n for $\lambda > \lambda^{t+1}$* **then**

 Find a min-cut, which will be of the form $s \cup X_{t+1} \cup Y_{t+1}$, where X_{t+1} is a subset of the agent nodes and Y_{t+1} , a subset of the house nodes

for $i \in I$ **do**

if $i_{(b_i)} \in X_{t+1}$ *and* $i_{(n_i)} \notin X_{t+1}$ **then**
 $\mathcal{A}^{t+1} \leftarrow \mathcal{A}^{t+1} \setminus \{(i, h_{i(b_i)})\}$
 $b_i \leftarrow b_i + 1$

for $i \in I$ **do**

$n_i = \min_{\ell > b_i} \{\ell \mid u_{i(\ell)}^{t+1} > 0\}$

$t \leftarrow t + 1$

end

for any $t \geq 0$. As $u_{i_{(k)}}^0 = e_{i_{(k)}}$ for each k , the result is true by definition for $t = 0$. Suppose the result is true at the beginning of iterations $0, 1, 2, \dots, t$. We show that the result is true at the beginning of iteration $t + 1$, equivalently, at the end of iteration t . In iteration t , the only arc-capacity that potentially decreases is that of the arc from s to $i_{(n_i)}$; but from Expressions (1) and (2), this decrease, if any, is offset by a corresponding increase in the capacity of the arc from s to $i_{(b_i)}$. The result follows. \blacksquare

We now turn to properties concerning ordinal efficiency and no-envy. To this end we start with a few observations. First, the total capacity of the arcs emanating from the source node s is *always* n , the total number of agents (equivalently, houses). Second, the maximum s - t flow in the network is always n , which implies that the flow along the arcs emanating from the source is unique. These observations are useful in proving the following:

Lemma 1 *Suppose the minimum-cut found by the algorithm in iteration t is $s \cup X_{t+1} \cup Y_{t+1}$. Then:*

- (a) *The set Y_{t+1} is precisely the set of all houses that are adjacent to some agent in X_{t+1} .*
- (b) *If agent-node $i_{(k)} \in X_{t+1}$, then $i_{(\ell)} \in X_{t+1}$ for any $\ell < k$.*
- (c) *Let \bar{X}_{t+1} denotes the set of agents nodes that are not in X_{t+1} . In any maximum-flow from s to t (in the current network or in the future), the flow carried by any arc from a node in \bar{X}_{t+1} to a node in Y_{t+1} is zero.*

Proof. Observe that arcs from X_{t+1} to Y_{t+1} have infinite capacity, and a minimum-cut cannot contain any such arc; so it is clear that every house node adjacent to *some* agent-node in X_{t+1} must be included in Y_{t+1} . Any house node h for which there is no incoming arc from any node in X_{t+1} will contribute one unit to the cut-capacity (as the arc from h to t will be part of the cut); removing such a node from the cut will decrease the cut-capacity by 1 as no additional arcs will contribute to the cut-capacity. This verifies part (a) of the proposition. To verify part (b), observe that agent-node $i_{(k)}$ connects to all the house nodes that any other agent-node $i_{(\ell)}$ (with $\ell < k$) connects to; including such an agent-node $i_{(\ell)}$ in the cut *may* decrease the cut-capacity (as the arc from s to $i_{(\ell)}$ will no longer contribute), but will not increase it. To verify part (c), observe that every arc from a node in \bar{X}_{t+1} to a node in Y_{t+1} is a *backward* arc in the minimum-cut, and so cannot carry positive flow in any maximum-flow in the network at time λ^{t+1} . Moreover, by the definition of λ^{t+1} , we know that in the network at time λ^{t+1} , the total capacity of the arcs connecting s to X_{t+1} equals the total capacity of the arcs connecting Y_{t+1} to t . In all future

networks, the total capacity of the arcs connecting s to X_{t+1} cannot decrease, which implies that the total *flow* along these arcs cannot decrease either. But the only way for this flow to reach the sink node t is via the arcs from Y_{t+1} to t . ■

Armed with these observations, we are now ready to formally prove that the CC mechanism is ordinally efficient and satisfies no justified envy as well as no-envy for agents with equal endowments.

We begin with ordinal efficiency. Let P be the assignment found by the CC algorithm. The CC algorithm is a “simultaneous eating” algorithm (we can find eating speed functions such that the assignment found by the simultaneous eating algorithm with these eating speed functions is the assignment P); Bogomolnaia and Moulin [4] showed that any assignment found by such an algorithm is ordinally efficient and every ordinally efficient assignment can be found this way. It follows then that P is ordinally efficient. Nevertheless, we present a direct proof of this result. The proof uses an alternative characterization of ordinal efficiency, due to Bogomolnaia and Moulin [4], and extended to the full preference domain by Katta and Sethuraman [5]. Given an assignment matrix P and preference relations \succeq_i for each agent i , define the binary relation $\tau(P, \succeq)$ over the set of houses H as follows:⁹

$$h\tau h' \Leftrightarrow \{ \exists i \in I : h \succeq_i h' \text{ and } p_{i,h'} > 0 \}.$$

Say that the relation is *strict* if $h \succ h'$ in the definition above. The relation τ is *cyclic* if there exists a cycle of relations $h_1\tau h_2, h_2\tau h_3, \dots, h_{k-1}\tau h_k, h_k\tau h_1$. It is *strictly cyclic* if it is cyclic, and at least one of the relations in the cycle is strict. The following result is due to Bogomolnaia and Moulin [4].

Proposition 2 *Let P be a random assignment matrix for the preference profile \succeq . Then P is ordinally efficient if and only if the relation $\tau(P, \succeq)$ is not strictly cyclic.*

Proposition 3 *The CC mechanism is ordinally efficient.*

Proof. Suppose not. Consider an instance for which the assignment P found by the CC algorithm is not ordinally efficient. Then there is a set of agents and a set of houses such that the relation τ is strictly cyclic. Suppose without loss of generality that the set of agents is $\{1, 2, \dots, k\}$, the set of houses $\{h_1, h_2, \dots, h_k\}$ and suppose that $h_i \succ_i h_{i+1}$ for each agent i (interpreting h_{k+1} as h_1),

⁹Note that τ depends on both the given assignment and the preference relation, but we suppress this dependence because it is usually clear from the context.

and $p_{i,h_{i+1}} > 0$. As agent i prefers h_i to h_{i+1} , and as $p_{i,h_{i+1}} > 0$, house h_i becomes unavailable for agent i before house h_{i+1} does. Let λ^i be the time at which house h_i becomes unavailable for agent i , and suppose $\lambda^1 = \min_\ell \{\lambda^\ell\}$. We claim that the minimum s - t cut $s \cup X \cup Y$ at time λ^1 contains all the house nodes h_1, h_2, \dots, h_k . Clearly, it contains h_1 as h_1 becomes unavailable to agent 1 at exactly this time. Since agent k is later assigned a positive amount of h_1 , part (c) of Lemma 1 implies that the set X must contain some copy of agent k with an arc to h_1 (recall that agents with no copy in the cut will not get assigned any amount of houses in the cut); but this copy of agent k will have an arc to h_k as well, because $h_k \succ_k h_1$. Therefore $h_k \in Y$. Applying the same argument, we see that $\{h_1, h_2, \dots, h_k\} \subset Y$. However, house h_1 is declared unavailable to agent 1, which implies her *next* house at time λ_1 should be outside of S . But note that part (c) of Lemma 1 implies that agent 1 cannot be allocated any more of the houses in Y and so p_{1,h_2} must be zero, a contradiction. ■

Proposition 4 *The CC allocation satisfies no justified envy.*

Proof. Let P be the assignment found by the CC algorithm on an instance of the problem. Consider an agent i and let $h_1 \succ_i h_2 \succ_i \dots \succ_i h_n$. Let λ^{t+1} be the epoch at which house h_k is declared unavailable for agent i . By the definition of the CC algorithm, house h_k must have been the best house for agent i at some point (possibly only at λ^{t+1}). Let $s \cup X_{t+1} \cup Y_{t+1}$ be the cut found by the CC algorithm. Note that the only additional agents that i may potentially envy because of h_k 's unavailability (to him) should continue to “consume” or “claim” h_k after time λ^{t+1} . By part (c) of Lemma 1, any such agent i' must have both his best and next agent-nodes in X_{t+1} . However, the CC algorithm always satisfies the following invariant

$$\sum_{\ell=1}^{n_j} u_{j(\ell)}^{t+1} = \sum_{\ell=1}^{n_j} e_{j(\ell)}, \quad (4)$$

for *any* agent j , if n_j is his next house at time λ^{t+1} . In particular, this expression is valid for agent i' .

Now consider the allocation of agent i found by the CC algorithm. By part (c) of Lemma 1, we know that agent i receives exactly λ^{t+1} units from the houses in Y_{t+1} . The discussion in the preceding paragraph implies that any individually rational allocation for agent i' must allocate more than λ^{t+1} units from the set Y_{t+1} . So the allocation of agent i is not individually rational for agent i' . ■

3.5 Computational considerations

To analyze the number of iterations needed for the algorithm to terminate, observe that each iteration (except the last) ends with the occurrence of Event (a) or Event (b) (or both); Event (a) causes at least one agent's next house to change; Event (b) causes at least one agent's best house to change. As any agent prefers her best house to her next house, we see that the number of possibilities for each agent is $O(n)$, so the algorithm terminates in $O(n^2)$ iterations. Each iteration of the algorithm involves finding the smallest breakpoint of a parametric max-flow problem. Even though some of the capacities are nonlinear because of Expression (1), it is clear that we can find each λ_t by solving at most $(n + 1)$ maximum-flow problems from s to t . (Each capacity is a piecewise linear function with at most 2 pieces; and once an agent becomes a "consumer," she cannot become a "claimer" unless her best house changes.) Thus the entire algorithm can be implemented by solving $O(n^3)$ maximum flow problems in a network with $O(n^2)$ nodes and $O(n^3)$ arcs. Our analysis of the running time is very loose, and a more careful implementation will likely be substantially faster, but we do not investigate this aspect any further as it falls outside the scope of this paper.

4 Impossibility Results

The CC mechanism satisfies individual rationality, efficiency and no justified envy, but is not even weakly strategyproof. This is not a coincidence: The following result rules out the existence of a mechanism satisfying individual rationality, efficiency, no justified envy, and weak strategyproofness.

Theorem 2 *Consider the strict preference domain and fix $|I| \geq 3$. Any mechanism satisfying individual rationality, efficiency, and no justified envy cannot be strategyproof, even in the weak sense.*

We note that to prove such an impossibility result for $|I| \geq k$, it is enough to consider the case $|I| = k$ as long as individual rationality is required. Any instance with k agents can be extended to one with a greater number of agents by letting agents $k + 1, \dots, n$ own a distinct house, which they prefer to any other house in the market.

Proof. Let $I = \{1, 2, 3\}$, $H = \{a, b, c\}$, and consider the following preference and endowment profile:

$$\begin{aligned} 1 : & \quad a \succ b \succ c \quad \{c\} \\ 2 : & \quad c \succ a \succ b \quad \{a\} \\ 3 : & \quad a \succ c \succ b \quad \{b\} \end{aligned}$$

Individual rationality dictates that $p_{2b} = 0$. Furthermore, ordinal efficiency implies that $p_{2a} = 0$; otherwise agent 2 can exchange a part of her share of a for house c from agents 1 or 2, resulting in a Pareto improving assignment for all agents. From these two observations, we get $p_{2c} = 1$, and $p_{1c} = p_{3c} = 0$. Since any allocation that 3 obtains will be individually rational for 1 and vice versa, no justified envy implies that 1 and 3 need to receive identical allocations. Thus we obtain

	a	b	c
1	$\frac{1}{2}$	$\frac{1}{2}$	0
2	0	0	1
3	$\frac{1}{2}$	$\frac{1}{2}$	0

Now consider what happens if agent 1 changes her preferences to $a \succ c \succ b$. Applying individual rationality for agents 1 and 2, we get $p_{1b} = p_{2b} = 0$, by which $p_{3b} = 1$. Then, from ordinal efficiency we get $p_{1a} = p_{2c} = 1$. For agent 1, this allocation dominates the original one, so weak strategyproofness is violated. ■

Remarks.

1. Observe that this result does not exploit the full power of the no justified envy requirement. For example, the violation of weak-strategyproofness persists as long as agent 1 does not get all of a in the original preference profile: *any* criterion of fairness that rules out $p_{1a} = 1$ in the original preference profile is incompatible with individual rationality, ordinal efficiency, and weak strategyproofness.
2. Our definition of NJE assumes that an agent i justifiably envies j if her allocation does not dominate j 's allocation and a relevant IR condition is met. Suppose we weaken the dominance-related part of the definition and require that j 's allocation strictly dominate i 's.¹⁰ This new definition makes it easier to find an allocation that satisfies NJE, and therefore makes establishing an impossibility result harder. However, Theorem 2 remains

¹⁰That is, i justifiably envies j if $p_j \succ_i p_i$ and $p_i \succeq_j e_j$.

valid even under this weaker equity criterion: for the given profile, IR and OE force c to be allocated to agent 2. Since agents 1 and 3 have identical preferences on the remaining objects, and any allocation of these objects to the agents is individually rational for both, the only envy-free allocation according to either definition is for them to receive equal amounts of each object. When agent 1 submits the modified preference ordering, IR and OE force the allocation that is stated, so envy is not used in any form in that case.

3. Theorem 2 asserts the incompatibility of individual rationality (IR), ordinal efficiency (OE), no justified envy (NJE), and weak strategyproofness (WSP). We do not know if there are mechanisms that satisfy *every* proper subset of these properties: the CC mechanism satisfies IR, OE and NJE; the PS mechanism satisfies OE, NJE (actually, no-envy) and WSP; we do not know of mechanisms satisfying WSP, IR, and either OE or NJE.

The impossibility result can be strengthened if we insist on strategyproofness in the strong sense: the following result shows that (strong) strategyproofness is incompatible with individual rationality and efficiency. This is somewhat surprising as typically individual rationality and efficiency are viewed as fairly mild requirements. The proof adapts an elaborate construction of Bogomolnaia and Moulin [4] to an environment with endowments and the elegant reasoning they use in their proof of a related (but different) impossibility result.¹¹

Theorem 3 *Consider the strict preference domain and fix $|I| \geq 4$. There is no mechanism that satisfies individual rationality, ordinal efficiency, and strategyproofness.*

Proof. We start with the following fact about strategyproof mechanisms (we omit the easy proof, see Bogomolnaia and Moulin [4]):

Fact 1 *Consider two orderings of houses $\sigma_i = h_1 \succ_i h_2 \succ_i \dots \succ_i h_n$ and $\sigma_{i'} = h'_1 \succ_i h'_2 \succ_i \dots \succ_i h'_n$. Suppose for some k , $\{h_1, \dots, h_k\} = \{h'_1, \dots, h'_k\}$. Consider a mechanism ϕ and suppose p_i and p'_i are the allocations it finds when agent i reports the preference order σ_i and $\sigma_{i'}$ respectively, for some fixed preferences of the other agents. If ϕ is strategyproof, then $\sum_{l=1}^k p_{il} = \sum_{l=1}^k p'_{il}$.*

Our proof of Theorem 3 proceeds by examining a sequence of profiles, noting down, in each case, the implications of the various properties; eventually, we shall show that these implications

¹¹Bogomolnaia and Moulin [4] show that strategyproofness is incompatible with efficiency and equal treatment equals for the random assignment problem with strict preferences.

are inconsistent. In the rest of the proof, we suppress the \succ_i notation in describing an agent's preferences so that $a \succ_i b \succ_i c \succ_i d$ is simply denoted $abcd$; the identity of the agent is usually clear from the context. We also use IR for individual rationality, OE for ordinal efficiency, and SP for strategyproofness.

Consider an instance of the problem with $|I| = 4$ and suppose each agent owns $1/4$ of each house.

- **Profile 1:** In this profile agents 2, 3, and 4 have the preference order $abcd$ and agent 1's preference order is $badc$. By OE, $p_{1a} = 0$ (as every other agent prefers a to b and agent 1 prefers b to a); and by IR, $p_{1b} = p_{1a} + p_{1b} = 1/2$, for $i = 2, 3, 4$. By a similar reasoning, $p_{1d} = 1/2$.

	a	b	c	d
$badc$	0	$\frac{1}{2}$	0	$\frac{1}{2}$
$abcd(3)$				

- **Profile 2:** Consider the profile in which agent 1 has the preference $bdac$, and the others, $abcd$. Here OE implies that $p_{1a} = 0$. For agent 1, SP and Fact 1 applied to profiles 1 and 2, we get $p_{1b} = 1/2$ and $p_{1d} = 1/2$. Also, as $p_{1a} = 0$, we must have $\max\{p_{2a}, p_{3a}, p_{4a}\} \geq 1/3$. Assume, without loss of generality, that the maximum is attained by agent 4, so that $p_{4a} = \beta \geq 1/3 > 1/4$. These observations are summarized as

	a	b	c	d
$bdac$	0	$\frac{1}{2}$	0	$\frac{1}{2}$
$abcd(2)$				
$abcd$	β			

Agent 4 is now special as this agent has the largest amount of a in Profile 2. In the rest of the profiles, agents 2 and 3 always have the preference order $abcd$ (like in Profiles 1 and 2); agents 1 and 4 will have different preference orderings in different profiles and these will be specified in each case.

- **Profile 3:** Here the preference order of agents 1 and 4 is $abdc$ (agents 2 and 3 have the preference order $abcd$). By IR, we have $p_{ia} = p_{ib} = 1/4$ for all i . Then, by OE, we have $p_{1d} = p_{4d} = p_{2c} = p_{3c} = 1/2$. So we get

	a	b	c	d
$abdc$	$\frac{1}{4}$	$\frac{1}{4}$	0	$\frac{1}{2}$
$abcd(2)$	$\frac{1}{4}$	$\frac{1}{4}$	$\frac{1}{2}$	0
$abdc$	$\frac{1}{4}$	$\frac{1}{4}$	0	$\frac{1}{2}$

- **Profile 4:** Now, consider the profile in which agent 1's preference is $badc$ and agent 4's is $adbc$. By OE, $p_{1a} = 0$, and by IR $p_{1b} \geq 1/2$. Suppose $p_{1b} = 1/2 + \alpha$, for some $\alpha \geq 0$. By IR, agents 2 and 3 collectively own at least $1/2$ units each of the bundle $\{a, b\}$; and agent 4 owns at least $1/4$ units of the bundle $\{a, b\}$. These observations imply $\alpha \leq 1/4$.

	a	b	c	d
$badc$	0	$\frac{1}{2} + \alpha$		
$abcd(2)$				
$adbc$				

- **Profile 5:** Now, consider the profile in which 1's preference is $bdac$ and 4's is $adbc$. Clearly, OE implies $p_{1a} = 0$. Also, SP and Fact 1 applied to agent 1 and Profiles 4 and 5 yields $p_{1b} = 1/2 + \alpha$. Furthermore, applying Fact 1 to agent 4 and Profiles 2 and 5, we get $p_{4a} = \beta > 1/4$. Now let $p_{1d} = \gamma$ and $p_{1c} = \epsilon$ and $p_{4d} = \delta$. All these observations are summarized as

	a	b	c	d
$bdac$	0	$\frac{1}{2} + \alpha$	ϵ	γ
$abcd(2)$				
$adbc$	β			δ

From Profile 4, we already know that $\alpha \leq 1/4$. We shall now show a sharper bound on α . Note that $p_{1a} + p_{1b} = 1/2 + \alpha$; by IR, agents 2 and 3 own at least $1/2$ units each of the bundle $\{a, b\}$, and agent 4 owns at least $1/3$ units of a (as $\beta \geq 1/3$). Adding up all of these,

we find that these four agents collectively own at least $11/6 + \alpha$ units of the bundle $\{a, b\}$; this implies $\alpha \leq 1/6 < 1/4$ as only two units of the bundle $\{a, b\}$ are available.

We shall now argue that $\gamma > 1/4$. If $\epsilon = 0$ then $\gamma = 1 - 1/2 - \alpha \Rightarrow \gamma > \frac{1}{4}$ because, we just saw that $\alpha < 1/4$. If, on the other hand, $\epsilon > 0$, then OE implies that $p_{2d} = p_{3d} = 0$. As $\beta > 1/4$, we must have $\delta < 3/4$, which implies $\gamma = 1 - \delta > 1/4$. In either case, we have $\gamma > 1/4$.

- **Profile 6:** Suppose agent 1's preference is $abdc$ and agent 4's is $adbc$. By IR, $p_{ia} = 1/4$ for all i . Then SP and Fact 1 applied to agent 1 and Profiles 4 and 6 implies $p_{1b} = 1/4 + \alpha$. Furthermore SP and Fact 1 applied to agent 1 and Profiles 5 and 6 implies $p_{1d} = \gamma > 1/4$. Consider agent 4. By OE, $p_{4b} = 0$; by SP and Fact 1 applied to agent 4 and Profiles 3 and 6, we have $p_{4a} + p_{4b} + p_{4d} = 1$, which implies $p_{4d} = 3/4$. So we get

	a	b	c	d
$abdc$	$\frac{1}{4}$	$\frac{1}{4} + \alpha$	ϵ	γ
$abcd(2)$	$\frac{1}{4}$			
$adbc$	$\frac{1}{4}$	0		$\frac{3}{4}$

If $\gamma > \frac{1}{4}$ we have $\gamma + \frac{3}{4} > 1$, a contradiction.

■

Remark. The result in Theorem 3 is sharp in the sense that there are mechanisms satisfying every proper subset of the properties in the statement of the theorem. The CC mechanism satisfies OE and IR; the *identity* mechanism that sets the final allocation to the initial endowments satisfies SP and IR; and any serial dictatorship that ignores the endowments—order the agents in some way and let them choose their best available houses in that order—satisfies SP and OE.

Our next result shows that NJE and EENE are incompatible in the presence of ordinal efficiency and individual rationality.

Theorem 4 *Consider the strict preference domain and fix $|I| \geq 5$. Any mechanism satisfying individual rationality and ordinal efficiency cannot simultaneously satisfy no justified envy and equal-endowment no envy.*

Proof. Let $I = \{1, 2, 3, 4, 5\}$, $H = \{a, b, c, d, e\}$, and consider the following preference and endowment profile:

1 :	<i>adbe</i>	$\{1/2b, 1/2e\}$
2 :	<i>aed</i>	$\{d\}$
3 :	<i>abe</i>	$\{1/2b, 1/2e\}$
4 :	<i>bc</i>	$\{c\}$
5 :	<i>ca</i>	$\{a\}$

The agent preference lists are not complete, but they can be made complete by ranking the remaining houses arbitrarily—the IR constraint will ensure that no agent gets any fraction of the houses that are not in the given preference lists. For convenience, therefore, we work with the shortened preference lists.

We examine the implications of imposing IR, OE, and EENE on this example. The restrictions imposed by IR are summarized as:

	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>	<i>e</i>
1	p_{1a}	p_{1b}	0	p_{1d}	p_{1e}
2	p_{2a}	0	0	p_{2d}	p_{2e}
3	p_{3a}	p_{3b}	0	0	p_{3e}
4	0	p_{4b}	p_{4c}	0	0
5	p_{5a}	0	p_{5c}	0	0

We use OE to further refine the set of possible allocations. Suppose $p_{1b} > 0$. This implies $p_{4b} < 1$, which implies $p_{4c} > 0$. Then, we must have $p_{5c} < 1$, which implies $p_{5a} > 0$. But then agents 1, 4 and 5 can perform mutually beneficial trade, which violates OE. We conclude that $p_{1b} = 0$. A similar argument establishes that $p_{3b} = 0$. As b can only be allocated to agents 1, 3, or 4, we must have $p_{4b} = 1$. But this implies $p_{4c} = 0$, which implies $p_{5c} = 1$, which, in turn, implies $p_{5a} = 0$. Thus, the allocation matrix should be:

	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>	<i>e</i>
1	p_{1a}	0	0	p_{1d}	p_{1e}
2	p_{2a}	0	0	p_{2d}	p_{2e}
3	p_{3a}	0	0	0	p_{3e}
4	0	1	0	0	0
5	0	0	1	0	0

In any allocation that satisfies EENE, agents 1 and 3 must get the same amount of object a ; suppose $p_{1a} = p_{3a} = x$. Then $p_{3e} = 1 - x$, and by IR for agent 3, $x \geq 1/2$. Therefore $x = 1/2$ and so $p_{1a} = p_{3a} = p_{3e} = 1/2$, which implies $p_{2a} = 0$. Finally, agents 1 and 2 have opposite preferences on the objects d and e , and by OE, agent 1 cannot be allocated any amount of e when agent 2 gets a positive amount of d . Thus, we must have $p_{2e} = 1/2$, $p_{1e} = 0$, $p_{1d} = p_{2d} = 1/2$. The final allocation matrix is

	a	b	c	d	e
1	$\frac{1}{2}$	0	0	$\frac{1}{2}$	0
2	0	0	0	$\frac{1}{2}$	$\frac{1}{2}$
3	$\frac{1}{2}$	0	0	0	$\frac{1}{2}$
4	0	1	0	0	0
5	0	0	1	0	0

(5)

In this allocation agent 2's allocation is individually rational for agent 1, and yet agent 2 prefers 1's allocation to his own. Thus, agent 2 justifiably envies agent 1.¹² ■

Comparing EENE and NJE. The example considered in the proof of Theorem 4 displays the tension between NJE and EENE. If we run the CC algorithm on that particular instance of the problem the allocation we obtain is the following:

	a	b	c	d	e
1	$\frac{1}{4}$	0	0	$\frac{1}{4}$	$\frac{1}{2}$
2	$\frac{1}{4}$	0	0	$\frac{3}{4}$	0
3	$\frac{1}{2}$	0	0	0	$\frac{1}{2}$
4	0	1	0	0	0
5	0	0	1	0	0

(6)

This allocation satisfies NJE but fails EENE for agents 1 and 3. In comparing it to the unique allocation that satisfies EENE, it is not entirely clear which of the two is fairer—In allocation (5), agents 1 and 3 do not envy each other, but agent 2 is justified in being unhappy with her allocation, when comparing it to that of 1. Conversely, in allocation (6), agent 2 is not justified in envying agent 3 (even though 3 gets more of house a than she does) but agent 1 strictly prefers 3's

¹²This remains true even if we alter the NJE definition along the lines of the second remark following Theorem 2.

allocation to her own, even though she came to the market with an identical endowment. This example suggests an inherent difficulty in reconciling property rights with concepts of fairness when endowments and allocations are fractional.

Finally, as the model considered here generalizes several well-studied models in the literature, the impossibility results of these special cases automatically carry over. The most prominent of these are stated in the following theorem.

Theorem 5

- (i) (*Bogomolnaia and Moulin [4]*) *Consider the strict preference domain and fix $|I| \geq 4$. There is no mechanism that satisfies ordinal efficiency, equal treatment of equals, and strategyproofness.*
- (ii) (*Yilmaz [10]*) *Consider the strict preference domain and fix $|I| \geq 3$. There is no mechanism that satisfies individual rationality, no justified envy, and strategyproofness.*
- (iii) (*Katta and Sethuraman [5]*) *Consider the full preference domain and fix $|I| \geq 3$. There is no mechanism that satisfies ordinal efficiency, envy-freeness, and weak strategyproofness.*

5 Extensions

We consider three extensions of the basic model treated in §2. In each case we briefly discuss how the algorithm and results extend.

Full Domain. The CC algorithm can be generalized in a straightforward manner to the full preference domain. For every agent i we introduce a node $i_{(k)}$ representing the *set* of his or her k 'th most preferred houses (unlike the case of strict preferences this need not be a singleton set) and connect it to the source with an edge $(s, i_{(k)})$. We include edges from $i_{(k)}$ to every house that is included in agent j 's k most preferred sets of houses. The algorithm then extends naturally. For each agent we—in effect— treat her bundle of equally preferable houses as a single house. Her endowment over this “house” is the sum of her endowments over the original houses. In the algorithm it is this quantity that we treat as the endowment of the “house” representing her k th best set of houses. The proofs of individual rationality, ordinal efficiency, and no justified envy extend easily.

One difference worth mentioning is that, as the algorithm encounters breakpoints that force the maximum s - t flow below n , there may be many ways to redistribute the houses in the bottleneck set. For example, consider

$$\begin{array}{lll} 1 : & a \sim b \succ c & \{1/2b, 1/2c\} \\ 2 : & a \sim b \succ c & \{1/2b, 1/2c\} \\ 3 : & b \succ c \succ a & \{a\} \end{array}$$

At time $\lambda = 2/3$ the set $\{a, b\}$ will become part of the min-cut, and so will be fully allocated to the agents. Agent 3 must get $2/3$ units of b but the way in which agents 1 and 2 split house a and the remaining $1/3$ units of b does not matter.

Arbitrary Endowment Profiles. The CC algorithm can be naturally adapted to cover general endowment profiles where $\sum_h e_{ih} \neq 1$ for some i . In those cases the capacities of outgoing arcs $(s, i_{(\cdot)})$ would simply sum to $\sum_h e_{ih}$, a non-negative number which could be greater or smaller than 1. If this number is smaller than 1, then we proceed as before with the only difference that node $i_{(n)}$ is granted a capacity of $1 - \sum_{l < n} e_{i_{(l)}}$ rather than $e_{i_{(n)}}$. If, on the other hand, this quantity is greater than 1, then we do not change anything and assign capacities to arcs $(s, i_{(k)})$ in the usual manner. Of course, these changes may affect the quantity of a house available in the market, which leads us to the next extension.

Arbitrary Fractions of Houses Available in the Market. The CC algorithm can be adapted to the case in which a non-unit fraction of a house h , say w_h , is present in the market. If $w_h \leq 1$ then the only modification we need to make is to set the capacity of arc (h, t) to w_h instead of 1. If $w_h > 1$, then we may split this house into two or more identical copies (thus increasing the number of houses in the market), such that all but one of these have $w_h = 1$ and exactly one has $w_h \leq 1$. We then set the capacities of arcs (h, t) to w_h . Since our algorithm can deal with indifferences, this poses no problem. The breakpoints in the algorithm are subsequently arrived at when the max flow drops below the quantity $\min\{\sum_h w_h, n\}$.

A number of other extensions such as unequal numbers of agents and houses, agents declaring certain houses as “unacceptable” etc. can be accommodated as well. Since these changes are much more straightforward and they have been addressed before in the literature on random assignment problems (see Bogomolnaia and Moulin [4], Katta and Sethuraman [5], for example), we do not discuss them in more detail here.

6 Future Research

We have provided a computationally-efficient algorithm for finding an assignment satisfying individual rationality, ordinal efficiency and no justified envy in generalized house allocation markets. We have shown that these properties are inconsistent with strategyproofness even in the weak sense. We have also shown that strategyproofness in the strong sense is incompatible with the (very reasonable) requirements of individual rationality and ordinal efficiency. In light of this and other impossibility results, a natural question to ask is whether there exists a mechanism that is individually rational, ordinally efficient, and weakly strategyproof.

Furthermore, while we have shown that IR, OE, NJE, and EENE are incompatible, our knowledge of mechanisms satisfying proper subsets of these properties is scarce. In particular, additional work is needed to clarify EENE and its compatibility with proper subsets of IR, OE, and NJE.

We also wish to comment on what we consider to be a limitation of the CC algorithm. Consider a preference and endowment profile for which there exists an assignment that is individually rational, ordinally efficient, and envy-free. In such a case, it is not unreasonable to expect the mechanism to find such an assignment. That the CC mechanism fails to do so is shown in the following example:

Example 6. Consider the following endowment and preference profile

$$\begin{array}{ll}
 1 : & a \succ b \succ c \succ d \quad \{5/18 \ a, 11/18 \ b, 1/9 \ d\} \\
 2 : & c \succ a \succ b \succ d \quad \{7/18 \ b, 1/2 \ c, 1/9 \ d\} \\
 3 : & c \succ b \succ a \succ d \quad \{7/18 \ a, 1/2 \ c, 1/9 \ d\} \\
 4 : & a \succ d \succ b \succ c \quad \{1/3 \ a, 2/3 \ d\}
 \end{array}$$

It is clear that agents 1, 2, and 3 will not trade any of their endowment for any portion of house d . Thus, agent 4's allocation will necessarily be $(1/3, 0, 0, 2/3)$. Furthermore, agents 2 and 3 will each receive $1/2$ of c , while 1, 2 and 3's allocation of d will remain fixed at $1/9$. An allocation that is individually rational, ordinally efficient, and envy-free is the following:

	a	b	c	d
1	$\frac{1}{3}$	$\frac{10}{18}$	0	$\frac{1}{9}$
2	$\frac{1}{3}$	$\frac{1}{18}$	$\frac{1}{2}$	$\frac{1}{9}$
3	0	$\frac{7}{18}$	$\frac{1}{2}$	$\frac{1}{9}$
4	$\frac{1}{3}$	0	0	$\frac{2}{3}$

On the other hand, the algorithm will allocate $1/3$ units of a to agent 4 (at which point houses a, b, c become unavailable to her) while it will allocate more than $1/3$ units of a to agent 1. Specifically, the CC algorithm finds the following assignment:

	a	b	c	d
1	$\frac{7}{12}$	$\frac{11}{36}$	0	$\frac{1}{9}$
2	$\frac{1}{12}$	$\frac{11}{36}$	$\frac{1}{2}$	$\frac{1}{9}$
3	0	$\frac{7}{18}$	$\frac{1}{2}$	$\frac{1}{9}$
4	$\frac{1}{3}$	0	0	$\frac{2}{3}$

In this assignment agent 4 envies 1 since $p_{1a} > p_{4a}$. This leads us to the following natural open question: Among mechanisms that are individually rational and ordinally efficient, is there one that always finds an envy-free assignment whenever there is one?

An interesting topic of research is to consider the core of the associated cooperative game. The most appropriate way to define the core is not apparent; our preliminary investigation suggests mostly negative results, but much remains to be done here. Finally, an interesting (and challenging) open question is to generalize the TTC mechanism to this setting.

References

- [1] A.Abdulkadiroglu and T.Sonmez (1998), “Random serial dictatorship and the core from random endowments in house allocation problems”, *Econometrica*, 66, 689–701.
- [2] A.Abdulkadiroglu and T.Sonmez (1999), “House allocation with existing tenants”, *Journal of Economic Theory*, 88, 233-260.
- [3] Ahuja, Magnanti and Orlin ,“ Network Flows: Theory, Algorithms, and Applications”, Prentice Hall, 1993.
- [4] A.Bogomolnaia and H. Moulin (2001), “A new solution to the Random Assignment problem”, *Journal of Economic Theory*, 100, 295–328.
- [5] A. Katta, and J. Sethuraman (2006), “A Solution to the Random Assignment Problem on the Full Preference Domain”, *Journal of Economic Theory*, 131(1):231–250.
- [6] L. Lovasz and M. D. Plummer (1986), *Matching Theory*, Annals of Discrete Mathematics, Vol. 29, North-Holland.

- [7] J. Ma (1994), “Strategy-proofness and strict core in a market with indivisibilities”, *International Journal of Game Theory*, 23, 75–83
- [8] J. Sethuraman (2001), “A new solution to the house allocation problem with existing tenants,” unpublished manuscript.
- [9] L. Shapley and H. Scarf (1974), “On cores and indivisibility,” *Journal of Mathematical Economics*, 1, 23–28.
- [10] O. Yilmaz (2009), “Random assignment under weak preferences,” *Games and Economic Behavior*, 546–558.