HARVARD Kennedy School
JOHN F. KENNEDY SCHOOL OF GOVERNMENT

Endowments, Exclusion, and Exchange

Faculty Research Working Paper Series

Ivan Balbuzanov
University of Melbourne

Maciej H. Kotowski
Harvard Kennedy School

Updated September 2017

RWP17-016

Endowments, Exclusion, and Exchange*

Ivan Balbuzanov† Maciej H. Kotowski‡

September 2, 2017

Abstract

We propose a new cooperative solution for discrete exchange economies and resource allocation problems, the exclusion core. The exclusion core rests upon a foundational idea in the legal understanding of property, the right to exclude others. By reinterpreting endowments as a distribution of exclusion rights, rather than as bundles of goods, our analysis extends to economies with qualified property rights, joint ownership, and social hierarchies. The exclusion core characterizes a generalized top trading cycle algorithm in a large class of economies, including those featuring private, public, and mixed ownership; it is neither weaker nor stronger than the (strong) core.

Keywords: Exchange Economy, Property Rights, Core, Top Trading Cycles, House Exchange
JEL: C71, D47, K11

*We thank seminar and workshop audiences at Boston College, James Cook University, McGill University, Oxford University, Stanford University, University College London, and the University of Melbourne for constructive comments. We thank Omer Ali, Chris Avery, Yeon-Koo Che, Jerry Green, Sangram Kadam, Fuhito Kojima, Matt Leister, Jian Li, Simon Loertscher, Jordi Massó, Alex Nichifor, Aniko Öry, Bobby Pakzad-Hurson, Anmol Ratan, Michael Richter, Antonio Rosato, Al Roth, Soroush Saghafian, Ran Shorrer, Tayfun Sönmez, Alex Teytelboym, Utku Ünver, Rakesh Vohra, Bill Zame, and Richard Zeckhauser for helpful commentary and discussions of this research. Hansenard Piou provided excellent research assistance. Part of this research was carried out when Maciej H. Kotowski was visiting the Stanford University Economics Department. He thanks Al Roth and the department for their generous hospitality. The authors declare that they have no relevant or material financial interests that relate to the research described in this paper.

†Department of Economics, The University of Melbourne, Level 4 111 Barry Street, Carlton VIC 3010, Australia. E-mail: <ivan.balbuzanov@unimelb.edu.au>

‡John F. Kennedy School of Government, Harvard University, 79 JFK Street, Cambridge MA 02138, USA. E-mail: <maciej_kotowski@hks.harvard.edu>

There is a striking contrast between the simplicity of endowments in economic models and the complexity of property in practice. In the former, agents are presumed to own some goods with little elaboration; efficient exchange is the usual corollary. In practice, ownership is hardly so straightforward. Co-owners of a house may be tenants in common or joint tenants, subject to other bespoke arrangements. Socially-recognized, but formally undocumented, claims to land are common in the developing world. More abstract forms of property further muddy the waters. In an intellectual property dispute, parties may assert ownership of the same invention; preventing others from using your idea is a patent's chief purpose. In light of such cases, the definition of an endowment and its relation to common understandings of property is elusive, and little studied.

What does an "endowment" mean? How is it related to legal interpretations of property? Do solutions to exchange and allocation problems capture these characteristics correctly? To answer these questions, we study an exchange economy that places endowments and property at the forefront. While Shapley and Scarf's (1974) seminal "house market" is a benchmark case, our setting is far more general. An agent may own multiple goods, none at all, or be a co-owner with others. As in practice, property rights may be clearly defined, caught in a web of competing claims, or even determined by relationships or social obligations.

Our model reveals new weaknesses of classic solutions to allocation problems, such as the core. Thus, our key contribution is the development of a new cooperative solution concept for exchange economies and allocation problems, which we call the exclusion core. The exclusion core's foundation is a reinterpretation of endowments as a distribution of exclusion rights, rather than as bundles of things to trade. Drawing on a simple idea—the ability to exclude others from goods in one's own endowment—the exclusion core identifies intuitively-compelling outcomes, even when the core is empty, excessively large, or as yet undefined.

At a high level, the exclusion core bridges two foundational insights, one in the legal understanding of property and the other in the economic theory of exchange. First, the exclusion core draws heavily on one defining principle of property—the right to exclude others. This right is a classic characterization of property, with roots in the mid-eighteenth century writings of William Blackstone and, later, Jeremy Bentham (Merrill, 1998). The United States Supreme Court has called this right among "the most essential sticks in the bundle of rights that commonly characterize property."[1] We show how this principle not only defines the boundaries of each agent's endowment, but is also sufficient to guide a decentralized market toward an efficient outcome.

[1] *Kaiser Aetna v. United States*, 444 U.S. 164 (1979).

Second, the exclusion core has a close relationship with acclaimed trading protocols, particularly in our model. Agents in our economy have single-unit demand, goods are indivisible, and there are no transfers. This spartan setting insulates our analysis from confounds associated with agents' preferences and lets us focus squarely on the property and endowment variables of interest. We show that the exclusion core has a close association with David Gale's top trading cycle (TTC) algorithm (Shapley and Scarf, 1974). We describe this trading procedure after introducing our model, but we remark here that its theoretical and practical importance cannot be overstated. Its idea of cyclic market clearing is an alluring metaphor for trade. And, suitably generalized, it underpins proposed solutions to many recent market design problems, including transplant organ exchange (Roth et al., 2004), student-school assignment (Abdulkadiroğlu and Sönmez, 2003), airport landing slot allocation (Schummer and Vohra, 2013), and refugee resettlement (Delacrétaz et al., 2016). The exclusion core characterizes a generalization of the TTC algorithm in a large class of economies, including those with private, public, and mixed ownership. Thus, the exclusion core's relevance to the above applications is immediate, though its logic applies broadly.

In the following section we propose a simple example that conveys the essence of our solution while also highlighting the limitations of classic approaches. More importantly, however, we explain how the idea of exclusion governs many other allocation problems, even those not typically interpreted as exchange economies. By focusing on the distribution of exclusion rights, it becomes possible to analyze economies with well-defined, contested, and even conflicting claims to goods with a common toolkit.

We divide our formal analysis into two parts that differ in the relative complexity of the property rights studied within. In Section 2 we treat endowments as an exogenous primitive. This simplified setting allows us to introduce the direct exclusion core and its refinement, the exclusion core. The latter is the focus of our analysis. The exclusion core coincides with the strong core in Shapley and Scarf's (1974) "house market," a benchmark case.[2] Generally, however, the exclusion core is neither a subset nor a superset of the strong core. Unlike the strong core, the exclusion core is never empty in our model and, unlike the weak core, its outcomes are always efficient.

In Section 3 we apply the concept of the exclusion core to situations where social or legal constraints introduce conflicting claims to goods. To model these cases, we introduce

[2]We define the strong and weak cores in Section 1. More formal definitions are presented in Section 2. Care is required as both strong and weak cores have been called "the core" by different authors. The strong core is defined with weak domination. The weak core is defined with strict domination. What we call the strong core is sometimes called the strict core.

relational economies where priorities over objects encode relationships among agents and conditional endowments describe an endogenous distribution of exclusion rights. We adopt the term "priorities" to highlight a technical parallel with discrete allocation problems, particularly those concerning student-school assignment (Abdulkadiroğlu and Sönmez, 2003). However, priorities play a novel role in our model. They are not a rationing device in a centralized assignment scheme. Rather, they constrain a decentralized market's endogenous endowment system and, thus, indirectly govern exclusion rights. We propose three versions of our solution applicable to relational economies, the strong, weak, and unconditional exclusion cores. These all stand on the same behavioral foundation as developed in Section 2, but differ only in how priorities map into endowments and exclusion rights. When the priority structure is acyclic, the strong and weak exclusion cores coincide. In this case, and unlike the strong core, they characterize a generalized TTC algorithm and are stable in the sense of von Neumann and Morgenstern (1944). Thus, the exclusion core embodies a self-enforcing "standard of behavior." Acyclic priority structures are common in practice, and include economies with private, public, and mixed ownership.

While we reference the related literature throughout our exposition, we offer a more structured survey in Section 4. Our paper contributes to the study of exchange economies by proposing the exclusion core, a new solution clarifying the foundations of exchange. Our results also have implications for the practice of market design, which often involves the identification of "good" centralized allocation procedures. We do not pursue this normative objective, although we do offer a new rationale for the use of trading cycle algorithms in applications, including those cited above. We show that such trading procedures identify outcomes that are robust to the exercise of direct and indirect exclusion rights, which are closely tied to a classic characterization of property.

Our solution's motivation also lets us contribute to a debate hitherto limited to legal scholars concerning the *in rem* and *in personam* interpretations of property (Merrill and Smith, 2001b). As explained in Section 4, this debate should also be of interest to economists as it addresses the foundations of impersonal exchange, commonly presumed in economic analysis, and the informational complexity of markets. An *in rem* interpretation of property is grounded in a few universal principles, with the right to exclude chief among them. Section 2 captures this paradigm. An *in personam* interpretation grounds property in a web of inter-agent relationships and obligations. We introduce relational economies with priority structures in Section 3 to model this perspective. Our analysis provides the first formal framework bridging both paradigms.

Though property rights are a touchstone for our analysis, the questions we consider are distinct and unrelated to the bilateral externalities examined by Coase (1960) or the incentive implications investigated by Grossman and Hart (1986) and Hart and Moore (1990). Our focus on economy-wide allocations and trade sets our analysis apart from this literature, which studies contractual arrangements.

We summarize our contributions and conclude in Section 5. With the exception of some immediate corollaries, we relegate all proofs to the Appendix.

1 Motivating Examples

To motivate our argument, it is helpful to first consider a straightforward allocation problem, which is an instance of our model. It highlights the limitations of existing theories and hints at the power of our interpretation emphasizing exclusion.

Example 1 (The Kingdom)**.** There are three agents—i, j, and k—and two indivisible goods, called houses—h_1 and h_2. At most one agent can live in a house and each agent has use for at most one house. Everyone strictly prefers h_1 to h_2 and there is no other medium of exchange. Assume that agent k, whom we call the King, initially owns *both* houses.

Which final allocation of houses will, or should, arise in this economy? First, since the King owns both houses, he will surely live in h_1. As he cannot live in more than one house, h_2 should be occupied by either i or j. Either outcome is efficient. Finally, one agent, again either i or j, will remain homeless as there are fewer houses than agents. Thus, either of two allocations is intuitive, justifiable, and efficient.[3]

It is surprising that neither the strong core nor the weak core is able to converge on the preceding assignments. Along with competitive equilibrium, these are the two most prominent solutions for exchange and assignment economies. An allocation belongs to the strong core if there does not exist a coalition of agents that can reallocate the houses they own such that no coalition member is made worse off and at least one coalition member becomes strictly better off. In the above example, *the strong core is empty.* Every arrangement can be improved upon, or "blocked," by some coalition. For example, if k is assigned to h_1, i to h_2, and j is homeless, j and k can together reallocate h_2 to benefit j. If instead house h_1 is

[3]The example's phrasing follows that of our model. An alternative framing is inspired by kidney exchange, a notable application (Roth et al., 2004). Agent k has two kidneys, h_1 and h_2, and has volunteered to be a live organ donor. There are two compatible recipients, i and j, who are equally deserving to receive a donated organ. Clearly, k will keep one kidney and either i or j will get the transplant.

occupied by k and h_2 is occupied by j, i and k can together reallocate h_2 to benefit i.[4]

The weak core is not empty, but it is also dissatisfying. An allocation belongs to the weak core if there does not exist a coalition of agents that can reallocate the houses they own such that all coalition members become strictly better off. In the above example, *the weak core is too large.* In fact, any assignment where agent k inhabits house h_1 belongs to the weak core. This includes the odd situation where house h_2 is vacant and both i and j are homeless. Neither i nor j can access h_2 since that house's owner gains nothing from the move. This inefficient outcome is dispiriting and is unlikely to arise in practice.[5]

The Kingdom's troubles are neither special to the example, nor are they technical anomalies. Rather, they are symptoms of a mis-calibration between the power of ownership rights and agents' desire to revise assignments. The strong core presumes agents have greater power than they do in practice, while the weak core fails to recognize the power that agents actually hold. Furthermore, neither core concept credits a coalition for its ability to *obstruct* outcomes, which is a "powerful lever in bargaining" (Shapley and Shubik, 1971, p. 128).

The exclusion core avoids these shortcomings by viewing endowments through the lens of exclusion. An allocation is in the (direct) exclusion core if no coalition can strictly gain from a reassignment of houses that might exclude (i.e. evict) non-coalition members from houses in the coalition's endowment. In the example above, only the two intuitive and efficient outcomes pass this test. In that economy, exclusion rights are vested in agent k. If i or j occupies house h_2, k gains nothing by evicting him and thus is unwilling to do so. Conversely, if house h_2 is vacant, k has no reason to prevent its occupancy. In each case, the King's conduct accords with intuition.

Associating endowments with exclusion rights proved insightful in the preceding example. Importantly, this reinterpretation extends to economies where the rules surrounding property contain ambiguities. The complexities (and headaches) surrounding joint, collective, or ill-defined ownership immediately come to mind. Aside from legal prescriptions, status, relationships, and social conventions also influence how goods are exchanged or allocated. These variables define property rights in practice, often implicitly.

---

[4]A possible remedy for the strong core's emptiness in the Kingdom is to assume, additionally, that the King prefers h_2 to be given to a specific agent. Regrettably, allowing for preferences over allocations, i.e. externalities, begets more problems in general. Even the weak core can be empty (Mumcu and Saglam, 2007).

[5]The core concepts' deficiencies are not driven by the economy's housing shortage. Adding a third, universally least-preferred house h_3, which is also owned by the King, does not change the example's conclusions. Problems occur even if i and j disagree about the relative merits of h_2 and h_3 with, say, j preferring h_3 over h_2. The inefficient allocation where h_2 is assigned to j and h_3 to i is in the weak core.

As an example, consider the allocation of seats on a city bus. This problem is not typically viewed as one of property or trade. Nevertheless, a custom built on hierarchical exclusion applies in most places around the world and the implied property rights are easy to spot. When open seats are plentiful, anyone can sit down. However, an informal priority rule kicks in when seats become scarce. A passenger in need may claim a seat provided the person inconvenienced does not command comparable recognition. A teenager is expected to defer to an elderly passenger; whether a blind man should defer to a pregnant woman, or vice versa, is less obvious. Ambiguities notwithstanding, final outcomes are generally efficient.

An even more complex property regime pertains to transplant organs, which occupy a grey zone between a donor's personal property and a societal resource (Truog, 2005; Cronin and Price, 2008; Cronin and Douglas, 2010). Organs from living donors are typically viewed like personal property, though monetary compensation for donation is generally prohibited. Organs from cadavers, in contrast, are treated as social resources to be distributed to the persons with the "greatest need," as determined by a medical authority. Yet, social conventions muddle this dichotomy. For example, doctors often seek family members' permission before transplanting organs from deceased relatives even though the deceased had consented to donation prior to death (Downie et al., 2008).[6] In Ontario, Canada, around 20 percent of willing donors have their wishes vetoed by their family postmortem (Bigham, 2016). In the United Kingdom, family objections blocked 547 transplants from 2010–16 despite the deceased donor's prior consent (Quinn, 2016). Though the next of kin did not inherit their relative's organs, they sometimes have a right to exclude others from benefiting from them.

We argue that the examples above—the Kingdom, the bus, and the transplant center—are instances of the same economic problem. Each case's peculiarity is due to the prevailing endowment, which is best understood as a distribution of exclusion rights. The exclusion core unifies these and other situations under a common umbrella. It identifies allocations where no coalition can gainfully exercise their rights of exclusion over others. Though the three examples above conform to our formal model, the ideas we propose apply broadly.

2 Simple Economies

A simple economy $\langle I, H, \succ, \omega \rangle$ consists of agents, goods, preferences, and an endowment system. $I = \{i_1, \ldots, i_n\}$ is a finite set of agents whom we sometimes denote by i, j, or k.

[6]We thank Al Roth for bringing this practice to our attention via his blog. Downie et al. (2008) report that 69 percent of health care providers in Canada believed that a bereaved family's wishes would be respected despite the deceased donor's consent to transplantation.

$H = \{h_1, \ldots, h_m\}$ is a finite set of indivisible objects, called houses, that can be allocated among the agents. Each agent may live in at most one house and each house $h \in H$ may shelter at most one agent. A house may be vacant and an agent need not be assigned to a house. We model this latter outcome by the agent's assignment to an outside option $h_0 \notin H$, which has unlimited capacity.[7] An *allocation* $\mu\colon I \to H \cup \{h_0\}$ is an assignment of agents to houses such that $|\mu^{-1}(h)| \leq 1$ for all $h \in H$. We interpret an allocation as the outcome of some centralized or decentralized assignment, bargaining, or exchange process, which we do not model directly. For brevity, we write $\mu(C)$ to denote $\bigcup_{i \in C} \mu(i)$ for any $C \subseteq I$.

Each agent has a strict and rational preference defined over $H \cup \{h_0\}$. If agent i prefers h to h', then $h \succ_i h'$. We write $h \succeq_i h'$ if $h \succ_i h'$ or $h = h'$. For convenience, we sometimes define $\succ_i$ by listing houses in preferred order, i.e. $\succ_i\colon h, h', \ldots$. Unlisted houses are worse than the outside option.

An *endowment system* specifies the houses in each coalition's endowment. It is a function $\omega\colon 2^I \to 2^H$ satisfying three properties.

(A1) Agency: $\omega(\varnothing) = \varnothing$.

(A2) Monotonicity: $C' \subseteq C \implies \omega(C') \subseteq \omega(C)$.

(A3) Exhaustivity: $\omega(I) = H$.

Condition (A1) is an innocuous resolution of a degenerate case. It restricts ownership to agents or groups. (A2) states that a coalition has in its endowment anything that belongs to any sub-coalition. Finally, (A3) says that everything belongs to the grand coalition.

In this section, we further assume that the endowment system satisfies

(A4) Non-Contestability: For each $h \in H$, there exists $C^h \subseteq I$, such that for all $C \subseteq I$, $h \in \omega(C) \iff C^h \subseteq C$.

We call C^h the *minimal controlling coalition* of house h. Condition (A4) guarantees that each house has a well defined set of one or more co-owners without opposing and mutually exclusive claims. We relax (A4) in Section 3.

Many economies satisfy (A1)–(A4), including those examined by Shapley and Scarf (1974) and Hylland and Zeckhauser (1979), which we discuss below. These two cases bracket a class of economies where each house's minimal controlling coalition is either a singleton (and the house is privately owned) or the grand coalition (and the house is part of the social

[7] The outside option is not required for our conclusions when there are sufficiently many acceptable houses.

endowment). Economies in this class have been used to model the allocation of dormitory rooms (Abdulkadiroğlu and Sönmez, 1999) and transplant organs (Roth et al., 2004).

2.1 The Strong and Weak Cores

Which allocations will, or should, arise given the agents' preferences and endowments? Two classic answers to this question are provided by the strong and weak cores. As explained in Section 1, both consist of allocations that cannot be "blocked" by any coalition.

Definition 1. A non-empty coalition $C \subseteq I$ can *weakly block* the allocation μ with allocation σ if

1. $\sigma(i) \succeq_i \mu(i)$ for all $i \in C$ and $\sigma(i) \succ_i \mu(i)$ for some $i \in C$; and
2. $\sigma(C) \subseteq \omega(C) \cup \{h_0\}$.

The *strong core* is the set of allocations that cannot be weakly blocked by any coalition.

Definition 2. A non-empty coalition $C \subseteq I$ can *strongly block* the allocation μ with allocation σ if

1. $\sigma(i) \succ_i \mu(i)$ for all $i \in C$; and
2. $\sigma(C) \subseteq \omega(C) \cup \{h_0\}$.

The *weak core* is the set of allocations that cannot be strongly blocked by any coalition.

Strong core allocations are *Pareto efficient*. That is, no agent can be made strictly better off without harming anyone. The strong core is a subset of the weak core.

Even in simple cases, the strong and weak cores often fail to offer satisfactory guidance, as illustrated by Example 1 above. Their deficiencies can be traced to faults in the two variants of blocking. Often it is too easy for a coalition to weakly block an allocation. This is because weak blocking presumes agents who are indifferent among allocations always agree to join a blocking coalition. Two arguments try to justify this behavior. The first is altruism—an unaffected agent should help others. This is at best an incomplete behavioral justification. Aiding one party very often harms another, which is hardly an altruistic disposition. The second is not-modeled side payments. If an agent benefits from a reallocation, the reasoning goes, he could bribe those who remain indifferent to enforce the reassignment. This argument is unconvincing. Equally well a side payment could be extorted from a potentially harmed agent to prevent a blocking coalition's formation. More fundamentally, any relevant transfers should be modeled directly. In our case, they are absent to preserve the model's simplicity. In some applications, such as kidney exchange, transfers are prohibited (Roth et al., 2004).

Strong blocking is immune to the questionable incentives that plague weak blocking, but it suffers from the opposite ailment. Often, it is too difficult for a blocking coalition to form because agents who benefit from a reallocation of houses cannot induce those who remain indifferent to cooperate. Consequently, unintuitive and inefficient outcomes persist.

2.2 The Direct Exclusion Core

Acknowledging the problems encountered by classic versions of the core, we propose an alternative solution. Our proposal reverts to a fundamental tenet of property—the right to exclude others. An agent exercising this right can prevent others from using property in his endowment, thus securing and preserving his wellbeing. We explain this idea's implications in two steps. As a heuristic, we first define the direct exclusion core to show the immediate consequences of the right to exclude. We then refine this solution by considering exclusion's indirect implications. This refinement leads to the exclusion core.

As a motivating case, consider an economy with three agents and three houses. Each house h_k is owned by agent i_k and the agents' preferences are

$$\succ_{i_1}: h_2, h_3, h_1 \qquad \succ_{i_2}: h_1, h_2 \qquad \succ_{i_3}: h_1, h_3 \ .$$

Consider the allocation

$$\mu(i_1) = h_3 \qquad \mu(i_2) = h_2 \qquad \mu(i_3) = h_1 \ .$$

Though Pareto efficient, i_1 and i_2 can strongly block μ with the allocation

$$\sigma(i_1) = h_2 \qquad \sigma(i_2) = h_1 \qquad \sigma(i_3) = h_3 \ .$$

The traditional interpretation of the move from μ to σ is that the coalition $C = \{i_1, i_2\}$ strictly gains by reallocating the houses in its endowment, $\omega(C) = \{h_1, h_2\}$. This is true, but another feature of this reallocation is noteworthy. The only agent harmed by the change was i_3. He was *excluded* from $\mu(i_3) = h_1$—a house in the coalition's endowment. In fact, the eviction of i_3, or the repossession of house h_1, is a prerequisite for i_1 and i_2 to reallocate h_1 among themselves. This feature hints at an alternative feasibility condition for blocking. A coalition can block an assignment whenever its members gain from an alternative and when those harmed by any reallocation were excluded from houses belonging to the coalition.

Definition 3. A non-empty coalition $C \subseteq I$ can *directly exclusion block* the allocation μ with allocation σ if

1. $\sigma(i) \succ_i \mu(i)$ for all $i \in C$; and

2. $\mu(j) \succ_j \sigma(j) \implies \mu(j) \in \omega(C)$.

The *direct exclusion core* is the set of allocations that cannot be directly exclusion blocked by any coalition. Thus, no coalition can gainfully destabilize or obstruct a direct exclusion core allocation by drawing on their collective exclusion rights. This logic differs subtly from the rhetoric of "enforcement" ascribed to classic definitions of blocking.

The direct exclusion core's non-emptiness will be implied by Theorem 1 below. Here we briefly highlight some of its properties.

Lemma 1. *The direct exclusion core is a subset of the weak core.*

Furthermore, direct exclusion core allocations are efficient.[8] Any reallocation of houses that benefits some agents but leaves all others unharmed is admissible. Pareto efficiency is an immediate consequence. In the Kingdom economy (Example 1), the direct exclusion core coincides with the two intuitive and focal allocations, as explained above.

2.3 The Exclusion Core

Like strong blocking, direct exclusion blocking insists that all blocking coalition members strictly benefit from a proposed reallocation. This requirement is seemingly constraining as many desirable reallocations require the acquiescence of unaffected third parties who coincidentally (co-)own a reassigned house. The usual way to relax this constraint is to replace the strict incentive condition (1) in Definition 3 with its weaker cousin. This approach is misguided. The resulting solution would be stronger than the strong core and vulnerable to the same criticisms concerning incentives. Instead, we rationalize the cooperation of third parties by extending the logic of exclusion. An example illustrates the idea.

[8]The direct exclusion core does not generally coincide with the set of Pareto-efficient, weak core allocations. In Example 3, μ is a Pareto-efficient weak core allocation, but it can be directly exclusion blocked.

Example 2. There are six agents and six houses. Each house h_k is owned by agent i_k. The agents' preferences are:

$$\begin{array}{lll} \succ_{i_1}: h_3, h_4, h_1 & \succ_{i_2}: h_1, h_2 & \succ_{i_3}: h_2, h_5, h_3 \\ \succ_{i_4}: h_2, h_4 & \succ_{i_5}: h_6, h_5 & \succ_{i_6}: h_3, h_6 \end{array}.$$

Figures 1(a) and 1(b) illustrate this economy's two direct exclusion core allocations, μ and σ. In each figure, there is a directed link from each house to its owner and from each agent to his assignment. Both allocations belong to the weak core; σ is the only strong core allocation.

Agents i_1 and i_3 strictly prefer their assignment under σ over their assignment under μ. To directly exclusion block μ with σ, agent i_1 needs to move to h_3 and agent i_3 needs to move to h_2, as illustrated in Figure 1(c). The first move is feasible for the coalition. Agent i_3 owns h_3 and he can veto i_6's assignment to that house as mandated by μ. Thereafter, h_3 is available for i_1. The second move is not feasible. While $h_2 = \mu(i_4) \succ_{i_4} \sigma(i_4)$, $h_2 \notin \omega(\{i_1, i_3\})$. Thus, i_1 and i_3 cannot directly exclusion block μ.

Whereas i_1 and i_3 do not own house h_2, we argue that they enjoy a form of indirect control over the house at μ. House h_2 is owned by i_2 for whom $\mu(i_2) = \sigma(i_2) = h_1$ and h_1 is in the coalition's endowment. While i_2 is indifferent between μ and σ, his wellbeing depends on the coalition's continued accommodation. Agents i_1 and i_3 can press i_2 to evict i_4 from h_2 by threatening to displace him from h_1. Acknowledging the power asymmetry at μ, agent i_2 would reasonably accept this demand. By exploiting i_2's dependency, i_1 and i_3 can forge a *repossession chain* giving them an indirect veto over h_2's assignment at μ.

The story is entirely different when the prevailing allocation is σ (Figure 1(d)). The coalition $\{i_4, i_5, i_6\}$ would like to block. However, houses h_2 and h_3 are inaccessible since the coalition lacks leverage over those houses' owners, i_2 and i_3, at σ. The coalition $\{i_4, i_5, i_6\}$ is too isolated. Therefore, σ seems more robust and compelling as a final assignment than μ.

Example 2 shows that the right to exclude can be a powerful, though subtle, stick. It secures an agent's own goods and it offers a channel through which he can access more options. Importantly, the chain of exclusion and repossession need not stop with one link, as in the example. By exploiting the interdependencies implied by exchange, a coalition can inductively relay credible threats of exclusion and eviction to all agents who are indirectly linked to its endowment $\omega(C)$. First, $(\mu^{-1} \circ \omega)(C)$ is the set of agents who are assigned by μ to houses in $\omega(C)$. Thus, with one step of influence, coalition C secures direct and indirect control over $\omega(C_1)$ where $C_1 = C \cup (\mu^{-1} \circ \omega)(C)$. At two steps of influence, it secures control over $\omega(C_2)$ where $C_2 = C_1 \cup (\mu^{-1} \circ \omega)(C_1)$. And so on. The recursive form ensures that a

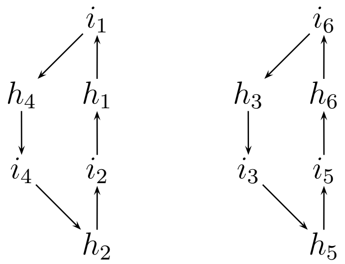

(a) Allocation μ.

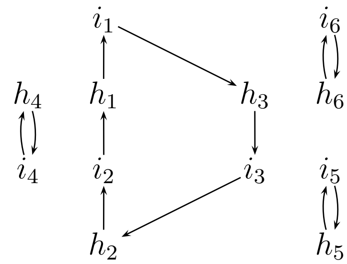

(b) Allocation σ.

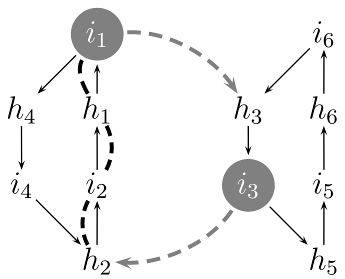

(c) A repossession chain at μ. The coalition $\{i_1, i_3\}$ can indirectly access h_2.

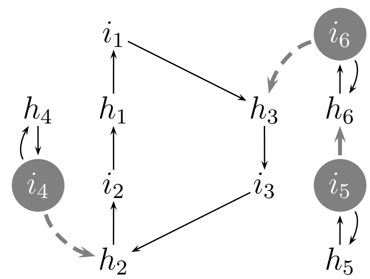

(d) Absence of a repossession chain at σ. The coalition $\{i_4, i_5, i_6\}$ cannot access h_2 or h_3.

Figure 1: Direct exclusion core allocations in Example 2.

collectively-owned house is included once all co-owners are deemed (indirectly) dependent on the coalition's endowment.

Definition 4. The *extended endowment* of coalition C (at μ) is

$$\Omega(C|\omega, \mu) := \omega\left(\bigcup_{k=0}^{\infty} C_k\right)$$

where $C_0 = C$ and $C_k = C_{k-1} \cup (\mu^{-1} \circ \omega)(C_{k-1})$ for every $k \geq 1$.[9]

A coalition's extended endowment does not bestow upon the group more property rights. Rather, it better reflects the relative power agents enjoy when exclusion, or threats thereof, underpin interaction. It also strengthens our previous blocking definition in a natural way.

Definition 5. A non-empty coalition $C \subseteq I$ can *indirectly exclusion block* the allocation μ with allocation σ if

[9]The infinite union in this definition simplifies notation. Since $C_{k-1} \subseteq C_k \subseteq I$ for all k, the union is finite in practice.

1. $\sigma(i) \succ_i \mu(i)$ for all $i \in C$; and,

2. $\mu(j) \succ_j \sigma(j) \implies \mu(j) \in \Omega(C|\omega, \mu)$.

The *indirect exclusion core* or, for simplicity, the *exclusion core* is the set of allocations that cannot be indirectly exclusion blocked by any coalition.

Definition 5 has two components. First, all coalition members must strictly benefit from the proposed reallocation of houses. We explained the importance of this requirement when discussing the strong core above. Second, agents harmed by a blocking action must have been excluded from houses in the coalition's extended endowment. Paralleling Definition 3, indirect exclusion blocking does not presume the imposition of a new allocation on the market as a whole. Rather, a coalition's interest and ability—through direct and indirect influence—to impede a particular outcome is all that matters. At an exclusion core allocation, everyone's desire or ability to veto others' assignments is neutralized.

Theorem 1. *For any economy $\langle I, H, \succ, \omega \rangle$, the exclusion core is not empty.*

We defer the proof of Theorem 1 to Section 3 where we generalize our model. Before considering that generalization, we remark on the exclusion core's properties and we examine two important special cases. The exclusion core is a subset of the direct exclusion core and its elements are Pareto efficient and belong to the weak core. The exclusion core's relation to the strong core is more nuanced. Since the strong core may be empty, the exclusion core is *not* necessarily a subset of the strong core. The strong core is also *not* necessarily a subset of the exclusion core, as confirmed by the next example.

Example 3. There are four agents and four houses. For each $k \in \{1, 2, 3\}$, $\omega(i_k) = \{h_k\}$. House h_4 is owned collectively, i.e. $h_4 \in \omega(I)$ and $h_4 \notin \omega(C)$ for all $C \subsetneq I$.[10] The agents' preferences are:

$$\succ_{i_1}: h_2, h_1 \qquad \succ_{i_2}: h_4, h_3, h_2 \qquad \succ_{i_3}: h_2, h_3 \qquad \succ_{i_4}: h_1, h_4, h_3, h_2 \ .$$

There are three strong core allocations, μ, ν, and σ, as illustrated in Figure 2. In the figure, each house is pointing to its owner (if it has one) and each agent is pointing to his assigned house. Only ν and σ belong to the exclusion core. The coalition $C = \{i_1, i_2, i_4\}$ can directly (and, hence, indirectly) exclusion block μ with the allocation σ. This coalition cannot weakly block μ with σ since $\sigma(i_2) = h_4$ and $h_4 \notin \omega(C)$.

[10] This economy is an example of a house allocation problem with existing tenants (Abdulkadiroğlu and Sönmez, 1999).

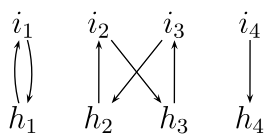

(a) Allocation μ.

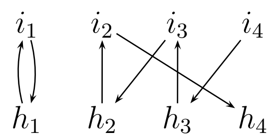

(b) Allocation ν.

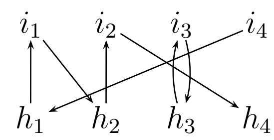

(c) Allocation σ.

Figure 2: Strong core allocations in Example 3. Only ν and σ are exclusion core allocations.

2.4 Private and Public Ownership

Private and public ownership are two dichotomous cases often considered in economic analysis. In a private ownership economy, every house has a single owner. That is, for every $h \in H$ there exists an agent i such that $h \in \omega(i)$. An agent may own multiple houses, as in Example 1, but no house is owned collectively.

Proposition 1. *In a private ownership economy, the strong core is a (possibly empty) subset of the exclusion core.*

Shapley and Scarf (1974) analyze a particular private ownership economy where each agent i_k owns exactly one house, i.e. $\omega(i_k) = \{h_k\}$, and $h \succ_i h_0$ for all i and $h \neq h_0$. They present an algorithm, attributed to David Gale, that constructs a strong core allocation in their market. We will generalize this algorithm in the sequel and we summarize it here.

Algorithm 1 (Top Trading Cycles)**.** Initially, all agents and houses are unassigned. In step $t \geq 1$ of the algorithm, each unassigned house points to its owner and each unassigned agent points to his most preferred house that remains in the market. As there is a finite number of agents and houses, there is at least one cycle of the form $h \to i \to \cdots \to h' \to i' \to h$. (A cycle may be formed by one agent and one house.) Pick any cycle and to each agent in the cycle assign the house that he is pointing to. Remove the assigned agents and houses from the market. This process continues until all agents and houses have been assigned.

The TTC algorithm identifies the economy's unique strong core allocation (Roth and Postlewaite, 1977) and this allocation can be supported as a competitive equilibrium (Shapley and Scarf, 1974). Roth (1982) shows that the TTC mechanism is strategy-proof[11] and Ma (1994) proves that it is the unique mechanism satisfying individual rationality,[12] Pareto

[11] A (direct) mechanism is strategy-proof if it is a dominant strategy for each agent to truthfully communicate his preferences to the mechanism. In each step of the TTC mechanism, each agent is assumed to point to his most-preferred available house. He cannot improve his final assignment by pointing elsewhere.

[12] If μ is an individually rational allocation, then $\mu(i) \succeq_i h$ for all $h \in \omega(i) \cup \{h_0\}$ and for every agent i.

optimality, and strategy-proofness. Furthermore, the strong core allocation is "stable" under multiple definitions (Roth and Postlewaite, 1977; Wako, 1984, 1991; Kawasaki, 2015). All things considered, the strong core allocation is this market's most compelling outcome.

Proposition 2. *The exclusion core and the strong core coincide in Shapley and Scarf's (1974) economy.*

The polar opposite of a private ownership economy is the public ownership economy. Hylland and Zeckhauser (1979) consider this assignment problem.[13] In this case, all houses belong only to the social endowment, i.e. $\omega(C) = \varnothing$ for all $C \subsetneq I$ and $\omega(I) = H$. Given the situation's ex ante symmetry, any Pareto efficient assignment is as a reasonable outcome.[14]

Proposition 3. *The exclusion core and the set of Pareto efficient allocations coincide in Hylland and Zeckhauser's (1979) assignment problem.*

We present direct proofs of Propositions 1–3 in the Appendix. In Section 3.4 we provide indirect proofs of the preceding (and stronger) results. Those arguments rely on a generalized environment that we examine in the following section.

3 Relational Economies

In the previous section, the endowment system ω was an exogenous and unqualified distribution of the right to exclude. Though simple, this formulation overlooks the nuance accompanying ownership rights in practice. These are often layered with caveats and ambiguities. Status and relationships may impart implicit property rights, as in the bus and transplant examples above. Thickets of conditional and competing claims can readily arise.[15] A variant of a prior example illustrates the difficulty of capturing such cases with the rigid interpretation of endowments presumed thus far.

Example 4 (The Diarchy)**.** Recall the Kingdom from Example 1. There are three agents and two houses. Everyone agrees that h_1 is the best house, and h_2 is second best. Suppose, however, that agents j and k are "co-kings" and "co-own" everything. Agent i, the peasant, owns nothing. Now, there are two focal allocations:

$$\mu(i) = h_0 \quad \mu(j) = h_1 \quad \mu(k) = h_2 \qquad \text{and} \qquad \sigma(i) = h_0 \quad \sigma(j) = h_2 \quad \sigma(k) = h_1 \ .$$

[13]See also Koopmans and Beckmann (1957).

[14]Unlike Hylland and Zeckhauser (1979), we focus exclusively on deterministic outcomes.

[15]The thicket metaphor is due to Shapiro (2000), who used it to describe the complex and interdependent claims typifying patents and intellectual property.

In each case, the kings split the two houses while the peasant receives his outside option. Both allocations are efficient and either is equally plausible given the economy's symmetry.

Is there an endowment system that reasonably captures this situation? The problem's symmetry suggests two natural candidates. The first places all houses only in the kings' joint endowment: $\omega(i) = \omega(j) = \omega(k) = \varnothing$ but $\omega(\{j, k\}) = \{h_1, h_2\}$. This endowment system satisfies (A1)–(A4) and the preceding section's analysis applies. Regrettably, the exclusion core includes outcomes that are implausible given the context. For instance, the allocation

$$\nu(i) = h_2 \qquad \nu(j) = h_0 \qquad \nu(k) = h_1$$

belongs to the exclusion core even though a king is homeless. Despite his royal pedigree, j cannot reclaim h_2 without k's assistance. But k is satiated and has no reason to aid j.

An alternative endowment system places each house in each king's personal endowment: $\omega(j) = \omega(k) = \{h_1, h_2\}$ and $\omega(i) = \varnothing$. This endowment system violates (A4) and every allocation can be blocked by the king who does not receive h_1. The exclusion core is empty.

The Diarchy's troubles stem from the proposed endowment systems' immutability and insensitivity to the agents' identities and relationships. More conditionality seems warranted. For instance, if the peasant occupies h_1 or h_2, either king should be able to expel him. However, a king should not have the same right when a house is occupied by a co-monarch, lest a civil war is to follow.

3.1 Priorities

To analyze economies with competing and conditional claims, such as the Diarchy, we appeal to the exclusion core, but we presume that endowments, i.e. exclusion rights, are endogenously determined. To further this idea, we first amend our definition of an economy. A *relational economy* $\langle I, H, \succ, \rhd \rangle$ consists of agents, houses, preferences, and a priority structure. The first three components are defined as before. The new primitive is the priority structure $\rhd = (\rhd_h)_{h \in H}$, which is a family of orders that describe pre-existing social, legal, or economic relationships among agents in relation to the economy's goods. We are agnostic about the priority structure's origin. It may be formally codified by law or it may be informally set by inter-personal relations, relative status, or historical context.[16] We emphasize from the outset that priorities in our model are not synonymous with property or exclusion rights per se. These will be derived below, naturally favoring agents with higher priorities.

[16] A formally-codified priority structure is found among creditors holding an issuer's debt.

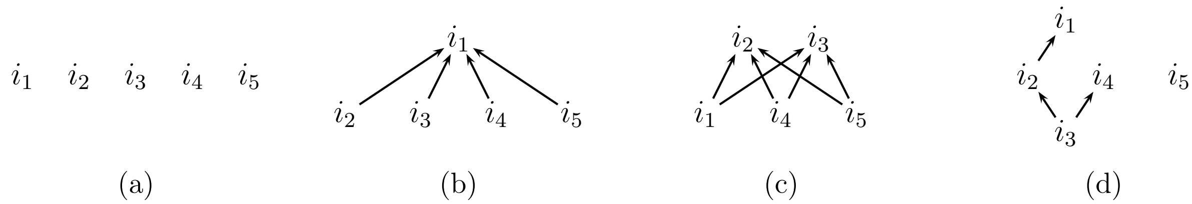

Figure 3: Hasse diagrams of example priority structures.

Formally, each $\rhd_h$ is a strict partial order of the set of agents. We write $i \trianglerighteq_h j$ if $i \rhd_h j$ or $i = j$. Many situations can be modeled with priorities and Figure 3 presents several examples.[17] If a house is part of the social endowment, no agent has priority over others, i.e. $i \ntriangleright_h j$ for all i and j, as in panel (a). If one agent $\rhd_h$-dominates all others, it will prove natural to call him the house's "owner." A diarchic structure occurs if two agents equally dominate others, but not one another. More exotic cases, like in Figure 3(d), may describe hierarchical relationships within families or across social groups. As conventions, we assume that $i \ntriangleright_{h_0} j$ for all $i, j \in I$ and $i \rhd_h \varnothing$ for all $i \in I$ and $h \in H$.

Priorities feature in many assignment models, particularly those concerning student-school matching. We adopt the same terminology to highlight a technical parallel that will be evident below. However, priorities carry a different interpretation in our setting than in a school-choice or centralized assignment problem. In the latter, priorities are administratively-defined rankings of students (the agents) that help ration places at desirable schools (the houses). Abdulkadiroğlu and Sönmez (2003) offer two interpretations of priorities in this context. First, they may impose an inviolable fairness requirement, "no justified envy," on the final assignment.[18] In this case, Gale and Shapley's (1962) deferred acceptance algorithm is the preferred assignment method. Priorities do not have this meaning in our model. Second, priorities may define students' relative opportunities. A student with a higher priority at a school should have a "better opportunity" to attend that school than someone with a lower priority (Abdulkadiroğlu and Sönmez, 2003, p. 736). Our use of priorities is closer to this second meaning. But, as we emphasize, priorities in our model should not be regarded as rationing devices within a centralized assignment problem. Rather, they describe social, legal, or economic relations that shape endowments and exclusion rights, as described below.

Our use of priorities in a decentralized economy bears some similarity Piccione and Rubinstein's (2007) strength relation in their model of a "jungle economy." Their strength relation

[17] Ehlers and Erdil (2010) model some of these situations with non-strict partial orders. Priorities have a different interpretation in their analysis than in ours.

[18] A student would feel justified envy if he prefers to attend a school that enrolled a lower-priority student.

is a linear order of all agents, while priorities in our model are good-specific and possibly incomplete. Furthermore, the link between priorities and agents' rights in a relational economy is mediated through conditional endowments, which we turn to next.

3.2 Conditional Endowments

An endowment system specifies a distribution of exclusion rights. In a relational economy, it must address two requirements, both illustrated by the Diarchy. First, exclusion rights are often qualified by the prevailing allocation. Thus, the endowment system should adjust accordingly. We adopt the term *conditional endowment system* to emphasize this conditionality. And second, conditional endowments should reflect the context conveyed by the priority structure. Intuitively, if $i \rhd_h j$, then i should enjoy rights no less than j with respect to house h.[19] We investigate two approaches formalizing both desiderata.

Weak Conditional Endowments

A natural starting point places a house in a coalition's conditional endowment if one of its members dominates that house's assigned occupant. A relational economy's *weak conditional endowment system at* μ, $\omega_\mu \colon 2^I \to 2^H$, is an endowment system defined as follows. For every house $h \in H$ and coalition $C \subseteq I$, $h \in \omega_\mu(C)$ if and only if there exists an agent $i \in C$ such that $i \trianglerighteq_h \mu^{-1}(h)$. Weak conditional endowments plug-in seamlessly into the definition of exclusion blocking, without otherwise changing its behavioral rationale. The following definition parallels Definition 5, with "ω_μ" replacing "ω" in the second point.

Definition 6. A non-empty coalition $C \subseteq I$ can *indirectly exclusion block the allocation* μ *with allocation* σ *given* ω_μ if

1. $\sigma(i) \succ_i \mu(i)$ for all $i \in C$; and,
2. $\mu(j) \succ_j \sigma(j) \implies \mu(j) \in \Omega(C|\omega_\mu, \mu)$.

The *strong exclusion core* of a relational economy is the set of allocations that cannot be indirectly exclusion blocked given ω_μ.

The intuitive derivation of ω_μ gives the strong exclusion core great appeal. Regrettably, it is easy to see that the strong exclusion core can be empty.

[19] See Campbell (1992) for a discussion of how hierarchical relations may qualify property rights.

Example 5. Suppose $I = \{i, j, k\}$ and $H = \{h_1, h_2\}$. Assume that $i \rhd_{h_1} j \rhd_{h_1} k$, $k \rhd_{h_2} j \rhd_{h_2} i$ and

$$\succ_i : h_2, h_1, h_0 \qquad \succ_j : h_1, h_0 \qquad \succ_k : h_1, h_2, h_0 \ .$$

Any assignment μ where $\mu(j) = h_1$ can be exclusion blocked by either i or k. But, if $\mu(j) = h_0$, efficiency demands that $\mu(i) = h_2$ and $\mu(k) = h_1$. This assignment can be exclusion blocked at ω_μ by j.

The root of the preceding example's problem is the economy's cyclic priority structure. Whether encountered in a consumer's preference or in committee voting, cyclic relations are a well-known challenge for economic analysis. The simplest strategy to address this complication is to restrict the priority structure accordingly. A priority structure $\rhd$ is *acyclic* if for all $h \in H$ and agents i, j, and k,

$$i \rhd_h j \ \& \ i \not\trianglerighteq_h k \implies k \rhd_{h'} j \qquad \forall h' \neq h, h_0. \tag{1}$$

Our definition of acyclicity is specifically phrased to accommodate incomplete priority structures; however, (1) reduces to the more-familiar Ergin (2002) acylicity when each $\rhd_h$ is a linear order of all agents.[20] It is related to strong acyclicity, which was proposed by Ehlers and Erdil (2010) as an extension Ergin's (2002) definition. Stronger forms of acyclicity are also proposed by Kesten (2006). We discuss the following theorem's proof in Section 3.3.

Theorem 2. *For any relational economy with an acyclic priority structure, the strong exclusion core is not empty.*

While acyclicity appears to be a demanding requirement, it is satisfied by many common situations. If house h is privately owned there is an agent i—the house's owner—such that $i \rhd_h j$ for all $j \neq i$ and $j \not\rhd_h k$ for all $j, k \in I \setminus \{i\}$, as in Figure 3(b). Conversely, if house h is part of the social endowment, $i \not\rhd_h j$ for all $i, j \in I$, as in Figure 3(a). Any economy featuring a combination of houses that are privately owned or belong to the social endowment has an acyclic priority structure.[21] Even the Diarchy of Example 4 can be modeled with an acyclic priority structure: $j \rhd_h i$ and $k \rhd_h i$ for all $h \in H$. In this case, the strong exclusion core coincides with the two focal allocations where the kings split the houses among themselves.

[20] An Ergin (2002) cycle occurs if for distinct houses h and h' and distinct agents i, j, and k, $k \rhd_h i \rhd_h j \rhd_{h'} k$. A priority structure is Ergin (2002) acyclic if it does not contain an Ergin (2002) cycle. If $\rhd_h$ is a linear order for each h, then (1) becomes $k \rhd_h i \rhd_h j \implies k \rhd_{h'} j$. Hence, an Ergin (2002) cycle cannot occur.

[21] To confirm this fact, note that the antecedent in (1), $i \rhd_h j$ & $i \not\trianglerighteq_h k$, is never satisfied.

Strong Conditional Endowments

In economies with cyclic priority structures, a strengthening of the preceding definitions offers a route to positive results. The proposed strengthening ensures the derived conditional endowment system does not inherit the priority structure's potentially problematic cycles. A relational economy's *strong conditional endowment system at* μ, $\omega^*_\mu \colon 2^I \to 2^H$, is an endowment system defined as follows. For every house $h \in H$ and coalition $C \subseteq I$, $h \in \omega^*_\mu(C)$ if and only if for every sequence of distinct houses $(h^0, h^1, \ldots, h^{k-1}) \ni h$ such that

$$\mu^{-1}(h^0) \ntrianglerighteq_{h^0} \mu^{-1}(h^1) \ntrianglerighteq_{h^1} \cdots \ntrianglerighteq_{h^{k-2}} \mu^{-1}(h^{k-1}) \ntrianglerighteq_{h^{k-1}} \mu^{-1}(h^0), \tag{2}$$

there exists $i \in C$ such that $i \trianglerighteq_{h^{\ell-1 \pmod{k}}} \mu^{-1}(h^\ell)$ for some $\ell \in \{0, \ldots, k-1\}$.

If house h belongs to a coalition's strong conditional endowment, the coalition must $\trianglerighteq$-dominate other agents who *may* directly or indirectly acquire rights to house h at μ. If $h \in \omega^*_\mu(C)$, then there exists an agent $i \in C$ such that $i \trianglerighteq_h \mu^{-1}(h)$.[22] Thus, $\omega^*_\mu(C) \subseteq \omega_\mu(C)$ for all $C \subseteq I$. Going further, coalition C must command a dominating position against agents whose claim to h is more indirect. This fact is easiest to see in the special case where each $\triangleright_h$ is a linear order. In this case, (2) becomes

$$\mu^{-1}(h^0) \triangleright_{h^{k-1}} \mu^{-1}(h^{k-1}) \triangleright_{h^{k-2}} \cdots \triangleright_{h^1} \mu^{-1}(h^1) \triangleright_{h^0} \mu^{-1}(h^0),$$

which is a "cycle" formed by successive comparisons of the form $\mu^{-1}(h^\ell) \triangleright_{h^{\ell-1}} \mu^{-1}(h^{\ell-1})$. When $h \in \omega^*_\mu(C)$, at least one $\triangleright_{h^{\ell-1}}$-dominating agent in this sequence is further dominated by some $i \in C$, i.e. $i \trianglerighteq_{h^{\ell-1}} \mu^{-1}(h^\ell) \triangleright_{h^{\ell-1}} \mu^{-1}(h^{\ell-1})$.

Replacing ω_μ with ω^*_μ in Definition 6 leads to a corresponding version of the exclusion core. The *weak exclusion core* of a relational economy is the set of allocations that cannot be indirectly exclusion blocked given ω^*_μ.

Lemma 2. *(a) The strong exclusion core is a subset of the weak exclusion core. (b) If the priority structure is acyclic, the weak and strong exclusion cores coincide.*

In Example 5, the priority structure was not acyclic and the strong exclusion core was empty. The unique weak exclusion core allocation assigns i to h_2 and k to h_1. More generally, the weak exclusion core identifies plausible final allocations even if the strong exclusion core offers little guidance.

Theorem 3. *For any relational economy, the weak exclusion core is not empty.*

[22] Consider the case where the sequence referenced in the definition has length $k = 1$.

3.3 Generalized Top Trading Cycles

We prove Theorems 2 and 3 in Appendix A. Our proofs are constructive and rely on an algorithm introduced below. The algorithm identifies a strong exclusion core assignment when the priority structure is acyclic (Theorem 2); otherwise, its output belongs to the weak exclusion core (Theorem 3). Necessarily, this algorithm builds upon several precursors given the exclusion core's coincidence with certain assignments in benchmark cases. First, the TTC algorithm identifies the unique exclusion core assignment in Shapley and Scarf's (1974) economy. Similarly, a Pareto efficient allocation in Hylland and Zeckhauser's (1979) market can be identified with a serial dictatorship.[23] A mechanism that nests both the TTC algorithm and the serial dictatorship is the "You Request My House—I Get Your Turn" (YRMH-IGYT) mechanism of Abdulkadiroğlu and Sönmez (1999). This mechanisms was proposed to solve the house allocation problem with existing tenants, where some houses have an owner and others belong only to the social endowment. Our mechanism builds on the "TTC variant" of the YRMH-IGYT mechanism, as presented by Abdulkadiroğlu and Sönmez (1999) and Sönmez and Ünver (2010). It is is also an immediate descendant of the TTC algorithm as applied to the school-choice problem with the tie-breaking of coarse priorities (Abdulkadiroğlu and Sönmez, 2003; Abdulkadiroğlu et al., 2009).[24] Our algorithm reduces to each of the above cases when the environment is appropriately restricted. This connection is encouraging and binds the exclusion core to well-known economic models. Consequently, our analysis adds a new justification for the use of trading cycle procedures in applications.

Algorithm 2 (Generalized Top Trading Cycles (GTTC))**.** Given $\langle I, H, \succ, \rhd \rangle$, let $\tilde{\rhd}_h$ be a complete linear order of agents such that $i \rhd_h j \implies i \tilde{\rhd}_h j$ for each $h \in H$. We call $\tilde{\rhd}$ a completion of $\rhd$.[25] Let $I^1 := I$ and $H^1 := H$. In step $t \geq 1$ the algorithm proceeds as follows with inputs I^t and H^t.

Step t. Let I^t and H^t be the sets of unassigned agents and houses, respectively, at step t. Construct a directed graph a follows. The set of vertices is $I^t \cup H^t \cup \{h_0\}$. Draw an arc from $i \in I^t$ to $h \in H^t \cup \{h_0\}$ if and only if h is agent i's most preferred house among those in

[23] In a serial dictatorship all agents are ordered. The first agent is assigned his most-preferred object. The second agent is assigned his most-preferred object from those remaining. And so on. The resulting assignment is Pareto efficient if preferences are strict.

[24] The main difference is our algorithm's accommodation of a more general class of priority structures than typically encountered in school-choice problems.

[25] It is tempting to regard $\tilde{\rhd}$ as a priority structure supplemented by a tie breaking rule (Ehlers, 2014). While formally compatible with our model, we hesitate to emphasize this interpretation. If $i \not\rhd_h j$ and $j \not\rhd_h i$, then i and j are not $\rhd_h$-comparable but do not necessarily have "equal claim" to house h.

$H^t \cup \{h_0\}$. For each $h \in H^t$, draw an arc from h to the $\tilde{\rhd}_h$-maximal agent in I^t.

(a) If there exists an agent i who is pointing to h_0, assign him to the outside option, i.e. set $\mu(i) = h_0$, and remove him from the market. Set $\tilde{I}^t = \{i\}$ and $\tilde{H}^t = \varnothing$.

(b) Otherwise, the constructed graph contains at least one cycle. Choose any cycle and carry out the implied assignments. That is, if $i \to h$ in the cycle then set $\mu(i) = h$. Remove the associated agents, $\tilde{I}^t \subseteq I^t$, and their assigned houses, $\tilde{H}^t \subseteq H^t$, from the market.

Define $I^{t+1} := I^t \setminus \tilde{I}^t$ and $H^{t+1} := H^t \setminus \tilde{H}^t$.

The above process continues until $I^t = \varnothing$. Any remaining houses are left unassigned.

An example in Appendix B illustrates Algorithm 2's step-by-step operation. As there is a finite number of agents and at least one agent is removed from the market in each step, the algorithm terminates in a finite number of steps.

For a given economy, the Algorithm 2 is parameterized by the employed completion $\tilde{\rhd}$. By varying this completion, we can identify a family of exclusion core outcomes.

Theorem 4. *Every strong exclusion core allocation in the relational economy $\langle I, H, \succ, \rhd \rangle$ can be identified by the GTTC algorithm with some completion $\tilde{\rhd}$ of $\rhd$.*

While Algorithm 2 can find all strong exclusion core allocations, it cannot find all weak exclusion core allocations.[26] The next corollary follows from Lemma 2 and Theorems 2–4.

Corollary 1. *Denote the weak exclusion core by WEC, the strong exclusion core by SEC, and the range (over all completions of the priority structure) of Algorithm 2 by $GTTC$.*

(a) Given an arbitrary priority structure, $SEC \subseteq GTTC \subseteq WEC$.

(b) If the economy's priority structure is acyclic, $SEC = GTTC = WEC$.

Since weak exclusion core allocations are Pareto efficient, all assignments identified by Algorithm 2 also have this property. Additionally, the algorithm is strategy-proof. No agent can improve his assignment by strategically misreporting his preference. This fact is a direct

[26] Consider the following example with three agents and three houses. The agents' preferences are $\succ_i: h_1, h_2, h_3$, $\succ_j: h_1, h_2, h_3$, and $\succ_k: h_1, h_3, h_2$. Agents j and k jointly own h_1 in the sense that $j \rhd_{h_1} i$ and $k \rhd_{h_1} i$. Agent i is the sole owner of h_2 and h_3: $i \rhd_h j$ and $i \rhd_h k$ for $h \in \{h_2, h_3\}$. Otherwise, the agents are not $\rhd$.-comparable. There are three weak exclusion core allocations. In two allocations, i always takes h_2 and either j or k claims h_1. In the third allocation—$\nu(i) = h_1$, $\nu(j) = h_2$, and $\nu(k) = h_3$—agents j and k trade away h_1 to i in exchange for h_2 and h_3. The assignment ν cannot be identified by Algorithm 2.

corollary to prior results derived by Roth (1982), Abdulkadiroğlu and Sönmez (1999), Roth et al. (2004), and (in particular) Abdulkadiroğlu and Sönmez (2003), who extend Gale's TTC algorithm to an assignment problem with priorities. Though priorities have a different meaning in our model, the argument is essentially identical and we omit the proof.

Drawing on the preceding analysis, we can highlight several further properties of the exclusion core. First, the strong exclusion core is stable in the sense of von Neumann and Morgenstern (1944) when the priority structure is acyclic.[27,28] Thus, echoing their interpretation, it defines a consistent "standard of behavior." Trivially, the strong exclusion core is internally stable since its outcomes cannot be indirectly exclusion blocked. Its external stability is confirmed by the following proposition.

Proposition 4. *Consider a relational economy with an acyclic priority structure. If μ is not a strong exclusion core allocation, there exists a coalition C that can indirectly exclusion block μ given ω_μ with some strong exclusion core allocation σ.*

Given Corollary 1, the preceding stability properties also apply to the weak exclusion core.

Second, the exclusion core exhibits intuitive comparative statics with respect to $\rhd$. Changes in $\rhd$ may reflect changing legal or social norms. Given I and H, the priority structure $\rhd'$ is a *coarsening* of $\rhd$ if for all $h \in H$ and $i, j \in I$, $i \rhd'_h j \implies i \rhd_h j$. Intuitively, $\rhd'$ coincides with $\rhd$ except some hierarchal relations among agents are possibly expunged.

Proposition 5. *If $\rhd'$ is a coarsening of $\rhd$, the strong (weak) exclusion core of $\mathcal{E}' = \langle I, H, \succ, \rhd' \rangle$ contains the strong (weak) exclusion core of $\mathcal{E} = \langle I, H, \succ, \rhd \rangle$.*

Agent-level implications can be derived too. We say that the set of allocations A $\succ_i$-*dominates* the set A' if for every $\mu \in A$ there exists $\mu' \in A'$ such that $\mu(i) \succeq_i \mu'(i)$ and for every $\mu' \in A'$ there exists a $\mu \in A$ such that $\mu(i) \succeq_i \mu'(i)$. The next proposition formalizes how Algorithm 2 respects a "priority improvement" for agent i (Balinski and Sönmez, 1999). Given Corollary 1, it also applies to the exclusion core when the priority structure is acyclic.

Proposition 6. *Let $\mathcal{E} = \langle I, H, \succ, \rhd \rangle$ and $\mathcal{E}' = \langle I, H, \succ, \rhd' \rangle$ be two economies where (i) for all $h \in H$ and $k \neq i$, $k \not\rhd_h i \implies k \not\rhd'_h i$; and, (ii) for all $h \in H$ and $k, j \neq i$,*

[27] A set of outcomes A is *von Neumann-Morgenstern (vNM) stable* if it is (i) internally stable: every $\mu \in A$ is not "dominated" by any $\sigma \in A$; and, (ii) externally stable: every $\mu \notin A$ is "dominated" by some $\sigma \in A$. In our context, the allocation σ "dominates" μ if there exists a coalition that can indirectly exclusion block μ with σ given ω_μ (Definition 6). Different definitions of "dominance" lead to different stable sets.

[28] In fact, a stronger conclusion follows. When indirect exclusion blocking (Definition 6) defines the dominance relation, the strong exclusion core is the *unique* vNM stable set. Other dominance relations can lead to multiple disjoint stable sets. For example, if weak blocking (Definition 1) defines the dominance relation, the Kingdom economy (Example 1) has two disjoint vNM stable sets.

$k \not\triangleright_h j \iff k \not\triangleright'_h j$.[29] *The set of allocations identified by the GTTC algorithm in $\mathcal{E}$ $\succ_i$-dominates the corresponding set in $\mathcal{E}'$.*

3.4 Conditional and Unconditional Endowments

We introduced relational economies to better model joint or qualified ownership. We conclude by linking relational economies to the setting of Section 2 where an endowment system ω, rather than a priority structure $\triangleright$, is the economic primitive.

First, consider the relational economy $\langle I, H, \succ, \triangleright \rangle$. We define its *unconditional endowment system*, $\omega^{**}\colon 2^I \to 2^H$, as follows. For all $h \in H$ and $C \subseteq I$, $h \in \omega^{**}(C)$ if and only if for every allocation μ there exists some $i \in C$ such that $i \trianglerighteq_h \mu^{-1}(h)$.[30] By replacing ω_μ with ω^{**} in Definition 6, we can define the *unconditional exclusion core* of a relational economy as the set of allocations that cannot be indirectly exclusion blocked given ω^{**}. It can be shown that $\omega^{**}(C) \subseteq \omega^*_\mu(C) \subseteq \omega_\mu(C)$ for all μ and C. Thus, a relational economy's unconditional exclusion core is not empty and contains its weak and strong exclusion cores.

Now, recalling Section 2, consider the simple economy $\langle I, H, \succ, \omega \rangle$ with an endowment system ω satisfying (A1)–(A4). The priority structure $\triangleright$ *represents* ω if for each $h \in H$, $i \triangleright_h j$ if and only if $i \in C^h$ and $j \notin C^h$. Theorem 1 is a corollary to the next lemma.

Lemma 3. *Let $\langle I, H, \succ, \omega \rangle$ be an economy with an endowment system ω satisfying (A1)–(A4). Suppose $\triangleright$ represents ω.*

(a) The exclusion core of the simple economy $\langle I, H, \succ, \omega \rangle$ coincides with the unconditional exclusion core of the relational economy $\langle I, H, \succ, \triangleright \rangle$.

(b) If $\langle I, H, \succ, \omega \rangle$ is a simple economy where every house is either privately owned or part of the social endowment, then its exclusion core coincides with the strong, weak, and unconditional exclusion cores of the relational economy $\langle I, H, \succ, \triangleright \rangle$.

Lemma 3 lets us revisit the cases of private and public ownership introduced in Section 2.4. The following result, due to Sönmez (1999), points to a closer connection between the exclusion core and the strong core in a private ownership economy.

Theorem 5 (Sönmez (1999))**.** *Suppose there exists a Pareto efficient, individually rational, and strategy-proof mechanism f in a private-ownership economy.*

[29]Conditions (i) and (ii) adapt to our context Balinski and Sönmez's (1999) definition of a priority improvement for agent i.

[30]Equivalently, $h \in \omega^{**}(C)$ if and only if C includes every $\triangleright_h$-maximal agent in the economy.

(a) The strong core of the economy is either empty or a singleton.

(b) If the strong core is not empty, the unique element of the strong core is identified by f.

Noting that the GTTC algorithm satisfies the conditions of Theorem 5, two corollaries follow.

Corollary 2. *In a private ownership economy, if the exclusion core contains more than one allocation, the strong core is empty.*

Corollary 3. *In a private ownership economy, the exclusion core equals the strong core whenever the latter is not empty.*

Corollary 3 also implies the coincidence of the exclusion core and the strong core in Shapley and Scarf's (1974) economy.

An economy with both private and social endowments is the house allocation problem with existing tenants, introduced by Abdulkadiroğlu and Sönmez (1999). In this problem, every house either belongs to the social endowment or is owned by exactly one agent. No agent owns more than one house. Such an economy's exclusion core may differ from its strong core (see Example 3). In this setting, the GTTC algorithm reduces to Abdulkadiroğlu and Sönmez's (1999) YRMH-IGYT mechanism. The following is a corollary to Theorem 4.

Corollary 4. *In the house allocation problem with existing tenants, the exclusion core coincides with the set of all possible allocations identified by Abdulkadiroğlu and Sönmez's (1999) YRMH-IGYT mechanism.*

Corollary 4 provides a new characterization of the YRMH-IGYT mechanism, complementing the axiomatization proposed by Sönmez and Ünver (2010).

Finally, the GTTC mechanism reduces to a serial dictatorship when all houses belong only to the social endowment. Thus, the exclusion core coincides with the set of Pareto efficient allocations, as shown directly by Proposition 3 above.

4 Related Literature

Our analysis bridges two previously segregated literatures. First, we contribute to the study of discrete exchange economies. And second, we complement scholarship in law and economics on the nature of property. We address each domain in turn.

Trading Cycles and Discrete Exchange Economies

Shapley and Scarf (1974) were the first to study the core of a discrete exchange economy and they introduced David Gale's TTC algorithm, which Algorithm 2 generalizes. Formally, the algorithm belongs to the class of "hierarchical exchange" mechanisms introduced by Pápai (2000). Such mechanisms rely on an alternative definition of endowments, termed inheritance trees. Pycia and Ünver (2017) introduce inheritance structures as part of their generalization of Pápai's (2000) model. An inheritance structure defines how unassigned houses are inherited or transferred during a multi-step assignment process, like the TTC procedure. For example, if agent i owns house h and exits the market with house h', the inheritance structure specifies the new owner of house h who may exchange it in the continuation of the trading process. Svensson and Larsson (2005) introduce endowment rules, which are similar.

The analogues of an inheritance structure in our model are the completions "$\tilde{\trianglerighteq}_h$" operating within Algorithm 2. We regard these completions as purely technical devices and they should not be conflated with endowments or property rights in our analysis. An inheritance structure is a specification of contingent control within a sequential assignment procedure. In contrast, an endowment system in our model describes a distribution of exclusion rights in general and is logically independent of any particular trading protocol.

Two recent studies draw on the dynamics of a hierarchical exchange to propose new variants of the core. Both combine weak blocking (Definition 1) with alternative definitions of endowments. Ekici (2013) calls an allocation reclaim proof if it cannot be weakly blocked by any coalition whose endowment is a combination of the pre-trade endowment and the ex post allocation. Starting with Pápai's (2000) model, Tang and Zhang (2016) define an agent's contingent endowment at allocation μ as the maximal set of houses he would have feasibly inherited during trade leading to μ given the prevailing inheritance structure. In contrast to these studies, our definitions overcome the problematic incentives underlying weak blocking. Our derivations and interpretations of endowments are distinct as well.

Many variants of the Shapley and Scarf (1974) economy have been considered. Konishi et al. (2001) show that the weak core may be empty if agents can consume multiple goods. As the exclusion core is a subset of the weak core, we cannot offer new positive results for this class of problems. To limit confounds, we have assumed a strict preference domain, integral endowments, and deterministic final outcomes. Each of these assumptions has been relaxed by many authors.[31] Farsighted solutions have also been considered (Klaus et al.,

[31] For richer preference domains, see Alcalde-Unzu and Molis (2011), Jaramillo and Manjunath (2012), and Saban and Sethuraman (2013). Kesten (2009), Athanassoglou and Sethuraman (2011), and Aziz (2015) study

2010). We defer investigating these extensions to future research.

Competitive or Walrasian equilibrium is another prominent solution often applied to exchange economies. The (possible) absence of personal endowments precludes the meaningful application of standard price equilibrium definitions to our setting, but generalized equilibrium notions are applicable. Richter and Rubinstein (2015) have recently introduced the notion of a "primitive equilibrium." Their definition does not rely on budget sets and it deliberately eschews endowments. Instead, they observe that equilibria in exchange economies induce an ordering of goods, from more to less desirable. Exclusion core allocations identified by the GTTC algorithm, which orders goods based on the step in which the good is assigned, satisfy Richter and Rubinstein's (2015) equilibrium definition.

Endowments and Property Rights

The interpretation of endowments and property that we advance is narrow. It is derived from a basic principle, the right to exclude others. We cannot hope to account for this principle's philosophical, historical, and legal development here. Penner (1997), Merrill (1998), and Merrill and Smith (2001b), among many others, elaborate on these points in detail. Klick and Parchomovsky (2017) document the importance of the right to exclude for land values. Its impact on patent and intellectual property law is undeniable (Mossoff, 2009).

We have sidestepped the familiar "bundle of rights" interpretation of property, which is central in Coase's (1960) analysis and in subsequent research emphasizing investment incentives and control rights (Merrill and Smith, 2001b; Segal and Whinston, 2012). Instead, our analysis suggests that efficient allocations depend on a small subset of the rights one may associate with property. The right to exclude is the only stick necessary to ensure efficient outcomes in our model. We have also set aside all implications associated with the allocation of property rights. Formally, our model inhabits Coase's hypothetical setting without transaction costs. Reassuringly, all exclusion core outcomes are efficient.

To simplify exposition, we split our analysis into two parts. Section 2 examined simple economies with exogenous endowments; Section 3 introduced relational economies and conditional endowments. Importantly, this division also reflects a debate among legal scholars concerning the nature of property. This debate is often framed as a spectrum running between *in rem* and *in personam* paradigms.[32] Section 2 presumes an *in rem* interpretation of property. An agent's ownership rights are shaped by his relationship with a specific thing

fractional endowments. Abdulkadiroğlu and Sönmez (1998) and Carroll (2014) examine random outcomes.

[32]Merrill and Smith (2001b) examine this distinction with particular reference to economic analysis. See also Campbell (1992), Merrill (1998), and Klick and Parchomovsky (2017) and the citations therein.

and are universal in character. The Kingdom (Example 1) and Shapley and Scarf's house market operate naturally under this paradigm. Arruñada (2012) argues that *in rem* rights are essential for impersonal exchange to be possible.

Conditional endowment systems formalize an *in personam* interpretation of property and ownership. An agent's rights are defined by his relationships with others, which we model with priorities. Exclusion rights are constrained by the priority structure and qualified by the prevailing assignment. Situations with overlapping claims or conditionality, such as the Diarchy (Example 4) or cadaveric organ transplantation, follow this archetype. An agent's ability to participate in these markets anonymously is limited, and often impossible.

Our model provides a new formal setting for the comparison of *in rem* and *in personam* paradigms. In fact, we provide an embedding of the former into the latter in Section 3.4 (Lemma 3). Still, our model's sparsity masks some notable differences. Merrill and Smith (2001a), for example, observe that *in rem* rights have a lower informational burden than their *in personam* counterparts. This consideration is absent from our model but affects a market's operation and scalability in practice. Its further investigation is likely to be fruitful.

5 Concluding Remarks

Property plays a pivotal role in markets; however, its interpretation and relation to endowments within economic analysis has been taken for granted. Drawing on a simple principle, the right to exclude others, we propose a new solution for exchange economies and assignment problems. The exclusion core rests on an interpretation of endowments as a distribution of exclusion rights, which may be shared, sensitive to competing claims, or qualified by relationships. Classic solutions, such as the strong and weak cores, cannot be readily applied in these cases; they miss the significance of exclusion for exchange. Our solution is compatible with recognized property paradigms and highlights subtle, though powerful, incentives that sustain efficient outcomes.

We have presented our theory in a simple setting and we hope that this exposition does not mask our analysis' broader conceptual message. The logic of indirect exclusion and the role of repossession chains generalize. Nevertheless, even the simple model we have analyzed is of great practical importance since it serves as a foundation for many market design applications. Examples include the allocation of college dormitories, the exchange of transplant organs, the allocation of airport arrival slots, and the assignment of students to schools. The exclusion core's tight connection with the TTC algorithm is a new justification for the latter's use in these and related applications.

Favoring brevity, we have deferred investigating many extensions. Notably, we have focused on deterministic allocations; "fair" outcomes may require randomization. It may also prove interesting to extend our analysis to the case of club goods, which are excludable but non-rivalrous. Finally, our analysis separates a good's use and ownership, which are sometimes conflated in models of exchange. Use dictates welfare while acquired ownership or exclusion rights govern outside options, as the difference between a tenant and a landlord illustrates. This distinction is immaterial in a one-period model but is salient in a dynamic setting, which is another promising extension worth pursuing.

A Proofs

Proof of Lemma 1. Let μ be a direct exclusion core allocation. To derive a contradiction, suppose μ can be strongly blocked by coalition C with allocation σ. Clearly, $\mu(i) \succeq_i h$ for all $h \in \omega(i) \cup \{h_0\}$ for all i. (Else, μ can be directly exclusion blocked by a single agent.) Thus, $\sigma(i) \succ_i \mu(i)$ for all $i \in C$ and $\sigma(C) \subseteq \omega(C)$. Let

$$\hat{\sigma}(i) = \begin{cases} \sigma(i) & \text{if } i \in C \\ h_0 & \text{if } i \notin C \ \& \ \mu(i) \in \sigma(C) \\ \mu(i) & \text{otherwise} \end{cases}.$$

Observe that $\hat{\sigma}(i) \succ_i \mu(i)$ for all $i \in C$. Moreover, if $\mu(i) \succ_i \hat{\sigma}(i)$, then $\hat{\sigma}(i) = h_0$. Thus, $\mu(i) \in \sigma(C) \subseteq \omega(C)$. Hence, C can directly exclusion block μ with $\hat{\sigma}$—a contradiction. □

Proof of Proposition 1. Let μ be a strong core allocation. Assume toward a contradiction that μ can be indirectly exclusion blocked by $C \subseteq I$ with σ. Thus, $\sigma(i) \succ_i \mu(i)$ for all $i \in C$. Moreover, if $i \in C$, then $\sigma(i) = \mu(j) \in H$ for some $j \in I$. (If $\sigma(i)$ was not occupied or $\sigma(i) = h_0$, the allocation μ would not be Pareto efficient, a contradiction.)

To derive a contradiction, we will identify a coalition K that will be able to strongly block μ. Start with any $i^0 \in C$ and define a sequence $(i^0, i^1, \ldots)$ of agents as follows:

1. If $i^\ell \in C$, there exists $h \in H$ such that $\sigma(i^\ell) = h$. Set $i^{\ell+1} = \omega^{-1}(h)$.

2. If $i^\ell \notin C$, then $\mu(i^\ell) \in H \cup \{h_0\}$. If $\mu(i^\ell) = h_0$, set $i^{\ell+1} = i^0$; otherwise, if $\mu(i^\ell) = h \in H$, set $i^{\ell+1} = \omega^{-1}(h)$.

As there is a finite number of agents, the defined sequence must eventually cycle. Without loss of generality and relabeling if necessary, let $K := (i^0, \ldots, i^k)$ be the cycle. Next we show that coalition K can strongly block μ with the allocation

$$\hat{\sigma}(i) = \begin{cases} \sigma(i) & \text{if } i \in K \cap C \\ \mu(i) & \text{if } i \in K \setminus C \\ h_0 & \text{otherwise} \end{cases}.$$

By construction, $\hat{\sigma}(i) \succeq_i \mu(i)$ for all $i \in K$ and $\hat{\sigma}(i^0) \succ_{i^0} \mu(i^0)$. Moreover, for each agent i^ℓ, $\ell < k$, $\hat{\sigma}(i^\ell) \in \omega(i^{\ell+1}) \in \omega(K)$. And for agent i^k, there either exists some $i^\ell \in \{i^0, \ldots, i^{k-1}\}$ such that $\hat{\sigma}(i^k) \in \omega(i^\ell)$ or $\hat{\sigma}(i^k) = \mu(i^k) = h_0$. Thus, $\hat{\sigma}(K) \subseteq \omega(K) \cup \{h_0\}$. Therefore, coalition K can strongly block μ, which is a contradiction. □

Remark A.1. When each house has at most one owner, it is simple to verify that $\Omega(C|\omega,\mu) = \bigcup_{k=0}^{\infty}(\omega \circ \mu^{-1})^k(\omega(C))$. We use this simplified expression in the following proof.

Proof of Proposition 2. The TTC assignment is this economy's unique strong core allocation. By Proposition 1 it belongs to the exclusion core. Conversely, suppose μ is an exclusion core allocation. The allocation μ can be represented as a directed graph where each house $h_k \in H$ points to its owner, say $h_k \to \omega^{-1}(h_k) = i_k$, and each agent points to his assignment, i.e. $i_k \to \mu(i_k)$. As all houses are acceptable and $|I| = |H|$, $\mu(I) = H$. The resulting graph partitions the set of agents and houses into disjoint cycles $\{K_1, \ldots, K_T\}$. Observe that $i \in K_t \iff \omega(i) \in K_t \iff \mu(i) \in K_t$. Hence, if $i \in K_t$ and $h \in K_t$, $h \in \bigcup_{k=0}^{\infty}(\omega \circ \mu^{-1})^k(\omega(i))$.

Suppose coalition $C = \{i'_1, \ldots, i'_k\}$ can weakly block μ with σ. Thus, $\sigma(i) \succeq_i \mu(i)$ for all $i \in C$, $\sigma(i) \succ_i \mu(i)$ for some $i \in C$, and $\sigma(C) \subseteq \omega(C) \cup \{h_0\}$. Clearly, the final condition can be strengthened to $\sigma(C) = \omega(C)$. Furthermore, without loss of generality we may assume that $\sigma(\cdot)$ assigns the houses in $\omega(C)$ cyclicly among the members of C. That is,

$$h'_1 \to i'_1 \to \cdots \to h'_k \to i'_k \to h'_1 \tag{A.1}$$

where $h'_\ell = \omega(i'_\ell)$ and $\sigma(i'_\ell) = h'_{\ell+1}$. for $\ell \leq k-1$ and $\sigma(i'_k) = h'_1$.[33] Let $W = \{i \in C \mid \sigma(i) \succ_i$

[33]If σ induces multiple cycles among agents in C, they are necessarily disjoint and without loss of generality we may focus on any one of them involving an agent $i \in C$ such that $\sigma(i) \succ_i \mu(i)$.

$\mu(i)\}$ and define

$$\hat{\sigma}(i) = \begin{cases} \sigma(i) & \text{if } i \in W \\ \mu(i) & \text{if } i \notin W \text{ \& } \mu(i) \notin \sigma(W) \\ h_0 & \text{otherwise} \end{cases} .$$

Clearly, $\hat{\sigma}(i) \succ_i \mu(i)$ for all $i \in W$. Next, pick any j that $\mu(j) \succ_j \hat{\sigma}(j)$. There exists some cycle K_j, as defined above, such that $j \in K_j$ and $\mu(j) \in K_j$. Furthermore, there exists some $i \in W$ such that $\hat{\sigma}(i) = \mu(j)$. Without loss of generality, suppose $i = i'_1$, according to the enumeration in (A.1). There are two cases. First, if $i'_1 \in K_j$, then $\mu(j) \in \bigcup_{k=0}^{\infty}(\omega \circ \mu^{-1})^k(\omega(i'_1))$. Alternatively, and second, if $i'_1 \notin K_j$, then there exists some $i'_t \in W \subseteq C$ such that $i'_t \in K_j$. (If this was not the case, then for all $i \in (i'_2, \ldots, i'_k)$, $\hat{\sigma}(i) = \mu(i)$. This implies $\mu(i'_k) = h'_1 \in K_j$. But, $h'_1 = \omega(i'_1)$ and hence $i'_1 \in K_j$, which is a contradiction.) Together, the preceding cases imply that $\mu(j) \in \bigcup_{k=0}^{\infty}(\omega \circ \mu^{-1})^k(\omega(W))$. As the choice of j was arbitrary, we conclude that coalition W can indirectly exclusion block μ with $\hat{\sigma}$, which contradicts μ being an exclusion core allocation. □

Proof of Proposition 3. As exclusion core allocations are Pareto optimal, it is sufficient to show that no Pareto optimal allocation μ can be indirectly exclusion blocked. Suppose the contrary. If coalition C can indirectly exclusion block μ with σ, there exists an agent j who is harmed by the reallocation. Thus, $\mu(j) \succ_j \sigma(j)$ and $j \notin C$. Indirect exclusion blocking implies that $\mu(j) \in \Omega(C|\omega, \mu)$. Necessarily, $C \subsetneq I$, which implies $\omega(C) = \varnothing$, $\omega(C \cup (\mu^{-1} \circ \omega)(C)) = \varnothing$, and so on. But then $\Omega(C|\omega, \mu) = \varnothing$—a contradiction. □

Proof of Lemma 2. Part (a) is immediate since $\omega^*_\mu(C) \subseteq \omega_\mu(C)$ for all $C \subseteq I$. To confirm part (b), it is sufficient to verify that $\omega_\mu(C) \subseteq \omega^*_\mu(C)$ when $\rhd$ is acyclic. Let $h \in \omega_\mu(C)$. To show that $h \in \omega^*_\mu(C)$, we will confirm that the following statement is true for all $k \geq 1$.

> ($\star$) For every (finite) sequence of distinct houses $(h^0, \ldots, h^{k-1}) \ni h$ such that $\mu^{-1}(h^0) \not\trianglerighteq_{h^0} \cdots \not\trianglerighteq_{h^{k-2}} \mu^{-1}(h^{k-1}) \not\trianglerighteq_{h^{k-1}} \mu^{-1}(h^0)$ there exists $i \in C$ such that $i \trianglerighteq_{h^{\ell-1 \pmod k}} \mu^{-1}(h^\ell)$ for some $\ell \in \{0, \ldots, k-1\}$.

Suppose $k = 1$. In this case, $h^0 = h$ and $0 - 1 \pmod 1 = 0$. Since $h \in \omega_\mu(C)$, there exists $i \in C$ such that $i \trianglerighteq_h \mu^{-1}(h)$. But this implies $i \trianglerighteq_{h^0} \mu^{-1}(h^0)$, as desired.

Now, suppose $k \geq 2$. Let $(h^0, \ldots, h^{k-2}, h^{k-1})$ be a sequence of distinct houses such that

$$\mu^{-1}(h^0) \not\trianglerighteq_{h^0} \cdots \not\trianglerighteq_{h^{k-2}} \mu^{-1}(h^{k-2}) \not\trianglerighteq_{h^{k-2}} \mu^{-1}(h^{k-1}) \not\trianglerighteq_{h^{k-1}} \mu^{-1}(h^0).$$

Without loss of generality, let $h^0 = h$. Since $h \in \omega_\mu(C)$, there exists $i \in C$ such that $i \trianglerighteq_h \mu^{-1}(h)$. Equivalently, $i \trianglerighteq_{h^0} \mu^{-1}(h^0)$. There are two cases. If $i = \mu^{-1}(h^0)$, then $i \trianglerighteq_{h^{k-1}} i = \mu^{-1}(h^0)$ and $(\star)$ is satisfied. Otherwise $i \triangleright_{h^0} \mu^{-1}(h^0)$ and there are two possibilities. If $i \trianglerighteq_{h^0} \mu^{-1}(h^1)$, then we are done and $(\star)$ is satisfied. Else, $i \ntrianglerighteq_{h^0} \mu^{-1}(h^1)$ and by acylcicity,

$$i \triangleright_{h^0} \mu^{-1}(h^0) \;\&\; i \ntrianglerighteq_{h^0} \mu^{-1}(h^1) \implies \mu^{-1}(h^1) \triangleright_{h'} \mu^{-1}(h^0)$$

for all $h' \neq h^0$. In particular, $\mu^{-1}(h^1) \triangleright_{h^1} \mu^{-1}(h^0)$. Since $\mu^{-1}(h^1) \ntrianglerighteq_{h^1} \mu^{-1}(h^2)$, acyclicity implies that $\mu^{-1}(h^2) \triangleright_{h^2} \mu^{-1}(h^0)$. Continuing in this manner by induction leads us to conclude that $\mu^{-1}(h^{k-1}) \triangleright_{h^{k-1}} \mu^{-1}(h^0)$, which is a contradiction. Thus, $(\star)$ is true for all sequences of length $k' \leq k$. □

Proof of Theorem 2. Let μ be the assignment identified by Algorithm 2 for some completion $\tilde{\triangleright}$ of $\triangleright$. We note that Algorithm 2 constructs μ sequentially by removing cycles of agents $(\tilde{I}^1, \tilde{I}^2, \ldots)$ and associated houses $(\tilde{H}^1, \tilde{H}^2, \ldots)$. To derive a contradiction, suppose coalition C can indirectly exclusion block μ with σ. Thus, $\sigma(i) \succ_i \mu(i)$ for all $i \in C$ and

$$\mu(j) \succ_j \sigma(j) \implies \mu(j) \in \Omega(C|\omega_\mu, \mu). \tag{A.2}$$

We organize the proof's remainder as a series of claims.

Claim 1. *Suppose $i \in C$ and $i \in \tilde{I}^{t_i}$. Then $\sigma(i) \in \tilde{H}^t$ for some $t < t_i$.*

Proof of Claim 1. At each step of the algorithm, each remaining agent points to his favorite house that has not been removed from the market. Thus, if house $\sigma(i)$ was not yet assigned at step t_i and $\sigma(i) \succ_i \mu(i)$, agent i should have been pointing to some $h \succeq_i \sigma(i) \succ_i \mu(i)$ at step t_i rather than at $\mu(i)$. Hence, $\sigma(i) \in \tilde{H}^t$ for some $t < t_i$. ◇

Claim 2. *Let $J \subseteq I$ and suppose $h \in \omega_\mu(J)$. If $h \in \tilde{H}^t$, then there exists $i \in J$ such that $i \in \tilde{I}^{t_i}$ for some $t_i \leq t$.*

Proof of Claim 2. Suppose the contrary and that $t_i > t$ for all $i \in J$. Thus, each agent $i \in J$ remains in the market at step t when house h is assigned. First suppose that $\mu(j) = h$ and j is $\tilde{\triangleright}_h$-maximal at step t. Since $h \in \omega_\mu(J)$, there exists $i \in J$ such that $i \trianglerighteq_h \mu^{-1}(h)$. As every agent in J is assigned after step t, $i \neq j$. Thus, $i \triangleright_h \mu^{-1}(h) = j \implies i \tilde{\triangleright}_h j$, which contradicts j being $\tilde{\triangleright}_h$-maximal at step t.

Suppose instead that house h is assigned in step t as part of a cycle that involves two or more agents and houses. Let

$$j^0 \to h^0 \to j^1 \to h^1 \to \cdots \to j^{k-1} \to h^{k-1} \to j^0$$

be this cycle of k agents and k houses. For each ℓ, $\mu(j^\ell) = h^\ell$ and $j^{\ell+1 \pmod k}$ is $\tilde{\rhd}_{h^\ell}$-maximal among agents in I^t. Without loss of generality, let $h = h^0$. Clearly, $j^0 \not\trianglerighteq_{h^0} j^1$. Since $h^0 \in \omega_\mu(J)$, there exists $i \in J$ such that $i \trianglerighteq_{h^0} \mu^{-1}(h) = j^0$. Since i is assigned a house in a later step, $i \rhd_{h^0} j^0$. Since $h^0 \to j^1$ at step t and $i \in I^t$, $i \not\trianglerighteq_{h^0} j^1$. Acyclicity implies that $j^1 \rhd_{h'} j^0$ for all $h' \neq h^0$. In particular, this implies that $j^1 \rhd_{h^{k-1}} j^0 \implies j^1 \tilde{\rhd}_{h^{k-1}} j^0$, which is a contradiction since j^0 was $\tilde{\rhd}_{h^{k-1}}$-maximal among agents in I^t. $\diamond$

Claim 3. *Let t_i be the step of Algorithm 2 where agent i is assigned a house and removed from the market, i.e. $i \in \tilde{I}^{t_i}$. Let $J \subseteq I$. There exists $i \in J$ such that $t_i \leq t_j$ for all $j \in (\mu^{-1} \circ \omega_\mu)(J)$.*

Proof of Claim 3. Let $j \in (\mu^{-1} \circ \omega_\mu)(J)$. Observe that $\mu(j) \in \omega_\mu(J)$ and $\mu(j) \in \tilde{H}^{t_j}$. By Claim 2, there exists $i \in J$ such that $i \in \tilde{I}^{t_i}$ and $t_i \leq t_j$. As the number of agents is finite, there exists some $i \in J$ who is assigned before all agents in $(\mu^{-1} \circ \omega_\mu)(J)$. $\diamond$

Henceforth, consider the earliest cycle occurring in Algorithm 2 that contains an agent j such that $\mu(j) \succ_j \sigma(j)$. Without loss of generality, $\mu(j) \in H$. Suppose this cycle is removed at step t_j. Let $\tilde{I}^{t_j}$ and $\tilde{H}^{t_j}$ be the sets of agents and houses, respectively, involved.

Claim 4. *If $i \in C$, then $i \in \bigcup_{t > t_j} \tilde{I}^t$, i.e. $C \cap (\bigcup_{t \leq t_j} \tilde{I}^t) = \varnothing$.*

Proof of Claim 4. First, suppose there exists $j_1 \in C \cap \tilde{I}^{t_j}$. By Claim 1, $\sigma(j_1) \in \tilde{H}^{t_2}$ for some $t_2 < t_j$. It follows that there exists $j_2 \in C \cap \tilde{I}^{t_2}$. Otherwise, $\sigma(j_1) = \mu(i') \succ_{i'} \sigma(i')$ for some $i' \in \tilde{I}^{t_2}$, which contradicts t_j being the earliest cycle involving an agent who strictly preferred their assignment under μ to that under σ. By Claim 1, $\sigma(j_2) \in \tilde{H}^{t_3}$ for some $t_3 < t_2$. Clearly, we may continue this reasoning by induction without end, which is a contradiction as there is a finite number of cycles. Thus, $C \cap \tilde{I}^{t_j} = \varnothing$. Analogous reasoning, starting the argument at any $t' < t_j$, confirms that $C \cap \tilde{I}^{t'} = \varnothing$. Thus, $C \cap \left(\bigcup_{t \leq t_j} \tilde{I}^t \right) = \varnothing$. $\diamond$

Continuing with the same agent j and his assignment $\mu(j)$ as above, we now argue that $\mu(j) \notin \Omega(C|\omega_\mu, \mu)$. This will contradict (A.2) and therefore prove the theorem. Recall that

$\Omega(C|\omega_\mu, \mu) = \omega_\mu\left(\bigcup_{k=0}^{\infty} C_k\right)$ where $C_0 = C$ and $C_k = C_{k-1} \cup (\mu^{-1} \circ \omega_\mu)(C_{k-1})$. As $C_{k-1} \subseteq C_k$, it is sufficient to show that $\mu(j) \notin \omega_\mu(C_k)$ for each k.

First, suppose $\mu(j) \in \omega_\mu(C)$. Thus, there exists some $i \in C$ such that $i \trianglerighteq_{\mu(j)} j$. Furthermore, Claim 2 implies that there is some $i' \in C_1 = C$, such that $i' \in \tilde{I}^{t_{i'}}$ and $t_{i'} \leq t_j$ where t_j is the step of Algorithm 2 where house $\mu(j)$ is assigned. By Claim 4, no members of C are assigned a house at step t_j, or earlier. Hence, we have arrived at a contradiction.

Continuing by induction, suppose $\mu(j) \notin \omega_\mu(C_{k'})$ for all $k' < k$. Suppose $\mu(j) \in \omega_\mu(C_k)$. By definition, $C_k = C_{k-1} \cup (\mu^{-1} \circ \omega_\mu)(C_{k-1})$. Again, Claim 2 implies that there is some $i' \in C_k$, such that $i' \in \tilde{I}^{t_{i'}}$ and $t_{i'} \leq t_j$ where t_j is the step of Algorithm 2 where house $\mu(j)$ is assigned. However, repeated application of Claim 3 implies that the agent in C_k who is assigned a house earliest is necessarily a member of $C_1 = C$. By Claim 4, no members of C are assigned a house at step t_j, or earlier, of Algorithm 2—a contradiction. □

Proof of Theorem 3. Let μ be the assignment identified by Algorithm 2 for some completion $\tilde{\triangleright}$ of $\triangleright$. To prove Theorem 3, it is sufficient to modify the proof of Theorem 2 as follows. First, replace ω_μ with ω_μ^* throughout. And second, replace Claim 2 with the following.

Claim 2′. *Let $J \subseteq I$ and suppose $h \in \omega_\mu^*(J)$. If $h \in \tilde{H}^t$, then there exists $i \in J$ such that $i \in \tilde{I}^{t_i}$ for some $t_i \leq t$.*

Proof of Claim 2′. If $h \in \tilde{H}^t$, there exists a cycle of agents and houses such that

$$j^0 \to h^0 \to j^1 \to h^1 \to \cdots \to j^{k-1} \to h^{k-1} \to j^0.$$

For each ℓ, $\mu(j^\ell) = h^\ell$, $j^{\ell+1 \pmod k}$ is $\tilde{\triangleright}_{h^\ell}$-maximal among agents in I^t, and $h \in \{h^0, \ldots, h^{k-1}\}$. In this cycle, each agent is pointing to the house that he is assigned and each house h' is pointing to the $\tilde{\triangleright}_{h'}$-maximal agent who remains in the market. It follows that $j^0 \not\triangleright_{h^1} j^2 \not\triangleright_{h^2} \cdots \not\triangleright_{h^{k-1}} j^{k-1} \not\triangleright_{h^{k-1}} j^0$. If $h \in \omega_\mu^*(J)$, there exists $i \in J$ such that $i \trianglerighteq_{h^{\ell-1 \pmod k}} \mu^{-1}(h^\ell) = j^\ell$. As j^ℓ was the $\tilde{\triangleright}_{h^{\ell-1 \pmod k}}$-maximal agent remaining in the market and $i \trianglerighteq_{h^{\ell-1 \pmod k}} j_\ell$, house $h^{\ell-1 \pmod k}$ must have pointed to agent i at some step $t' \leq t$ of Algorithm 2. Hence, agent i was not in the market at step t or $j_\ell = i$. In either case, $i \in \tilde{I}^{t_i}$ for some $t_i \leq t$. ◇

□

Proof of Theorem 4. Let μ be a strong exclusion core allocation. We will construct a completion $\tilde{\triangleright}$ such that μ is the assignment generated by the algorithm. The argument proceeds as follows. We define a sequence of graphs. In each graph, we identify a set of agents who are assigned their most-preferred house among those in the graph. We then order the agents such that each set is "cleared" together by the GTTC algorithm (if this set forms a certain type of cycle) or in sequence (if this set forms a certain type of open path).

Let $I^1 := I, H^1 := H$. Construct a graph Γ^1 with vertices $I^1 \cup H^1 \cup \{h_0\}$. Let there be an arc from $i \in I^1$ to $h \in H^1 \cup \{h_0\}$ if and only if $\mu(i) = h$. Let there be an arc from each $h \in H^1$ to $i \in I^1$ if and only if i is $\rhd_h$-maximal among agents I^1. Finally, there is an arc from the outside option h_0 to every $i \in I^1$.

Let $\tau^1(i)$ denote the highest-ranked house in i's preference order among $H^1 \cup \{h_0\}$.

Claim 1. *The graph Γ^1 contains at least one cycle in which each agent i points to $\tau^1(i)$.*

Proof of Claim 1. If $\tau^1(i) = h_0$ for some i, then $\mu(i) = h_0$. If the outside option is an agent's most-preferred assignment, he is always able to block any allocation that does not assign him to h_0. Thus, the cycle $i \to h_0 \to i$ satisfies the claim.

Instead, suppose $\tau^1(i) \neq h_0$ for all i. Note that for each i, house $\tau^1(i)$ must be occupied by some agent at μ. Otherwise, if $\tau^1(i)$ is vacant, agent i would be able to indirectly exclusion block μ unilaterally. Assume toward contradiction that there is no cycle satisfying the above claim. Construct an alternating sequence of agents and houses as follows. First, fix some enumeration of all agents in $I^1 = \{i_1, i_2, \ldots\}$. (This index can be arbitrary, but it must be fixed.) Start with some agent $i^0 \in I^1$ and let $h^0 = \tau^1(i^0)$. Continuing by induction, given a sequence $(i^0, h^0, \ldots, i^{k-1}, h^{k-1})$, let i^k be the agent with the lowest index number (given the fixed enumerate) such that (a) $i^k \trianglerighteq_{h^{k-1}} \mu^{-1}(h^{k-1})$ and (b) i^k is $\rhd_{h^{k-1}}$-maximal among agents in I^1. Let $h^k := \tau^1(i^k)$. As there is a finite number of agents and houses, the sequence $(i^0, h^0, \ldots)$ must eventually flow into a cycle. Relabeling as necessary, and without loss of generality, let $(i^0, h^0, \ldots, i^{k-1}, h^{k-1})$ be that cycle. Thus, $i^0 \trianglerighteq_{h^{k-1}} \mu^{-1}(h^{k-1})$.

Let C be the set of agents in this cycle such that $\tau^1(i^\ell) \succ_{i^\ell} \mu(i^\ell)$. It follows that $C \neq \varnothing$.[34] Given the cycle $(i^0, h^0, \ldots, i^{k-1}, h^{k-1})$ and the fact that $h^\ell \in \omega_\mu(i^{\ell+1 \pmod k})$ for all ℓ, it follows that $\{h^0, \ldots, h^{k-1}\} \subseteq \Omega(C|\omega_\mu, \mu)$. Thus, coalition C can indirectly exclusion block μ by reallocating their most preferred houses among themselves, which is a contradiction. Therefore, we conclude that there exists at least one cycle in Γ^1 where each agent i points to $\tau^1(i)$. $\diamond$

Noting Claim 1, if Γ^1 contains a cycle where $i \to h_0 \to i$ and $\mu(i) = \tau^1(i) = h_0$, let $K^1 = (i, h_0)$. Otherwise, let $K^1 = (i^0, h^0, \ldots, i^{k-1}, h^{k-1})$ be a cycle in Γ^1 in which each agent i points to $\tau^1(i) \neq h_0$. By definition of Γ^1, $\tau^1(i) = \mu(i)$ for each agent i in K^1.

[34] Otherwise the cycle $(i^0, h^0, \ldots, i^{k-1}, h^{k-1})$ would form a cycle in the graph Γ^1 where each agent i points to $\tau^1(i)$. This situation has been ruled out by assumption.

Now, define $I^2 := I^1 \setminus (K^1 \cap I^1)$ and $H^2 := H^1 \setminus (K^1 \cap H^1)$. We can construct a graph Γ^2, with vertices $I^2 \cup H^2 \cup \{h_0\}$, using the same procedure as for Γ^1. It is straightforward to adapt the argument of Claim 1 to conclude that Γ^2 has a cycle K^2 where each agent $i \in K^2 \cap I^2$ is pointing to $\tau^2(i) = \mu(i)$ and $\tau^2(i)$ is agent i's most preferred house among those in $H^2 \cup \{h_0\}$. Continuing in this manner we can define a sequence of cycles $(K^1, K^2, \ldots, K^T)$ until no agents remain in Γ^T. (The outside option h_0 is always a member of Γ^T.)

Next we will use the sequence of cleared cycles to define a completion $\tilde{\rhd}_h$ for each $h \in H$. Consider cycle K^1. If $K^1 = (i, h_0)$, there is nothing to do and we can move to K^2. Otherwise, suppose K^1 defines a cycle of the form $i^0 \to h^0 \to i^1 \to \ldots h^{k-1} \to i^0$. For each h^ℓ let $i^{\ell+1 \pmod k}$ be the (unique) maximal element under $\tilde{\rhd}_{h^\ell}$, i.e. $i^{\ell+1 \pmod k} \tilde{\rhd}_{h^\ell} j$ for all $j \neq i^{\ell+1 \pmod k}$. The remainder of $\tilde{\rhd}_{h^\ell}$ can be defined in any manner not violating $\rhd_{h^\ell}$.

Continuing by induction, consider cycle K^t. If K^t includes the outside option, i.e. $K^t = (i, h_0)$, there is nothing to do and we can move to K^{t+1}. Otherwise, suppose K^t defines a cycle of the form $i^0 \to h^0 \to i^1 \to \ldots h^{k-1} \to i^0$. For each h^ℓ define $\tilde{\rhd}_{h^\ell}$ as follows. First, identify all agents $j \in I \cap (\cup_{\tau<t} K^\tau)$ such that $j \rhd_{h^\ell} i^{\ell+1 \pmod k}$. Let J be this set. Order these agents in an arbitrary manner not violating $\rhd_{h^\ell}$. Place $i^{\ell+1 \pmod k}$ in the $\tilde{\rhd}_{h^\ell}$ order immediately after all agents in J. Rank all remaining agents $J' = I \setminus (J \cup \{i^{\ell+1 \pmod k}\})$ in an arbitrary manner after $i^{\ell+1 \pmod k}$ such that $\rhd_{h^\ell}$ is not violated. The constructed completion $\tilde{\rhd}_{h^\ell}$ should have the following structure:

$$\underbrace{j^1 \; \tilde{\rhd}_{h^\ell} \; \cdots \; \tilde{\rhd}_{h^\ell} \; j^k}_{J} \; \tilde{\rhd}_{h^\ell} \; i^{\ell+1 \pmod k} \; \tilde{\rhd}_{h^\ell} \; \underbrace{j' \; \tilde{\rhd}_{h^\ell} \cdots}_{J'}.$$

Finally, if house h has not been assigned a completion as part of the preceding steps (and thus it is unassigned under μ), we can let $\tilde{\rhd}_h$ be an arbitrary completion of $\rhd_h$.

One can now verify that the GTTC algorithm outputs μ when $\tilde{\rhd}_h$ is the linear completion of $\rhd_h$ for each $h \in H$. In particular, (up to the order of cleared simultaneous disjoint cycles or simultaneous cycles involving h_0) K^t is the cleared cycle in step t of the algorithm. □

Proof of Proposition 4. Given $\langle I, H, \succ, \rhd \rangle$ and μ, let $\tilde{\rhd}$ be any admissible completion of the priority structure $\rhd$ such that for each h, $i \tilde{\rhd}_h \mu^{-1}(h) \iff i \rhd_h \mu^{-1}(h)$.[35] Let σ be the allocation identified by the GTTC algorithm given $\tilde{\rhd}$. Let $\tilde{I}^t$ be the set of agents assigned to a house in step t of the algorithm. Since $\rhd$ is acyclic, σ is an exclusion core allocation. We will show that if μ is not an exclusion core allocation, the coalition $C = \{i \mid \sigma(i) \succ_i \mu(i)\}$ can indirectly exclusion block μ with σ given ω_μ.

[35] Recall that by convention $i \rhd_h \varnothing$ for every $h \in H$.

First, we observe that $C \neq \varnothing$. We know this because $\sigma \neq \mu$ and if $\mu(i) \succeq_i \sigma(i)$ for all i, then σ would not be Pareto-optimal, a contradiction.

To derive a contradiction, suppose that C cannot indirectly exclusion block μ with σ. Thus, $\exists\ j \notin C$ such that $\mu(j) \succ_j \sigma(j)$ and $\mu(j) \notin \Omega(C|\omega_\mu, \mu)$. Out of all the agents who satisfy these conditions, let j^0 be one who was assigned at the earliest step of the GTTC algorithm.[36] Suppose that j^0 is assigned in step t_0 of the GTTC algorithm, i.e. $j^0 \in \tilde{I}^{t_0}$. Since $\sigma(j^0) \succeq_{j^0} h_0$, it follows that $\mu(j) \neq h_0$. Thus, that $\mu(j^0) = h^1$ for some $h^1 \in H$. Since $h^1 \succ_{j^0} \sigma(j^0)$, h^1 must have been assigned before step t_0, say in step $t_1 < t_0$. Thus, there must exist some agent $j^1 \in \tilde{I}^{t_1}$ such that $j^1 \tilde{\trianglerighteq}_{h^1} i$ for all $i \in \bigcup_{t \geq t_1} \tilde{I}^t$. In particular, given the definition of $\tilde{\triangleright}$, $j^1 \tilde{\triangleright}_{h^1} j^0 = \mu^{-1}(j^0)$ if and only if $j^1 \triangleright_{h^1} j^0$ and thus $\mu(j^0) = h^1 \in \omega_\mu(j^1)$.

If $\sigma(j^1) \succ_{j^1} \mu(j^1)$, then $j^1 \in C$ and thus $\mu(j^0) \in \Omega(C|\omega_\mu, \mu)$, which is a contradiction. If $\mu(j^1) \succ_{j^1} \sigma(j^1)$ instead, then, since $j^0 \in \tilde{I}^{t_0}, j^1 \in \tilde{I}^{t_1}, t_1 < t_0$ and j^0 was chosen to be the earliest agent assigned by GTTC for whom both $\mu(j^0) \succ_{j^0} \sigma(j^0)$ and $\mu(j^0) \notin \Omega(C|\omega_\mu, \mu)$, it follows that $\mu(j^1) \in \Omega(C|\omega_\mu, \mu)$. Since $\mu(j^0) \in \omega_\mu(j^1)$, this means $\mu(j^0) \in \Omega(C|\omega_\mu, \mu)$, again a contradiction.

Thus, we conclude $h^2 = \mu(j^1) = \sigma(j^1)$. In this case, h^2 is assigned at step $t_2 = t_1$ of the GTTC algorithm given $\tilde{\triangleright}$.[37] In particular, there exists an agent $j^2 \in \tilde{I}^{t_2} = \tilde{I}^{t_1}$. Such that $j^2 \tilde{\trianglerighteq}_{h^2} i$ for all $i \in \bigcup_{t \geq t_2} \tilde{I}^t$. In particular, given the definition of $\tilde{\triangleright}_{h^2}$, this implies $j^2 \trianglerighteq_{h^2} \mu^{-1}(h^2) = j^1$. We know that $j^2 \neq j^1$. This is because the cycle formed by agents and houses assigned at step t_1 must include the agent who is assigned to h^1 by the algorithm.[38] Thus, $j^2 \triangleright_{h^2} j^1$, which implies $\mu(j^1) = h^2 \in \omega_\mu(j^2)$.

Thus, we can find a chain of agents $(j^0, j^1, \ldots, j^n)$ such that $\mu(j^k) = \sigma(j^k)$ for all $k = 1, \ldots, n-1$, $j^k \in \tilde{I}^{t_1}$ for $k = 1, 2, \ldots, n$, $\mu(j^k) \in \omega_\mu(j^{k+1})$ for $k = 1, \ldots, n-1$, and $\mu(j^n) \neq \sigma(j^n)$. The existence of such a j^n in step $\tilde{I}^{t_1}$ is guaranteed because $h^1 = \mu(j^0)$ and $j^0 \notin \tilde{I}^{t_1}$. Now, regardless of whether $\mu(j^n) \succ_{j^n} \sigma(j^n)$ or $\sigma(j^n) \succ_{j^n} \mu(j^n)$, using arguments similar to the ones we applied to j^1 above we can see that $\mu(j^n)$ is in the set $\Omega(C|\omega_\mu, \mu)$. But this implies $\mu(j^{n-1}), \ldots, \mu(j^1)$, and $\mu(j^0)$ belong to $\Omega(C|\omega_\mu, \mu)$ as well—a contradiction. □

Proof of Proposition 5. Let ω_μ be the weak conditional endowment system in $\mathcal{E} = \langle I, H, \succ \triangleright \rangle$; define ω'_μ analogously for $\mathcal{E}' = \langle I, H, \succ \triangleright' \rangle$. To prove that the strong exclusion core of $\mathcal{E}$ is a subset of the strong exclusion core of $\mathcal{E}'$, it is sufficient to show that for all $C \subseteq I$, $C \neq \varnothing$, and any allocation μ, $\omega'_\mu(C) \subseteq \omega_\mu(C)$. Suppose $h \in \omega'_\mu(C)$. Thus, there exists $i \in C$

[36] If multiple such agents are assigned in the same step, pick any of them.
[37] Since more than one house is assigned in the same step, $h^2 \in H$.
[38] If $j^2 = j^1$, the identified cycle would involve only one agent and one house.

such that $i \trianglerighteq'_h \mu^{-1}(h)$. And so, $i \trianglerighteq_h \mu^{-1}(h)$, which implies $h \in \omega_\mu(C)$. The analogous result for the weak exclusion cores follows similarly. □

Proof of Proposition 6. Consider the economy $\mathcal{E}$ with priority structure $\triangleright$. Let $\tilde{\triangleright}$ be a completion of $\triangleright$ and let μ be the allocation identified by the GTTC algorithm given $\tilde{\triangleright}$. Now consider economy $\mathcal{E}'$ with priority structure $\triangleright'$. For each $h \in H$, define $\tilde{\triangleright}'_h$ as any completion of $\triangleright'_h$ such that for all $k, j \neq i$, $k \tilde{\triangleright}_h j \iff k \tilde{\triangleright}'_h j$ and $i \tilde{\triangleright}_h j \implies i \tilde{\triangleright}'_h j$. $\tilde{\triangleright}' = (\tilde{\triangleright}'_h)_{h \in H}$ is an admissible completion for $\triangleright'$. Each $\tilde{\triangleright}'_h$ ranks i at least as high as $\tilde{\triangleright}_h$ without changing the order of the other agents. Let μ' be the allocation identified by the GTTC algorithm given $\tilde{\triangleright}'$. Because the top trading cycles algorithm with complete priority rankings respects improvements in priorities,[39] $\mu(i) \succeq_i \mu'(i)$. The lemma's second part is proved analogously, except agent i's position in the derived completion is downgraded, if necessary. □

Proof of Lemma 3. (a) When $\triangleright_h$ represents ω, $\omega^{**}(C) = \omega(C)$ for all C. The result follows.

(b) Generally, the weak exclusion core is contained in the unconditional exclusion core. Noting part (a), it is sufficient to show that the exclusion core of $\langle I, H, \succ, \omega \rangle$ is contained in the weak exclusion core of the corresponding relational economy.

If $\triangleright$ represents ω, then $h \in \omega(i) \iff i \triangleright_h j \ \forall j \neq i$; otherwise, agents are not $\triangleright$.-comparable. Let μ be an exclusion core allocation in $\langle I, H, \succ, \omega \rangle$. Given this allocation, $h \in \omega_\mu(i)$ if and only if (i) $h \in \omega(i)$, (ii) $h = \mu(i)$, or (iii) $\mu^{-1}(h) = \varnothing$. Suppose coalition C can indirectly exclusion block μ with σ in the relational economy $\langle I, H, \succ, \triangleright \rangle$ given ω_μ. If $\mu(j) \succ_j \sigma(j)$, then $\mu(j) = h \in \Omega(C|\omega_\mu, \mu)$ and $j \notin C$. Thus, there exists a sequence of agents $i^1, \ldots, i^K$ such that $h \in \omega_\mu(i^1)$, $\mu(i^1) \in \omega_\mu(i^2)$, $\ldots$, $\mu(i^{K-1}) \in \omega_\mu(i^K)$ and $i^K \in C$. But, this implies $h \in \omega(i^1)$, $\mu(i^1) \in \omega(i^2)$, $\ldots$, $\mu(i^{K-1}) \in \omega(i^K)$. Therefore, $h \in \Omega(C|\omega, \mu)$ and coalition C can indirectly exclusion block μ in $\langle I, H, \succ, \omega \rangle$, which is a contradiction. □

B The GTTC Algorithm: An Example

Example B.1. To illustrates the operation of the GTTC algorithm, suppose there are four agents and four houses. The agents' preferences are

$$\succ_{i_1}: h_2, h_1, h_4, h_3 \qquad \succ_{i_2}: h_4, h_1, h_3, h_2 \qquad \succ_{i_3}: h_3, h_4, h_1, h_2 \qquad \succ_{i_4}: h_3, h_1, h_2, h_4.$$

[39] The earliest reference to this property of the TTC algorithm we are aware of is by Abdulkadiroğlu and Che (2010), who state it without proof in an unpublished working paper. (The proof is straightforward and relies on the observation that when an agent's priority ranking improves he will be assigned in an earlier cycle by the algorithm.) Balinski and Sönmez (1999) introduced the notion of an assignment mechanism "respecting improvements." They examined the deferred acceptance algorithm in a college admissions setting.

(a) $\rhd_{h_1}$. (b) $\rhd_{h_2}$. (c) $\rhd_{h_3}$. (d) $\rhd_{h_4}$.

Figure B.1: Priority structures in Example B.1.

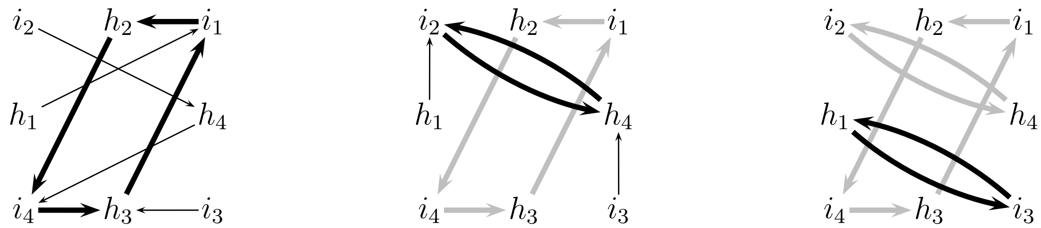

(a) Step 1. (b) Step 2. (c) Step 3.

Figure B.2: Operation of the GTTC algorithm given the completion (B.1).

Panels (a)–(d) of Figure B.1 define this economy's priority structure.

First, suppose the completion of $\rhd$ is given by

$$\tilde{\rhd}_{h_1} : i_1, i_2, i_3, i_4 \qquad \tilde{\rhd}_{h_2} : i_4, i_1, i_2, i_3 \qquad \tilde{\rhd}_{h_3} : i_1, i_2, i_3, i_4 \qquad \tilde{\rhd}_{h_4} : i_4, i_1, i_2, i_3. \tag{B.1}$$

Panels (a)–(c) of Figure B.2 illustrate the GTTC algorithm's operation. (We omit the outside option h_0 from the figure.) In step 1, i_1 is assigned h_2 and i_4 is assigned h_3. In step 2, agent i_2 is assigned h_4. In step 3, i_3 is assigned h_1. The final allocation is

$$\mu(i_1) = h_2 \qquad \mu(i_2) = h_4 \qquad \mu(i_3) = h_1 \qquad \mu(i_4) = h_3.$$

Suppose instead that the completion of $\rhd$ is given by

$$\tilde{\rhd}'_{h_1} : i_1, i_2, i_3, i_4 \qquad \tilde{\rhd}'_{h_2} : i_4, i_1, i_2, i_3 \qquad \tilde{\rhd}'_{h_3} : i_1, i_2, i_3, i_4 \qquad \tilde{\rhd}'_{h_4} : i_4, i_3, i_1, i_2. \tag{B.2}$$

(B.2) is identical to (B.1) except i_3 ranks ahead of i_1 and i_2 in $\tilde{\rhd}'_{h_4}$. Figure B.3 illustrates GTTC algorithm's operation. (Again, the figure omits h_0.) In step 1, i_1 is assigned h_2 and i_4 is assigned h_3. In step 2, i_3 receives h_4. In step 3, i_2 is assigned h_1. The final allocation is

$$\nu(i_1) = h_2 \qquad \nu(i_2) = h_1 \qquad \nu(i_3) = h_4 \qquad \nu(i_4) = h_3.$$

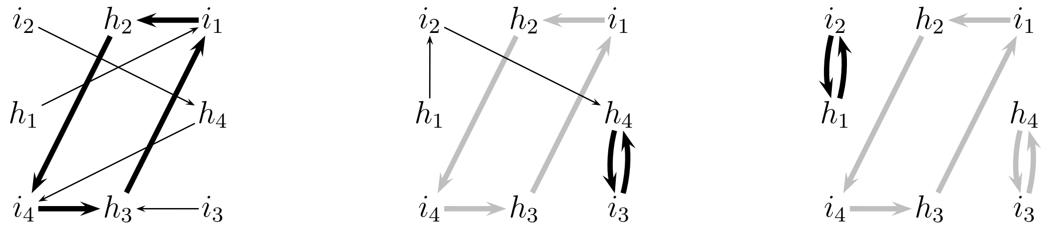

(a) Step 1. (b) Step 2. (c) Step 3.

Figure B.3: Operation of the GTTC algorithm given the completion (B.2).

It can be verified that μ and ν are the only possible assignments generated by the GTTC algorithm in this economy. Every other completion will lead to one of these two assignments.

References

Abdulkadiroğlu, A. and Che, Y.-K. (2010). The role of priorities in assigning indivisible objects: A characterization of top trading cycles. Mimeo.

Abdulkadiroğlu, A., Pathak, P. A., and Roth, A. E. (2009). Strategy-proofness versus efficiency in matching with indifferences: Redesigning the NYC high school match. *American Economic Review*, 99(5):1954–1978.

Abdulkadiroğlu, A. and Sönmez, T. (1998). Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica*, 66(3):689–701.

Abdulkadiroğlu, A. and Sönmez, T. (1999). House allocation with existing tenants. *Journal of Economic Theory*, 88(2):233–260.

Abdulkadiroğlu, A. and Sönmez, T. (2003). School choice: A mechanism design approach. *American Economic Review*, 93(3):729–747.

Alcalde-Unzu, J. and Molis, E. (2011). Exchange of indivisible goods and indifferences: The top trading absorbing sets mechanisms. *Games and Economic Behavior*, 73:1–16.

Arruñada, B. (2012). Property as an economic concept: Reconciling legal and economic conceptions of property rights in a Coasean framework. *International Review of Economics*, 59(2):121–144.

Athanassoglou, S. and Sethuraman, J. (2011). House allocation with fractional endowments. *International Journal of Game Theory*, 40:481–513.

Aziz, H. (2015). Generalizing top trading cycles for housing markets with fractional endowments. Mimeo.

Balinski, M. and Sönmez, T. (1999). A tale of two mechanisms: Student placement. *Journal of Economic Theory*, 84:73–94.

Bigham, B. (2016). One in five willing organ donors in Ontario have their wishes vetoed by family—and doctors don't argue. *National Post*, November 25, 2016. http://news.nationalpost.com/news/one-in-five-willing-organ-donors-in-ontario-have-their-wishes-vetoed-by-family-and-doctors-dont-argue (accessed December 7, 2016).

Campbell, K. (1992). On the general nature of property rights. *The King's College Law Journal*, 3:79–97.

Carroll, G. (2014). A general equivalence theorem for allocation of indivisible objects. *Journal of Mathematical Economics*, 51:163–177.

Coase, R. H. (1960). The problem of social cost. *Journal of Law and Economics*, 3:1–44.

Cronin, A. J. and Douglas, J. F. (2010). Directed and conditional deceased donor organ donations: Laws and misconceptions. *Medical Law Review*, 18:275–301.

Cronin, A. J. and Price, D. (2008). Directed organ donation: Is the donor the owner? *Clinical Ethics*, 3:127–131.

Delacrétaz, D., Kominers, S. D., and Teytelboym, A. (2016). Refugee resettlement. Mimeo.

Downie, J., Shea, A., and Rajotte, C. (2008). Family override of valid donor consent to postmortem donation: Issues in law and practice. *Transplantation Proceedings*, 40:1255–1263.

Ehlers, L. (2014). Top trading with fixed tie-breaking in markets with indivisible goods. *Journal of Economic Theory*, 151:64–87.

Ehlers, L. and Erdil, A. (2010). Efficient assignment respecting priorities. *Journal of Economic Theory*, 145:1269–1282.

Ekici, Ö. (2013). Reclaim-proof allocation of indivisible objects. *Games and Economic Behavior*, 81:1–10.

Ergin, H. (2002). Efficient resource allocation on the basis of priorities. *Econometrica*, 70(2):2489–2497.

Gale, D. and Shapley, L. S. (1962). College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):9–15.

Grossman, S. J. and Hart, O. D. (1986). The costs and benefits of ownership: A theory of vertical integration. *Journal of Political Economy*, 94(4):691–719.

Hart, O. and Moore, J. (1990). Property rights and the nature of the firm. *Journal of Political Economy*, 98(6):1119–1158.

Hylland, A. and Zeckhauser, R. J. (1979). The efficient allocation of individuals to positions. *Journal of Political Economy*, 87(2):293–314.

Jaramillo, P. and Manjunath, V. (2012). The difference indifference makes in strategy-proof allocation of objects. *Journal of Economic Theory*, 147:1913–1946.

Kawasaki, R. (2015). Roth-Postlewaite stability and von Neumann-Morgenstern stability. *Journal of Mathematical Economics*, 58:1–6.

Kesten, O. (2006). On two competing mechanisms for priority-based allocation problems. *Journal of Economic Theory*, 127:155–171.

Kesten, O. (2009). Why do popular mechanisms lack efficiency in random environments? *Journal of Economic Theory*, 144:2209–2226.

Klaus, B., Klijn, F., and Walzl, M. (2010). Farsighted house allocation. *Journal of Mathematical Economics*, 46(5):817–824.

Klick, J. and Parchomovsky, G. (2017). The value of the right to exclude: An empirical assessment. *University of Pennsylvania Law Review*, 165:917–966.

Konishi, H., Quint, T., and Wako, J. (2001). On the Shapley-Scarf economy: The case of multiple types of indivisible goods. *Journal of Mathematical Economics*, 35(1):1–15.

Koopmans, T. C. and Beckmann, M. (1957). Assignment problems and the location of economic activity. *Econometrica*, 25(1):53–76.

Ma, J. (1994). Strategy proofness and the strict core in a market with indivisibilities. *International Journal of Game Theory*, 23:75–83.

Merrill, T. W. (1998). Property and the right to exclude. *Nebraska Law Review*, 77(4):730–755.

Merrill, T. W. and Smith, H. E. (2001a). The property/contract interface. *Columbia Law Review*, 101(4):773–852.

Merrill, T. W. and Smith, H. E. (2001b). What happened to property in law and economics? *Yale Law Journal*, 111:357–398.

Mossoff, A. (2009). Exclusion and exclusive use in patent law. *Harvard Journal of Law and Technology*, 22(2):321–379.

Mumcu, A. and Saglam, I. (2007). The core of a housing market with externalities. *Economics Bulletin*, 3(55):1–5.

Pápai, S. (2000). Strategyproof assignment by hierarchical exchange. *Econometrica*, 68(6):1403–1433.

Penner, J. E. (1997). *The Idea of Property Law.* Oxford University Press, New York.

Piccione, M. and Rubinstein, A. (2007). Equilibirum in the jungle. *Economic Journal*, 117(522):883–896.

Pycia, M. and Ünver, M. U. (2017). Incentive compatible allocation and exchange of discrete resources. *Theoretical Economics*, 12(1):287–329.

Quinn, B. (2016). Families veto hundreds of organ donations in five years, figures show. *The Guardian*, January 15, 2016. https://www.theguardian.com/uk-news/2016/jan/15/nhs-moves-to-stop-bereaved-families-blocking-donation-of-relatives-organs (accessed January 4, 2017).

Richter, M. and Rubinstein, A. (2015). Back to fundamentals: Equilibrium in abstract economies. *American Economic Review*, 105(8):2570–94.

Roth, A. E. (1982). Incentive compatibility in a market with indivisible goods. *Economics Letters*, 9(2):127–132.

Roth, A. E. and Postlewaite, A. (1977). Weak versus strong domination in a market with indivisible goods. *Journal of Mathematical Economics*, 4:131–137.

Roth, A. E., Sönmez, T., and Ünver, M. U. (2004). Kidney exchange. *Quarterly Journal of Economics*, 119(2):457–488.

Saban, D. and Sethuraman, J. (2013). House alocation with indifferences: A generalization and a unified view. Mimeo.

Schummer, J. and Vohra, R. (2013). Assignment of arrival slots. *American Economic Journal: Microeconomics*, 5(2):164–185.

Segal, I. and Whinston, M. D. (2012). Property rights. In Gibbons, R. and Roberts, J., editors, *Handbook of Organization Economics*, chapter 3, pages 100–158. Princeton University Press.

Shapiro, C. (2000). Navigating the patent thicket: Cross licenses, patent pools, and standard setting. *Innovation Policy and the Economy*, 1:119–150.

Shapley, L. S. and Scarf, H. (1974). On cores and indivisibility. *Journal of Mathematical Economics*, 1(1):23–37.

Shapley, L. S. and Shubik, M. (1971). The assignment game I: The core. *International Journal of Game Theory*, 1(1):111–130.

Sönmez, T. (1999). Strategy-proofness and essentially single-valued cores. *Econometrica*, 67(3):677–689.

Sönmez, T. and Ünver, M. U. (2010). House allocation with existing tenants: A characterization. *Games and Economic Behavior*, 69(2):425–445.

Svensson, L.-G. and Larsson, B. (2005). Strategy-proofness, core, and sequential trade. *Review of Economic Design*, 9:167–190.

Tang, Q. and Zhang, Y. (2016). Hierarchical exchange rules and the core in indivisible objects allocation. Mimeo.

Truog, R. D. (2005). Are organs personal property or a societal resource? *The American Journal of Bioethics*, 5(4):14–16.

von Neumann, J. and Morgenstern, O. (1944). *Theory of Games and Economic Behavior*. Princeton University Press, Princeton, NJ.

Wako, J. (1984). A note on the strong core of a market with indivisible goods. *Journal of Mathematical Economics*, 13(2):189–194.

Wako, J. (1991). Some properties of weak domination in an exchange market with indivisible goods. *The Economic Studies Quarterly*, 42(4):303–314.