

Endowments, Exclusion, and Exchange

Faculty Research Working Paper Series

Ivan Balbuzanov

University of Melbourne

Maciej H. Kotowski

Harvard Kennedy School

Updated October 2018

RWP17-016

Visit the **HKS Faculty Research Working Paper Series** at:

<https://research.hks.harvard.edu/publications/workingpapers/Index.aspx>

The views expressed in the **HKS Faculty Research Working Paper Series** are those of the author(s) and do not necessarily reflect those of the John F. Kennedy School of Government or of Harvard University. Faculty Research Working Papers have not undergone formal review and approval. Such papers are included in this series to elicit feedback and to encourage debate on important public policy challenges. Copyright belongs to the author(s). Papers may be downloaded for personal use only.

Endowments, Exclusion, and Exchange*

Ivan Balbuzanov[†]

Maciej H. Kotowski[‡]

October 15, 2018

Abstract

We propose a new 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, social hierarchies, and production. The exclusion core is characterized by 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, Chinese University of Hong Kong, Deakin University, James Cook University, McGill University, National University of Singapore, Oxford University, Singapore Management University, Stanford University, Universitat Autònoma de Barcelona, University College London, University of Melbourne, and the University of Rochester for constructive comments. We thank Omer Ali, Chris Avery, Yeon-Koo Che, David Delacrétaz, 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 Saghaian, Ran Shorrer, Tayfun Sönmez, Alex Teytelboym, Utku Ünver, Rakesh Vohra, Bill Zame, and Richard Zeckhauser for helpful 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, Level 4 111 Barry Street, The University of Melbourne 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, endowments define what agents own with little elaboration. In practice, ownership is hardly so straightforward. Co-owners of a house may be tenants in common or joint tenants under bespoke arrangements. Socially recognized, but formally undocumented, claims to land are common in the developing world. Blocking patents, patent pools, and patent “thickets” typify the complexity of intellectual property (Shapiro, 2000). In light of such cases, the interpretation of an endowment and its relation to usual understandings of property is elusive, and little studied.

In this paper, we study an exchange economy that places complex endowments at the forefront. 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 qualified by relationships or social obligations. In the face of these complications, the economy’s core, a traditional benchmark solution, may be empty (for the strong core) or include unintuitive outcomes (weak core).¹

Our key contribution is the development of a new solution concept, which we call the exclusion core. The exclusion core’s foundation is a reinterpretation of endowments in an exchange economy as a distribution of exclusion rights, rather than as bundles of things to trade. A simple idea—the ability to exclude others from goods in one’s own endowment—offers reallocation possibilities that are absent from traditional core solutions and is at the heart of the exclusion core’s rationale and predictive power. We analyze the exclusion core’s properties in economies with single-unit demand, indivisible goods, and no transfers. Beyond its practical importance (see below), this setting lets us emphasize the variables of interest, including complex, hierarchical, and qualified forms of property rights.

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 on one defining principle of property—the right to exclude others. This right is a classic tenet of property, with roots in the mid-eighteenth century writings of William Blackstone and others (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.”²

Second, we show that the exclusion core has a close association with David Gale’s top

¹We define the strong and weak cores in Section 1 and more formally again 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 allowing for indifference. The weak core is defined with strong (or strict) domination.

²*Kaiser Aetna v. United States*, 444 U.S. 164 (1979).

trading cycle (TTC) algorithm (Shapley and Scarf, 1974). Beyond its theoretical elegance, the TTC algorithm is of substantial practical importance. Suitably generalized, it underpins implemented or proposed solutions to many 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). A generalized TTC algorithm characterizes the exclusion core 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, we also explain how the ability to exclude governs many allocation problems. By focusing on the distribution of exclusion rights, it is possible to analyze economies with well-defined, conditional, and even conflicting claims to goods using a common toolkit.

We divide our main analysis into two parts that differ in the relative complexity of the prevailing property rights. In Section 2 endowments are an exogenous primitive, the typical case in economic analysis. This familiar setting allows us to introduce the direct exclusion core and its refinement, the exclusion core. The latter is our focus. The exclusion core coincides with the strong core in Shapley and Scarf’s (1974) “house market,” a benchmark case. 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 exclusion core solution to situations where social or legal constraints introduce conflicting claims to goods. To model these cases, we introduce 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 acknowledge a technical parallel with centralized 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 indirectly govern exclusion rights by constraining the economy’s endowment system. We propose three versions of our solution applicable to relational economies, the strong, weak, and unconditional exclusion cores. These all stand on the behavioral foundation 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 or weak core, they are characterized by a generalized TTC algorithm and are stable solutions in the sense of von Neumann and Morgenstern (1944). Acyclic priority structures are common in practice, and include economies with private, public, and mixed ownership.

Though our analysis primarily considers exchange economies and allocation problems, our proposal is also relevant for production. In Section 4 we apply the exclusion core concept to a production economy by defining endowments and exclusion rights over existing and potential goods. We illustrate this application in a generalized “house exchange” economy where coalitions of agents additionally own firms that engage in production.

While we reference the related literature throughout our exposition, we offer a more structured survey in Section 5. Our paper contributes to the study of exchange economies by proposing the exclusion core, a new solution clarifying the foundations of exchange. Our solution’s inspiration in the right to exclude also lets us contribute to a debate primarily among legal scholars on the conceptual understanding of property. Finally, our results also have implications for the practice of market design. We do not pursue this objective, although we do offer a new rationale for the use of trading cycle algorithms in applications, including those cited above. Trading cycle procedures identify outcomes that are robust to the exercise of exclusion rights, which are closely tied to a classic characterization of property.

With the exception of some immediate corollaries, we relegate all proofs to the Appendix. An Online Appendix contains the formal analysis pertaining to production economies.

1 A Motivating Example

To motivate our argument, it is useful to examine a simple instance of our model. It highlights the limitations of existing theories and hints at the power of our emphasis on 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 the

two allocations is intuitive and efficient and, therefore, easiest to justify.³

It is surprising that neither the strong core nor the weak core, two prominent solutions for exchange and assignment economies, is able to converge on the preceding outcomes. An allocation belongs to the strong core if no coalition of agents can reallocate the goods 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 allocation 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 h_1 is occupied by k and h_2 is occupied by j instead, 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 no coalition of agents can reallocate the goods they own such that all coalition members are made strictly better off. In the above example, the weak core is too large. In fact, any assignment where agent k inhabits h_1 belongs to the weak core. This includes the odd situation where 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. Given the lack of externalities or information asymmetries, this inefficient outcome seems implausible.⁵

The Kingdom's troubles are neither special to the example, nor are they technical anomalies. The cores' deficiencies can be traced to faulty presumptions concerning agents' desire and ability to form a blocking coalition. The strong core is empty because 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 often harms another, which is hardly an altruistic disposition. The second is not-modeled side payments. An agent who benefits, the reasoning goes, can bribe those who remain indifferent to enforce a reassignment. This argument is unconvincing. Equally well a side payment can be extorted from a potentially harmed agent to prevent a reassignment, an often ignored possibility.

³The example's phrasing follows that of our model. An alternative framing is inspired by kidney exchange (Roth et al., 2004). Agent k has two kidneys, h_1 and h_2 , and has resolved 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.

⁴A possible remedy for the strong core's emptiness 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).

⁵The core concepts' deficiencies are not due to 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.

The weak core is immune to the questionable incentives that plague the strong core, 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 cannot induce those who remain indifferent to cooperate. Consequently, unintuitive and inefficient outcomes persist.

The exclusion core avoids the above shortcomings. An allocation is in the (direct) exclusion core if no coalition can strictly benefit from a reassignment of houses in which any agents hurt by the reassignment are excluded (i.e. evicted) from houses in the coalition's endowment. In the Kingdom, 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 h_2 , k gains nothing by evicting him and thus is unwilling to do so. Conversely, allowing either of i and j to occupy a previously empty h_2 does not harm k and he has no reason to prevent this move.

Associating endowments with exclusion rights proved insightful in the preceding example. Importantly, this reinterpretation extends to economies where the rules surrounding property are less definite. The complexities (and headaches) surrounding joint or ill-defined ownership immediately come to mind. Aside from legal prescriptions, status and social conventions also influence how goods are exchanged or allocated. These variables define property rights in practice, often implicitly. The role of exclusion rights is readily apparent. Bigger kids usually take the best toys unless an adult intervenes. An elderly man may expect a teenager to yield him a seat on a bus, though he would not demand similar deference from a blind passenger. And doctors often seek family members' permission before transplanting organs from deceased relatives, even when the deceased had consented to donation prior to death.⁶ Though the next of kin did not inherit their relative's organs, they are often conferred the right to exclude others from benefiting from them. By interpreting endowments as a distribution of exclusion rights, the exclusion core unifies many exchange and allocation problems under a common analytic umbrella.

2 Simple Economies

A *simple economy* $\langle I, H, \succ, \omega \rangle$ consists of agents, goods, preferences, and an endowment system. $I = \{i_1, \dots, i_n\}$ is a finite set of agents whom we sometimes denote by i , j , or k . $H = \{h_1, \dots, h_m\}$ 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

⁶We thank Al Roth for bringing this practice to our attention via his blog. In the United Kingdom, family objections blocked 547 transplants from 2010 to 2016 (Quinn, 2016). See also Downie et al. (2008).

take in 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* $\mu: I \rightarrow H \cup \{h_0\}$ 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. 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 $h = h'$. We often define $\succ_i$ by listing houses in preferred order, i.e. $\succ_i: h, h', \dots$. Unlisted houses are worse than the outside option.

An *endowment system* specifies the houses owned by each coalition. It is a function $\omega: 2^I \rightarrow 2^H$ satisfying three properties.

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

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

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

Condition (A1) restricts ownership to agents or groups. Condition (A2) states that a coalition has in its endowment anything that belongs to any sub-coalition. Finally, (A3) says that the grand coalition jointly owns everything.

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

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

We call C^h the *minimal controlling coalition* of house h . Condition (A4) guarantees that each house has a set of one or more “co-owners” without opposing and mutually exclusive claims. We relax (A4) in Section 3.⁸

Many economies satisfy (A1)–(A4), including the Kingdom (Example 1), as well as those examined by Shapley and Scarf (1974) and Hylland and Zeckhauser (1979), which we discuss below. These latter 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 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).

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

⁸Example 4 presents an economy that satisfies (A1)–(A3) but not (A4).

In the preceding section we explained why the strong and weak cores may fail to provide satisfactory guidance, even in simple problems. As suggested above, these solutions can be defined as allocations that cannot be “blocked” by any coalition. They rely on two variants of blocking, whose definitions we record for completeness.

Definition 1. A non-empty coalition $C \subseteq I$ can *weakly block* the allocation μ with allocation σ if (a) $\sigma(i) \succeq_i \mu(i)$ for all $i \in C$, (b) $\sigma(i) \succ_i \mu(i)$ for some $i \in C$, and (c) $\sigma(C) \subseteq \omega(C) \cup \{h_0\}$.

Definition 2. A non-empty coalition $C \subseteq I$ can *strongly block* the allocation μ with allocation σ if (a) $\sigma(i) \succ_i \mu(i)$ for all $i \in C$ and (b) $\sigma(C) \subseteq \omega(C) \cup \{h_0\}$.

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

2.1 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. By preventing others from using property in his endowment, an agent can secure and preserve his wellbeing. We explain this idea’s implications in two steps. First, we define the direct exclusion core to show the immediate power of the right to exclude. Second, we build on this solution by considering exclusion’s indirect implications, leading to the exclusion core.

As motivation, consider an economy with three agents and three houses. Each house h_k is owned by agent i_k and the agents’ preferences are

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

Consider the allocation

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

The coalition $C = \{i_1, i_2\}$ can strongly block μ with the allocation

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

The traditional interpretation of the move from μ to σ is that the coalition (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 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 each member strictly gains from an alternative and anyone harmed by the reallocation is excluded from a house belonging to the coalition.

Definition 3. A non-empty coalition $C \subseteq I$ can *directly exclusion block* the allocation μ with allocation σ if (a) $\sigma(i) \succ_i \mu(i)$ for all $i \in C$ and (b) $\mu(j) \succ_j \sigma(j) \implies \mu(j) \in \omega(C)$.

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

Though only a waypoint in our analysis, the direct exclusion core enjoys some appealing properties.

Lemma 1. *For any economy $\langle I, H, \succ, \omega \rangle$ where ω satisfies (A1)–(A4), the direct exclusion core is a non-empty subset of the weak core.*

Moreover, direct exclusion core allocations are Pareto efficient.⁹ This is because any Pareto-improving reallocation of houses can be invoked by its beneficiaries to directly exclusion block an allocation.¹⁰ In the Kingdom (Example 1), the direct exclusion core coincides with the two intuitive and focal allocations, as explained above. The strong core is not necessarily contained in the direct exclusion core, as demonstrated by Example 3 below.

2.2 The Exclusion Core

Direct exclusion blocking requires all blocking coalition members to strictly benefit from the new allocation. This requirement is seemingly constraining as many desirable reallocations require the acquiescence of unaffected third parties who coincidentally (co-)own a reassigned house. However, assuming indifferent parties always join a blocking coalition is misguided.

⁹The 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.

¹⁰If σ Pareto dominates μ , condition (b) in Definition 3 holds vacuously since $\sigma(i) \succeq_i \mu(i)$ for all i .

The resulting solution would be stronger than the strong core and vulnerable to the same criticisms concerning incentives. Instead, we can rationalize the cooperation of third parties by inductively extending the logic of exclusion. An example illustrates the idea.

Example 2. There are six agents and six houses. Agent i_k owns only house h_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 , i_1 must move to h_3 and i_3 must 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 can veto i_6 's assignment to h_3 as mandated by μ . Thereafter, h_3 is available for i_1 . The second move is not feasible since $h_2 = \mu(i_4) \succ_{i_4} \sigma(i_4)$ but $h_2 \notin \omega(\{i_1, i_3\})$. Thus, i_1 and i_3 cannot directly exclusion block μ .


Whereas i_1 and i_3 do not own h_2 , we argue that they enjoy a form of indirect control over it. 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. Agent i_2 is indifferent between μ and σ , but 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 , 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 σ . However, houses h_2 and h_3 are inaccessible since the coalition lacks leverage over those houses' owners. In fact, the pattern of exchange implied by σ insulates i_1 , i_2 , and i_3 from the coalition's direct and indirect exclusion power. Therefore, σ seems more compelling than μ as a final allocation in this economy.


Example 2 shows that the right to exclude can be a powerful, though subtle, stick. 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




(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.

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 formulation ensures that a 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 allocation μ is $\Omega(C|\omega, \mu) := \omega(\bigcup_{k=0}^{\infty} C_k)$ where $C_0 = C$ and $C_k = C_{k-1} \cup (\mu^{-1} \circ \omega)(C_{k-1})$ for every $k \geq 1$.

A coalition's extended endowment does not bestow upon the group more property rights. Rather, it better reflects agents' power when exclusion, or threats thereof, underpin interaction. By allowing a coalition to exclude others from houses in its extended endowment, we arrive at the following strengthening of Definition 3.

Definition 5. A non-empty coalition $C \subseteq I$ can *indirectly exclusion block* the allocation μ with allocation σ if (a) $\sigma(i) \succ_i \mu(i)$ for all $i \in C$ and (b) $\mu(j) \succ_j \sigma(j) \implies \mu(j) \in \Omega(C|\omega, \mu)$.

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

The following theorem is implied by results derived in Section 3 in a more general model.

Theorem 1. *For any economy $\langle I, H, \succ, \omega \rangle$ where ω satisfies (A1)–(A4), the exclusion core is not empty.*

Exclusion core allocations belong to the direct exclusion core. Thus, they are Pareto efficient and also belong to the weak core.

2.3 Private and Public Ownership

Both private and public ownership are common in practice and 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 $h \succ_{i_k} h_0$ for each i_k and $h \neq h_0$. They present an algorithm, attributed to David Gale, that selects a strong-core allocation in their market. We generalize this algorithm in Section 3 and summarize it here.

Algorithm 1 (Top Trading Cycles (TTC)). 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 unassigned house. As there is a finite number of agents and houses, there is at least one cycle of the form $h \rightarrow i \rightarrow \dots \rightarrow h' \rightarrow i' \rightarrow h$. (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). Ma (1994) proves that the TTC mechanism is the unique mechanism satisfying individual rationality,¹¹ Pareto efficiency, and strategy-proofness.¹² Furthermore,

¹¹If 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 .

¹²A (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.

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 antipode of a private-ownership economy is the *public-ownership economy* where all houses belong only to the social endowment, i.e. $\omega(C) = \emptyset$ for all $C \subsetneq I$ and $\omega(I) = H$. Hylland and Zeckhauser (1979) consider this assignment problem.¹³ This economy’s strong core equals the set of Pareto efficient allocations.

Proposition 3. *In a public ownership economy, the exclusion core equals the strong core. Equivalently, the exclusion core equals the Pareto frontier.*

In the above cases, the strong core was a subset of the exclusion core. However, the strong core is *not* necessarily a subset of the exclusion core, as confirmed by the next example. The example is an instance of a “house-allocation problem with existing tenants” (Abdulkadiroğlu and Sönmez, 1999) where some houses are privately owned and others belong only to the social endowment.

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 $C \subsetneq I$. The agents’ preferences are:

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

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 constitute the exclusion core. The coalition $C = \{i_1, i_2, i_4\}$ can directly (and, hence, indirectly) exclusion block μ with the allocation σ .

3 Relational Economies

In the previous section, we reinterpreted the endowment system ω as a distribution of exclusion rights. Despite the added generality, our analysis has so far glossed over much of

¹³See also Koopmans and Beckmann (1957).




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

the complexity of ownership rights in practice. These are often layered with caveats and qualifications. A variant of a prior example illustrates the practical limitations of a simple economy in describing these situations.

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 remains a property-less peasant. Two allocations are focal. In the first, j takes h_1 , k settles for h_2 , and i receives nothing. In the second, the kings swap houses. Both outcomes are efficient and equally plausible given the economy’s symmetry.

Perhaps surprisingly, this situation cannot be convincingly modeled as a simple economy. The problem concerns the correct endowment system. There are two natural options. The first places all houses only in the kings’ joint endowment: $\omega(i) = \omega(j) = \omega(k) = \emptyset$ and $\omega(\{j, k\}) = \{h_1, h_2\}$. 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 where k claims h_1 , i takes h_2 , and j —a king—is homeless belongs to the exclusion core.

The obvious alternative places both houses in each king’s personal endowment: $\omega(j) = \omega(k) = \{h_1, h_2\}$ and $\omega(i) = \emptyset$. This endowment system satisfies (A1)–(A3) but not (A4). The exclusion core is empty since every allocation can be blocked by the king who does not receive h_1 .

The Diarchy’s problems stem from the proposed endowment systems’ immutability and insensitivity to the agents’ identities and relationships. These can only be accounted for in a more contract-like arrangement: “If the peasant occupies a house, either king can evict him. However, a king cannot do likewise to a co-monarch.” Such conditional property rights are rare in economic models and differ from mere “co-ownership.” Each king’s exclusion rights are derived from his regal status, yet are qualified by the identity of a house’s occupant.

3.1 Priorities

To analyze situations with conditional and simultaneous claims, we appeal to the exclusion core, but we posit that endowments, i.e. exclusion rights, are endogenously determined. To further this idea, we first amend our definition of an economy. A *relational economy* $\langle I, H, \succ, \triangleright \rangle$ consists of agents, houses, preferences, and a priority structure. The first three components are defined as before. The new primitive is the *priority structure* $\triangleright = (\triangleright_h)_{h \in H}$, which is a family of orders that describe pre-existing social, legal, or economic relationships among agents in relation to the economy's goods. It may be formally codified by law or it may be informally set by social conventions, relative status, or historical context. In our understanding, priorities embody a minimal constraint that any property rights in a relational economy must respect in the following sense: If $i \triangleright_h j$, then i should enjoy rights no less than j with respect to house h .¹⁴

Formally, each $\triangleright_h$ is a strict partial order (i.e. an irreflexive and transitive relation) of the set of agents. We write $i \succeq_h j$ if $i \triangleright_h j$ or $i = j$. Many situations suggest natural priority structures.¹⁵ If a house is publicly owned, no agent has priority over others, i.e. $i \not\triangleright_h j$ for all i and j . If one agent $\triangleright_h$ -dominates all others, it is natural to call him the house's "owner." A diarchic structure occurs if two agents $\triangleright_h$ -dominate others, but not one another. More exotic cases are possible too. For instance, a case like $i \triangleright_h j \triangleright_h k$ and $\ell \triangleright_h k$ (with no other $\triangleright_h$ -dominances) may describe relationships within a family or a social group. We assume that the relation $\triangleright_{h_0}$ is empty and $i \triangleright_h \emptyset$ for all $i \in I$ and $h \in H$.

Priorities feature in many assignment problems, particularly those concerning student-school matching. We adopt the same terminology to highlight a technical parallel that will be evident below. However, our use of priorities to encode social, legal, or economic relationships in a (possibly decentralized) market differs from their use in a centralized assignment problem. For example, in a school-choice problem 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.¹⁶ Priorities do not have this meaning in our model. Second, priorities may define 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. Though

¹⁴See Campbell (1992) on how hierarchical relations may qualify property rights.

¹⁵Ehlers and Erdil (2010) model some of these situations with non-strict partial orders.

¹⁶A student feels justified envy if he prefers to attend a school that enrolled a lower-priority student.

still distinct, our use of priorities is closer in spirit to this second meaning.

Our use of priorities bears some similarity to Piccione and Rubinstein's (2007) strength relation in their model of a "jungle economy." Their strength relation 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 the prevailing endowment system, which we turn to next.

3.2 Endowments in Relational Economies

In a relational economy, exclusion rights should reflect the context conveyed by the priority structure. To do so convincingly, they may need to be qualified by the prevailing allocation, as suggested by the Diarchy (Example 4). We use the term *conditional endowment* to emphasize when this occurs. Noting these desiderata, we propose three natural definitions of an endowment system in a relational economy. These lead to the strong, weak, and unconditional exclusion cores.

The Strong Exclusion Core The simplest definition of an endowment system in a relational economy places house h in a coalition's (conditional) endowment if one of its members $\succeq_h$ -dominates that house's occupant. We define a relational economy's *weak conditional endowment system* at μ , $\omega_\mu: 2^I \rightarrow 2^H$, as follows. For every $h \in H$ and $C \subseteq I$, $h \in \omega_\mu(C)$ if and only if $i \succeq_h \mu^{-1}(h)$ for some $i \in C$. Weak conditional endowments plug 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 point (b).

Definition 6. A non-empty coalition $C \subseteq I$ can *indirectly exclusion block* the allocation μ with allocation σ given ω_μ if (a) $\sigma(i) \succ_i \mu(i)$ for all $i \in C$ and (b) $\mu(j) \succ_j \sigma(j) \implies \mu(j) \in \Omega(C|\omega_\mu, \mu)$.

The allocation μ belongs to a relational economy's *strong exclusion core* if and only if μ cannot be indirectly exclusion blocked given ω_μ .

The intuitive derivation of ω_μ gives the strong exclusion core great appeal. Regrettably, the strong exclusion core can be empty.

Example 5. Let $I = \{i, j, k\}$ and $H = \{h_1, h_2\}$. Suppose $i \triangleright_{h_1} j \triangleright_{h_1} k$, $k \triangleright_{h_2} j \triangleright_{h_2} i$ and

$$\succ_i: h_2, h_1 \quad \succ_j: h_1 \quad \succ_k: h_1, h_2 \ .$$

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 given 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 $\triangleright$ is *acyclic* if for all $h \in H$ and agents i, j , and k ,

$$[i \triangleright_h j \ \& \ i \not\triangleright_h k] \implies k \triangleright_{h'} j \text{ for all } h' \neq h, h_0. \quad (1)$$

Our definition of acyclicity is specifically phrased to accommodate incomplete relations; however, (1) reduces to Ergin (2002) acyclicity when each $\triangleright_h$ is a linear order of all agents.¹⁷ It is related to strong acyclicity, which was proposed by Ehlers and Erdil (2010) as an extension of Ergin's (2002) definition to non-strict priority rankings. See also Kesten (2006). The next result is implied by Lemma 2 and Theorem 3, which are stated below.

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

Acyclic priority structures are common. If house h is privately owned there is an agent i , the house's owner, such that $i \triangleright_h j$ for all $j \neq i$ and $j \not\triangleright_h k$ for all $j, k \in I \setminus \{i\}$. Conversely, if house h is publicly owned, $i \not\triangleright_h j$ for all i and j . Any economy featuring a combination of privately and publicly owned houses has an acyclic priority structure.¹⁸ The Diarchy (Example 4) can also be modeled with an acyclic priority structure: both j and k $\triangleright_h$ -dominate i for each $h \in H$, but not each other. Its strong exclusion core contains only the two focal allocations where the kings claim both houses.

The Weak Exclusion Core The strong exclusion core presumes a fairly liberal distribution of exclusion rights. A more conservative distribution may be preferable when the priority structure is not acyclic. To simplify notation, we say that $C \supseteq_h j$ if and only if $i \supseteq_h j$ for some $i \in C$. And, we write $C \triangleright_h j$ if $C \supseteq_h j$ but $j \notin C$. We define a relational economy's

¹⁷An Ergin (2002) cycle occurs if for distinct houses h and h' and distinct agents i, j , and k , $k \triangleright_h i \triangleright_h j \triangleright_{h'} k$. A priority structure is Ergin (2002) acyclic if it does not contain an Ergin (2002) cycle. If $\triangleright_h$ is a linear order for each h , then (1) becomes $k \triangleright_h i \triangleright_h j \implies k \triangleright_{h'} j$. Hence, an Ergin (2002) cycle cannot occur.

¹⁸To confirm this fact, note that the antecedent in (1), $i \triangleright_h j \ \& \ i \not\triangleright_h k$, is never satisfied.

strong conditional endowment system at μ , $\omega_\mu^*: 2^I \rightarrow 2^H$, as follows. For every $h \in H$ and $C \subseteq I$, $h \in \omega_\mu^*(C)$ if and only if

(a) $C \supseteq_h \mu^{-1}(h)$ and

(b) $[C \supset_h \mu^{-1}(h) \ \& \ C \not\supseteq_h k] \implies k \supset_{h'} \mu^{-1}(h)$ for all $h' \neq h, h_0$.

Condition (a) says that if $h \in \omega_\mu^*(C)$, then $i \supseteq_h \mu^{-1}(h)$ for some $i \in C$. This is identical to the condition defining the weak conditional endowment system above. Condition (b) ensures that the socially recognized exclusion rights, as defined by ω_μ^* , do not inherit the problematic cycles found in the priority structure. Intuitively, it can be interpreted as a “conditional acyclicity” requirement noting its resemblance to (1).¹⁹ The strong conditional endowment system does not recognize (potential) exclusion rights that are vulnerable to cyclic challenges or counterclaims at the prevailing allocation.

Replacing “ ω_μ ” with “ ω_μ^* ” in Definition 6 leads to the corresponding version of the exclusion core. The allocation μ belongs to a relational economy’s *weak exclusion core* if and only if μ cannot be indirectly exclusion blocked given ω_μ^* . Since $\omega_\mu^*(C) \subseteq \omega_\mu(C)$ for all C , the strong exclusion core is a subset of the weak exclusion core.

Lemma 2. *If the priority structure is acyclic, the weak and strong exclusion cores coincide.*

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

In Example 5, the strong exclusion core was empty. The unique weak exclusion core allocation assigns i to h_2 and k to h_1 .

The Unconditional Exclusion Core Both weak and strong conditional endowments are functions of the prevailing allocation μ . Whether a coalition has exclusion rights with respect to house h depends on its occupant. Though warranted in situations like the Diarchy (Example 4), we can naturally expunge this conditionality. We define a relational economy’s *unconditional endowment system*, $\omega^{**}: 2^I \rightarrow 2^H$, as follows. For all $h \in H$ and $C \subseteq I$, $h \in \omega^{**}(C)$ if and only if C includes every $\supset_h$ -maximal agent in the economy.²⁰ By replacing “ ω_μ ” with “ ω^{**} ” in Definition 6, we define the *unconditional exclusion core* of a relational economy as the set of allocations that cannot be indirectly exclusion blocked given ω^{**} . For all μ and C , $\omega^{**}(C) \subseteq \omega_\mu^*(C) \subseteq \omega_\mu(C)$.²¹ Thus, a relational economy’s unconditional exclusion

¹⁹For added intuition, suppose $C = \{i\}$. In this case, part (b) becomes $[i \supset_h \mu^{-1}(h) \ \& \ i \not\supseteq_h k] \implies k \supset_{h'} \mu^{-1}(h)$ for all $h' \neq h, h_0$. Now, agent $\mu^{-1}(h)$ plays the role of agent j from (1).

²⁰Agent $i \in I$ is $\supset_h$ -maximal if there is no $j \in I$ such that $j \supset_h i$.

²¹If $h \in \omega^{**}(C)$, then $C \supseteq_h \mu^{-1}(h)$. Point (b) in the definition of $\omega_\mu^*(\cdot)$ holds since $C \not\supseteq_h k$ is impossible.

core is not empty and contains its weak and strong exclusion cores. The unconditional exclusion core connects simple and relational economies, as explained in Section 3.4.

3.3 Generalized Top Trading Cycles

We prove Theorem 3 in Appendix A. Our proof is constructive and relies on the algorithm introduced below. The algorithm identifies a weak exclusion core allocation. When the priority structure is acyclic, Lemma 2 implies this algorithm's output is a strong exclusion allocation, thus proving Theorem 2. Our algorithm builds upon several precursors. 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.²² 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). Our algorithm is also a direct descendant of the TTC algorithm with a tie-breaker of coarse priorities, as applied to the school-choice problem (Abdulkadiroğlu and Sönmez, 2003; Abdulkadiroğlu et al., 2009).²³ Our algorithm reduces to each of the above cases when the environment is appropriately restricted. This connection binds the exclusion core to well-known applications.

Algorithm 2 (Generalized Top Trading Cycles (GTTC)). Given $\langle I, H, \succ, \triangleright \rangle$, let $\tilde{\triangleright}_h$ be a linear order of agents such that $i \triangleright_h j \implies i \tilde{\triangleright}_h j$ for each $h \in H$. We call $\tilde{\triangleright}$ an extension of $\triangleright$.²⁴ 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 as 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 $H^t \cup \{h_0\}$. For each $h \in H^t$, draw an arc from h to the $\tilde{\triangleright}_h$ -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 $\tilde{H}^t = \emptyset$.

²²In 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.

²³The main difference is our algorithm's accommodation of a more general class of priority structures than typically encountered in school-choice problems.

²⁴An extension $\tilde{\triangleright}$ always exists by the Szpilrajn Extension Theorem. It is tempting to view $\tilde{\triangleright}$ as a priority structure supplemented by a tie-breaking rule (Ehlers, 2014). While compatible with our model, we hesitate to emphasize this interpretation. If $i \not\triangleright_h j$ and $j \not\triangleright_h i$, agents i and j are not necessarily "equal" in our economy. For instance, it may be true that $i \triangleright_h k$ but $j \not\triangleright_h k$.

- (b) Otherwise, the constructed graph contains at least one cycle. Choose any cycle and carry out the implied assignments. That is, if $i \rightarrow h$ 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.

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

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

An example in Appendix B illustrates the operation of Algorithm 2. 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. Furthermore, the algorithm is strategy-proof. This fact is a corollary to results by Roth (1982), Abdulkadiroğlu and Sönmez (1999), Roth et al. (2004), and (in particular) Abdulkadiroğlu and Sönmez (2003), who extend the 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.

Algorithm 2 is parameterized by the extension $\tilde{\succ}$. By varying the extension, Algorithm 2 identifies a family of exclusion core outcomes.

Theorem 4. *Every strong exclusion core allocation in the relational economy $\langle I, H, \succ, \triangleright \rangle$ can be identified by Algorithm 2 with some extension of $\triangleright$.*

While Algorithm 2 can find all strong exclusion core allocations, it cannot find all weak exclusion core allocations.²⁵ Allocations identified by Algorithm 2 need not belong to the strong exclusion core when the economy's priority structure is not acyclic (see Example 5). The next corollary follows from Lemma 2 and Theorems 3 and 4.

Corollary 1. *Denote the weak exclusion core by WEC, the strong exclusion core by SEC, and the range (over all extensions 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, $SEC = GTTC = WEC$.*

We can highlight several further properties of the strong and weak exclusion cores. For instance, they enjoy a natural stability property when the priority structure is acyclic.

²⁵Consider the following example. 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 : $j \triangleright_{h_1} i$ and $k \triangleright_{h_1} i$. Agent i owns h_2 and h_3 : $i \triangleright_h j$ and $i \triangleright_h k$ for $h \in \{h_2, h_3\}$. Otherwise, the agents are not $\triangleright$ -comparable. There are three weak exclusion core allocations. In two allocations, i takes h_2 and either j or k claims h_1 . The third allocation— $\mu(i) = h_1$, $\mu(j) = h_2$, and $\mu(k) = h_3$ —cannot be identified by Algorithm 2.

Proposition 4. *The strong (or weak) exclusion core of a relational economy with an acyclic priority structure is stable in the sense of von Neumann and Morgenstern (1944).²⁶*

The exclusion core also exhibits intuitive comparative statics with respect to $\triangleright$. Changes in $\triangleright$ may reflect changing legal or social norms. We call $\triangleright'$ a *coarsening* of $\triangleright$ if for all $h \in H$ and $i, j \in I$, $i \triangleright'_h j \implies i \triangleright_h j$. Intuitively, $\triangleright'$ coincides with $\triangleright$ except some hierarchical relations among agents are possibly expunged.

Proposition 5. *If $\triangleright'$ is a coarsening of $\triangleright$, the strong/weak/unconditional exclusion core of $\langle I, H, \succ, \triangleright' \rangle$ contains the strong/weak/unconditional exclusion core of $\langle I, H, \succ, \triangleright \rangle$.*

3.4 Simple and Relational Economies

We introduced relational economies to offer a more sophisticated model of collective or qualified ownership. We conclude this section by linking relational economies to the simple economies of Section 2 where an endowment system, rather than a priority structure, is the primitive.

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 ω be an endowment system 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 or publicly owned, 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 2.3. The following result, due to Sönmez (1999), helps connect 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 φ for the class of private-ownership economies.*

²⁶A set of allocations A is *von Neumann-Morgenstern stable* if it is (a) internally stable: every $\mu \in A$ is not “dominated” by any $\sigma \in A$; and, (b) externally stable: every $\mu \notin A$ is “dominated” by some $\sigma \in A$. In our case, σ “dominates” μ if some coalition can indirectly exclusion block μ with σ given ω_μ (or ω_μ^*).

- (a) *The strong core of any such economy is either empty or a singleton.*
- (b) *If the strong core of that economy is not empty, its unique element is identified by φ .*

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.*

Corollaries 2 and 3 together imply Proposition 1. Corollary 3 also implies the coincidence of the exclusion and strong cores in Shapley and Scarf's (1974) economy (Proposition 2).

An economy with both private and public ownership is the house-allocation problem with existing tenants (Abdulkadiroğlu and Sönmez, 1999). In this problem, every house is either owned by exactly one agent or belongs to the social endowment. 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 and Lemma 3.

Corollary 4. *In the house-allocation problem with existing tenants, the exclusion core coincides with the set of all possible allocations identified by the YRMH-IGYT mechanism.*

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

4 Production Economies

By defining the distribution of exclusion rights, i.e. the endowment system, over existing and potential goods, we can apply the exclusion core concept to production economies. Now a coalition may modify a firm's production decision by leveraging the dependence of the firm's owners on the coalition's extended endowment.

Consider the following generalization of the simple economy from Section 2. A simple economy with production is a tuple $\langle (I, F), (H_0, \mathcal{H}), \succ, \omega \rangle$. The set of agents I is augmented by a finite set of firms $F = \{f_1, \dots, f_{n'}\}$. The set of houses is replaced by an existing housing stock H_0 and a collection $\mathcal{H} = \{\mathcal{H}_1, \dots, \mathcal{H}_{n'}\}$ of production sets, one for each firm. The

production set of firm f_k is a non-empty menu of production plans, $\mathcal{H}_k = \{H_{k,1}, \dots, H_{k,m_k}\}$.²⁷ The set of houses produced by firm f_k under plan j is $H_{k,j}$. Each firm's production technology satisfies the following assumption:

(A5) Differentiated production plans: If $(k, j) \neq (k', j')$, then $H_{k,j} \cap H_{k',j'} = \emptyset$.

Condition (A5) captures the idea that a firm faces tradeoffs in production due to a resource constraint. To produce a new house, a firm must modify all other houses that it produces to secure the required inputs. Even minor modifications lead to distinct outputs.²⁸

The set of all possible houses in the economy is $H := H_0 \cup \left(\bigcup_{k,j} H_{k,j}\right)$. Each agent's strict preference $\succ_i$ is defined over $H \cup \{h_0\}$. As before, $h_0 \notin H$ is an outside option.

Finally, $\omega: 2^I \rightarrow 2^H$ is an endowment system satisfying (A1)–(A4). Recall that assumption (A4) says that each house $h \in H$ has a minimal controlling coalition C^h . We extend this idea to production processes by assuming

(A6) Production non-contestability: For $k \neq 0$, $\{h, h'\} \subseteq \bigcup_j H_{k,j} \implies C^h = C^{h'}$.

Assumption (A6) says that all goods that can be produced by firm f_k have the same minimal controlling coalition. In everyday language, we may call this coalition the firm's owners. It is notable that the set $\omega(C)$ may include goods that exist in the market and goods that *can* exist, if they are produced.

An *outcome* (μ, ρ) consists of an allocation μ and a production plan ρ . An allocation $\mu: I \rightarrow H \cup \{h_0\}$ is defined as before. A production plan $\rho: F \rightarrow \bigcup_k \mathcal{H}_k$ is a function such that $\rho(f_k) \in \mathcal{H}_k$ for each $f_k \in F$. It defines what each firm produces. Let $H_\rho := H_0 \cup \left(\bigcup_k \rho(f_k)\right)$ be the set of available houses given ρ . An outcome (μ, ρ) is *feasible* if $\mu(I) \subseteq H_\rho \cup \{h_0\}$.

The only novelties in the next definition are the qualifications concerning feasibility.

Definition 7. A non-empty coalition $C \subseteq I$ can *indirectly exclusion block* the feasible outcome (μ, ρ) with the feasible outcome (σ, ρ') if (a) $\sigma(i) \succ_i \mu(i)$ for all $i \in C$ and (b) $\mu(j) \succ_j \sigma(j) \implies \mu(j) \in \Omega(C|\omega, \mu)$.

The *exclusion core of a simple economy with production* is the set of feasible outcomes that cannot be indirectly exclusion blocked. Note that Definition 7 allows a blocking coalition to change the economy-wide production plan, including altering the output of firms that it

²⁷For instance, it is natural to assume that inaction is possible, i.e. $\emptyset \in \mathcal{H}_k$. When $\mathcal{H}_k = \{\emptyset\}$ for each k , a simple economy with production reduces to the model of Section 2.

²⁸For example, a worker might spend less time working on a house, thus slightly degrading its quality.

strong exclusion core, which may be empty if priority cycles induce too many conflicting claims. These conflicts are settled by the weak exclusion core, which tightens the criteria for a coalition to “own” a house. Both strong and weak exclusion cores rely on conditional endowments, where the distribution of exclusion rights changes with the prevailing allocation. The unconditional exclusion core dismisses this conditionality and provides a logical connection to the paradigm of simple economies where priorities play no role. All four exclusion core variants coincide when each house is either privately owned or part of the social endowment, two familiar possibilities.

Our study of property rights with the aid of a house-exchange economy allows our analysis to bridge two 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 conclude by addressing each contribution in turn.

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, Algorithm 2 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 in their generalization of Pápai’s (2000) model. An inheritance structure defines how unassigned houses are inherited or transferred during a multi-step assignment process. Svensson and Larsson (2005) introduce endowment rules, which are similar.

The extensions $\tilde{\succ}$ in Algorithm 2 play the role of inheritance structures in our argument. Like inheritance structures, $\tilde{\succ}$ defines contingent control within a sequential assignment process. Despite this similarity, these extensions are purely technical devices in our analysis and characterize neither property nor endowments. Endowments in our model describe a distribution of exclusion rights and are independent of any particular trading protocol.

Several studies rely on hierarchical exchange mechanisms to propose new variants of the core. These definitions combine weak blocking (Definition 1) with alternative specifications 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. Svensson and Larsson (2005), and later Tang and Zhang (2016), define endowments in terms of houses a coalition would have feasibly inherited during trade given a prevailing endowment rule or inheritance structure. Unlike these studies, our definitions avoid the problematic incentives underlying weak blocking (see Section 1). Our derivations

and interpretations of endowments are distinct as well.

Many variants of Shapley and Scarf's (1974) model 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 assumed a strict preference domain, integral endowments, and deterministic final outcomes. Each of these assumptions has been relaxed by many authors.²⁹ Farsighted solutions have also been considered (Klaus et al., 2010). We defer investigating these extensions to future research.

Walrasian equilibrium is another solution commonly applied to exchange economies. The possible absence of personal endowments precludes the application of standard price equilibrium definitions to our setting. Richter and Rubinstein (2015) introduce the notion of a "primitive equilibrium," which does not rely on budget sets nor endowments. Instead, they observe that equilibria 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), Merrill and Smith (2001b), and Klick and Parchomovsky (2017), among many others, elaborate on these points in detail. Though property rights are a touchstone for our analysis, the questions we consider are distinct from the bilateral externalities examined by Coase (1960) or the literature on contracts and control following Grossman and Hart (1986).

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. This division corresponds to different conceptual understandings of property among legal scholars, which can be framed as a spectrum running between *in rem* and *in personam* paradigms (Merrill and Smith, 2001a,b).³⁰ The model of Section 2 can be interpreted as embodying an *in rem* interpretation of property. A coalition's exclusion rights are universal, valid against all others, and embodied in their ownership of a specific thing.

²⁹For 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 (2018) study fractional endowments. Abdulkadiroğlu and Sönmez (1998) and Carroll (2014) examine random outcomes.

³⁰Merrill 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.

Arruñada (2012) argues that *in rem* rights allow for impersonal exchange to be possible.

The conditional endowments introduced in Section 3 channel an *in personam* interpretation of property. Under this paradigm, an agent's rights are defined by his relationships with others, which we model with priorities. Exclusion rights are constrained by the priority structure, directed at specific persons, and qualified by the prevailing allocation, almost in a contract-like manner. An agent's ability to participate in such markets anonymously is necessarily limited, and often impossible.

Our analysis provides a new formal comparison of *in rem* and *in personam* paradigms. In fact, Lemma 3 provides a connection between them. Nevertheless, our model's sparsity masks some notable differences. For example, Merrill and Smith (2001a) argue that the information burden accompanying *in rem* and *in personam* rights is different. Though absent from our analysis, this consideration affects a market's operation and scalability in practice. Further investigation of this distinction is likely to be fruitful.

A Proofs

Proof of Lemma 1. Theorem 1 implies the direct exclusion core's non-emptiness. Suppose coalition C can strongly block the allocation μ with σ . This implies $\sigma(C) \subseteq \omega(C) \cup \{h_0\}$. It is sufficient to show that C can directly exclusion block μ . Consider the allocation $\hat{\sigma}$ where $\hat{\sigma}(i) = \sigma(i)$ for all $i \in C$, $\hat{\sigma}(i) = h_0$ if $i \notin C$ and $\mu(i) \in \sigma(C) \setminus \{h_0\}$, and $\hat{\sigma}(i) = \mu(i)$ otherwise. Observe that $\hat{\sigma}(i) \succ_i \mu(i)$ for all $i \in C$. Furthermore, if $\mu(j) \succ_j \hat{\sigma}(j)$, then $\mu(j) \in \sigma(C) \setminus \{h_0\} \subseteq \omega(C)$. Hence, C can directly exclusion block μ with $\hat{\sigma}$. $\square$

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 . Without loss of generality, $\sigma(i) \succ_i \mu(i)$ if and only if $i \in C$. Moreover, if $i \in C$, then $\sigma(i) = \mu(j) \neq h_0$ for some $j \in I$. Else, the allocation μ would not be Pareto efficient, a contradiction.

To derive a contradiction, we provide an algorithm that identifies a coalition C' that can weakly block μ . Choose $i^0 \in C$ and consider house $\sigma(i^0) = h^0$. This gives the sequence (i^0, h^0) , the algorithm's initial input. In step $k \geq 0$ the algorithm proceeds as follows with input $(i^0, h^0, \dots, i^k, h^k)$, which is a sequence of distinct agents and houses such that: (i) $h^\ell \in \omega(i^{\ell+1})$ for all $\ell < k$ and $h^k \neq h_0$ (the outside option), (ii) if $i^\ell \in C$, then $h^\ell = \sigma(i^\ell)$, and (iii) if $i^\ell \notin C$, then $h^\ell = \mu(i^\ell)$. (Note that (i^0, h^0) satisfies these conditions.)

Step k. Given $(i^0, h^0, \dots, i^k, h^k)$, one of four cases must apply.

Case 1. If $h^k \in \omega(\{i^0, \dots, i^k\})$, the coalition $C' = \{i^0, \dots, i^k\}$ can weakly block μ with any allocation $\hat{\sigma}$ such that $\hat{\sigma}(i^{k'}) = h^{k'}$. This is because, by construction, $\hat{\sigma}(C') = \{h^0, \dots, h^k\} \subseteq \omega(C')$, $\hat{\sigma}(i^\ell) \succeq_{i^\ell} \mu(i^\ell)$ for all $i^\ell \in C'$, and $\hat{\sigma}(i^0) \succ_{i^0} \mu(i^0)$ for $i^0 \in C'$.

Otherwise, if $h^k \notin \omega(\{i^0, \dots, i^k\})$, then there exists some $i^{k+1} \notin \{i^0, \dots, i^k\}$ such that $h^k \in \omega(i^{k+1})$. There are three remaining possibilities.

Case 2. If $i^{k+1} \notin C$ and $\mu(i^{k+1}) = h_0$, then we are done. Specifically, set $h^{k+1} = \mu(i^{k+1}) = h_0$. This gives the sequence $(i^0, h^0, \dots, i^k, h^k, i^{k+1}, h^{k+1})$. The coalition $C' = \{i^0, \dots, i^{k+1}\}$ can weakly block μ with any allocation $\hat{\sigma}$ such that $\hat{\sigma}(i^\ell) = h^\ell$ for each $i^\ell \in C'$.

Case 3. If $i^{k+1} \in C$, set $h^{k+1} = \sigma(i^{k+1})$. There are two sub-cases.

- (a) If $h^{k+1} \notin \{h^0, \dots, h^k\}$, append (i^{k+1}, h^{k+1}) to the original sequence to get $(i^0, h^0, \dots, i^{k+1}, h^{k+1})$, which satisfies conditions (i)–(iii). Go to step $k + 1$.
- (b) Else, if $h^{k+1} = h^{k'}$ for some $k' \leq k$, then we are done.

Specifically, consider the subsequence starting at $k' + 1$, $(i^{k'+1}, h^{k'+1}, \dots, i^{k+1}, h^{k+1})$. Now consider the coalition $C' = \{i^{k'+1}, \dots, i^{k+1}\}$ and an allocation $\hat{\sigma}$ such that $\hat{\sigma}(i^\ell) = h^\ell$ for each $i^\ell \in C'$. By construction, $h^\ell \in \omega(i^{\ell+1})$ for all $\ell = k' + 1, \dots, k$ and $h^{k+1} = h^{k'} \in \omega(i^{k'+1})$. Hence, $\hat{\sigma}(C') = \{h^{k'+1}, \dots, h^{k+1}\} \subseteq \omega(C')$. Moreover, $\hat{\sigma}(i^\ell) \succeq_{i^\ell} \mu(i^\ell)$ for all $i^\ell \in C'$ and $\hat{\sigma}(i^{k+1}) \succ_{i^{k+1}} \mu(i^{k+1})$. Thus, coalition C' can weakly block μ .

Case 4. If $i^{k+1} \notin C$ and $\mu(i^{k+1}) \in H$, set $h^{k+1} = \mu(i^{k+1})$. There are two sub-cases.

- (a) If $h^{k+1} \notin \{h^0, \dots, h^k\}$, append (i^{k+1}, h^{k+1}) to the original sequence to get $(i^0, h^0, \dots, i^{k+1}, h^{k+1})$, which satisfies conditions (i)–(iii). Go to step $k + 1$.
- (b) Else, if $h^{k+1} = h^{k'}$ for some $k' \leq k$, then we are done.

First, observe that $\sigma(i^{k'}) = h^{k'}$. (If $\mu(i^{k'}) = h^{k'}$, then μ would not be a valid allocation since $\mu(i^{k+1}) = \mu(i^{k'})$ and $i^{k+1} \neq i^{k'}$.) Since $h^{k+1} = h^{k'} = \mu(i^{k+1}) \succ_{i^{k+1}} \sigma(i^{k+1})$, $h^{k'} \in \Omega(C|\omega, \mu)$. Thus, starting at $h^{k'}$, the subsequence $(h^{k'}, i^{k'+1}, \dots, h^k, i^{k+1})$ must contain an agent $j \in C$. If this was not the case, $\mu(i^\ell) = h^\ell$ for each $\ell = k' + 1, \dots, k + 1$, which implies $h^{k'} \notin \Omega(C|\omega, \mu)$.

Now consider the coalition $C' = \{i^{k'+1}, \dots, i^{k+1}\}$ and an allocation $\hat{\sigma}$ such that $\hat{\sigma}(i^\ell) = h^\ell$ for each $i^\ell \in C'$. By construction, $h^\ell \in \omega(i^{\ell+1})$ for all $\ell = k' + 1, \dots, k$ and $h^{k+1} = h^{k'} \in \omega(i^{k'+1})$. Hence, $\hat{\sigma}(C') = \{h^{k'+1}, \dots, h^{k+1}\} \subseteq \omega(C')$.

Moreover, $\hat{\sigma}(i^\ell) \succeq_{i^\ell} \mu(i^\ell)$ for all $i^\ell \in C'$, and $\hat{\sigma}(j) \succ_j \mu(j)$ for at least one $j \in C' \cap C$. Thus, coalition C' can weakly block μ .

As there is a finite number of agents, the preceding algorithm must terminate. It does so only in Case 1, Case 2, Case 3b, and Case 4b after identifying a blocking coalition C' . $\square$

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. By Proposition 1, the unique strong-core allocation belongs to the exclusion core. Conversely, suppose μ is an exclusion core allocation. This allocation can be represented as a directed graph where each $h_k \in H$ points to its owner, say $h_k \rightarrow i_k$, and each agent points to his assignment, i.e. $i_k \rightarrow \mu(i_k)$. As all houses are acceptable and $|I| = |H|$, $\mu(I) = H$. The resulting graph partitions $I \cup H$ into disjoint cycles $K^1, \dots, K^T$. Observe that $i \in K^t$ if and only if $\omega(i) \in K^t$ if and only if $\mu(i) \in K^t$. 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 $C = \{i^0, \dots, i^{k-1}\}$ 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 σ assigns the houses in $\omega(C)$ cyclically among the members of C . That is,

$$i^0 \rightarrow h^1 \rightarrow \dots \rightarrow i^{k-1} \rightarrow h^0 \quad (\text{A.1})$$

where $h^\ell = \omega(i^\ell)$ and $\sigma(i^\ell) = h^{\ell+1 \pmod k}$.³¹

Let $C' = \{i \in C | \sigma(i) \succ_i \mu(i)\}$ and consider the allocation $\hat{\sigma}$ where $\hat{\sigma}(i) = \sigma(i)$ for all $i \in C'$, $\hat{\sigma}(i) = h_0$ if $i \notin C'$ and $\mu(i) \in \sigma(C')$, and $\hat{\sigma}(i) = \mu(i)$ otherwise.

By construction, $\hat{\sigma}(i) \succ_i \mu(i)$ for all $i \in C'$. Pick any j such that $\mu(j) \succ_j \hat{\sigma}(j)$. This implies $\mu(j) \in \hat{\sigma}(C') = \sigma(C')$. Let K^j be the cycle, as defined above, for which $j \in K^j$ and $\mu(j) \in K^j$. Furthermore, let $i \in C'$ be such that $\hat{\sigma}(i) = \mu(j)$. Without loss of generality, suppose $i = i^0$, according to the enumeration in (A.1). There are two cases. First, if $i^0 \in K^j$, then $\mu(j) \in \bigcup_{k=0}^{\infty} (\omega \circ \mu^{-1})^k(\omega(i^0))$. Alternatively, and second, if $i^0 \notin K^j$, then there exists $i^t \in C' \subseteq C$ such that $i^t \in K^j$. To see why, note that $\mu(j) \in \sigma(C') \subseteq \sigma(C) = \omega(C)$. Thus $\mu(j)$ points to some $j' \in C \cap K^j$. If $j' \in C'$, we are done. Otherwise, we must have $\mu(j') = \sigma(j')$, which in turn implies that $\mu(j')$ must point to another member of

³¹If σ induces multiple cycles, they are necessarily disjoint. Without loss of generality we may focus on any one of them involving an agent $i \in C$ such that $\sigma(i) \succ_i \mu(i)$.

C . Eventually, if none of the members of C forming this chain is also in C' , then for all $i \in (i^1, \dots, i^{k-1})$, $\sigma(i) = \mu(i)$. This implies $\mu(i^{k-1}) = h^0 \in K^j$. But, $h^0 = \omega(i^0)$ and hence $i^0 \in K^j$, which is a contradiction.

Together, the preceding cases imply that $\mu(j) \in \bigcup_{k=0}^{\infty} (\omega \circ \mu^{-1})^k(\omega(C'))$. As the choice of j was arbitrary, we conclude that coalition C' can indirectly exclusion block μ with $\hat{\sigma}$, which contradicts μ being an exclusion core allocation. $\square$

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

Proof of Lemma 2. It is sufficient to verify $\omega_\mu(C) \subseteq \omega_\mu^*(C)$ when the priority structure is acyclic. Let $h \in \omega_\mu(C)$. Thus, there exist $i \in C$ such that $i \succeq_h \mu^{-1}(h)$. Therefore, $C \succeq_h \mu^{-1}(h)$. Now suppose $C \triangleright_h \mu^{-1}(h)$ and $C \not\triangleright_h k$. This implies there exists $i \in C$ such that $i \triangleright_h \mu^{-1}(h)$ and $i \not\triangleright_h k$. Acyclicity implies that $k \triangleright_{h'} \mu^{-1}(h)$ for all $h' \neq h, h_0$. Hence, $h \in \omega_\mu^*(C)$. $\square$

Proof of Theorem 3. Let μ be the assignment identified by Algorithm 2 for some extension $\tilde{\succ}$ of $\triangleright$. We note that Algorithm 2 constructs μ sequentially by removing sets of agents $(\tilde{I}^1, \tilde{I}^2, \dots)$ and associated houses $(\tilde{H}^1, \tilde{H}^2, \dots)$. To derive a contradiction, suppose coalition C can indirectly exclusion block μ with σ given ω_μ^* . Thus, and without loss of generality, $\sigma(i) \succ_i \mu(i)$ if and only if $i \in C$ and

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

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 remains in 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 $t < t_i$. $\diamond$

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. Fix $h \in \tilde{H}^t$, but assume $t_i > t$ for all $i \in J$. That is, each agent in J is assigned (strictly) *after* house h by Algorithm 2 given $\tilde{\succ}$.

First suppose that $\mu(j) = h$ and j is $\tilde{\succ}_h$ -maximal at step t . Since $h \in \omega_\mu^*(J)$, there exists $i \in J$ such that $i \succeq_h \mu^{-1}(h) = j$. As every agent in J is assigned after step t , $i \neq j$. Thus, $i \triangleright_h \mu^{-1}(h) = j$ and therefore, as $\tilde{\succ}_h$ extends $\triangleright_h$, $i \tilde{\succ}_h j$, which contradicts j being $\tilde{\succ}_h$ -maximal at step t .

Suppose instead that house h is assigned as part of a cycle that involves multiple agents and houses. Let $\cdots \rightarrow h' \rightarrow j^0 \rightarrow h \rightarrow j^1 \rightarrow \cdots$ be part of this cycle, where $\mu(j^0) = h$, j^0 (resp. j^1) is $\tilde{\succ}_{h'}$ -maximal (resp. $\tilde{\succ}_h$ -maximal) among agents in I^t . (Recall that I^t is the set of agents who remain in the market at the beginning of step t of Algorithm 2.) Clearly, $j^0 \not\triangleright_h j^1$. Since $h \in \omega_\mu^*(J)$, $J \succeq_h j^0$. Since every agent in J is assigned a house at a later step, $j^0 = \mu^{-1}(h) \notin J$. This implies $J \triangleright_h j^0$. Since $h \rightarrow j^1$ at step t , $i \not\triangleright_h j^1$ for every $i \in J$. This implies $J \not\triangleright_h j^1$. Because $h \in \omega_\mu^*(J)$, $[J \triangleright_h j^0 \ \& \ J \not\triangleright_h j^1]$ imply $j^1 \triangleright_{h'} \mu^{-1}(h) = j^0$. So $j^1 \tilde{\succ}_{h'} j^0$, which is a contradiction since j^0 was $\tilde{\succ}_{h'}$ -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, 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 given $\tilde{\succ}$ that contains an agent j^0 such that $\mu(j^0) \succ_{j^0} \sigma(j^0)$. It has to be the case that $\mu(j^0) \in H$. Suppose this cycle is removed at step t_0 . Let $\tilde{I}^{t_0}$ and $\tilde{H}^{t_0}$ be the sets of agents and houses, respectively, removed from the market at this step.

Claim 4. If $i \in C$, then $i \in \bigcup_{t > t_0} \tilde{I}^t$, i.e. $C \cap (\bigcup_{t \leq t_0} \tilde{I}^t) = \emptyset$.

Proof of Claim 4. First, suppose there exists $j^1 \in C \cap (\bigcup_{t \leq t_0} \tilde{I}^t)$. If there are multiple such agents, choose one who is assigned during the earliest cycle occurring in Algorithm 2. Let the cycle containing j^1 be removed at step $t_1 \leq t_0$. By Claim 1, $\sigma(j^1) \in \tilde{H}^{t_2}$ for some $t_2 < t_1 \leq t_0$. Let $j^2 \in \tilde{I}^{t_2}$ be such that $\mu(j^2) = \sigma(j^1)$. In particular, this implies that $\mu(j^2) \neq \sigma(j^2)$. If $\mu(j^2) \succ_{j^2} \sigma(j^2)$, this would contradict the way j^0 was chosen. If $\sigma(j^2) \succ_{j^2} \mu(j^2)$, this would contradict the way j^1 was chosen. Thus, we must conclude that in fact $C \cap (\bigcup_{t \leq t_0} \tilde{I}^t) = \emptyset$. $\diamond$

Continuing with the same agent j^0 and his assignment $\mu(j^0)$ as above, we now argue that $\mu(j^0) \notin \Omega(C|\omega_\mu^*, \mu)$. This will contradict (A.2) and therefore prove the theorem. Recall that

$\Omega(C|\omega_\mu^*, \mu) := \omega_\mu^*(\bigcup_{k=0}^\infty C_k)$ 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 $\mu(j^0) \notin \omega_\mu^*(C_k)$ for each k .

First, suppose $\mu(j^0) \in \omega_\mu^*(C)$. Thus, there exists some $i \in C = C_0$ such that $i \succeq_{\mu(j^0)} j^0$. Furthermore, Claim 2 implies that there is some $i' \in C = C_0$, such that $i' \in \tilde{I}^{t_{i'}}$ and $t_{i'} \leq t_0$. By Claim 4, no members of C are assigned a house at step t_0 , or earlier. Hence, we have arrived at a contradiction.

Continuing by induction, suppose $\mu(j^0) \notin \omega_\mu^*(C_{k'})$ for all $k' < k$. Suppose $\mu(j^0) \in \omega_\mu^*(C_k)$. 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 $t_{i'} \leq t_0$. However, repeated application of Claim 3 implies that the agent in C_k who is assigned a house earliest is necessarily a member of $C_0 = C$. By Claim 4, no members of C are assigned at step t_0 , or earlier, of Algorithm 2—a contradiction. $\square$

Proof of Theorem 4. Let μ be a strong exclusion core allocation. We will construct an extension $\tilde{\succ}$ such that Algorithm 2 outputs μ . The argument proceeds as follows. We define a sequence of directed 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 Algorithm 2.

Let $I^1 := I$, $H^1 := H$. Construct a graph Γ^1 with vertices $I^1 \cup H^1 \cup \{h_0\}$. Draw an arc from $i \in I^1$ to $h \in H^1 \cup \{h_0\}$ if and only if $\mu(i) = h$. Draw an arc from each $h \in H^1$ to $i \in I^1$ if and only if i is $\triangleright_h$ -maximal among agents I^1 . Finally, draw 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 $i \rightarrow h_0 \rightarrow i$ 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 $I^1 = \{i_1, i_2, \dots\}$. (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 $(i^0, h^0, \dots, i^{k-1}, h^{k-1})$, let i^k be the agent with the lowest index number (given the fixed enumeration) such that (a) $i^k \succeq_{h^{k-1}} \mu^{-1}(h^{k-1})$ and (b) i^k is $\triangleright_{h^{k-1}}$ -maximal among

agents in I^1 .³² Let $h^k := \tau^1(i^k)$. As there is a finite number of agents and houses, the sequence $(i^0, h^0, \dots)$ must eventually flow into a cycle. Relabeling as necessary, and without loss of generality, let $(i^0, h^0, \dots, i^{k-1}, h^{k-1})$ be that cycle. Thus, $i^0 \succeq_{h^{k-1}} \mu^{-1}(h^{k-1})$.

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 $C \neq \emptyset$.³³ Given the cycle $(i^0, h^0, \dots, i^{k-1}, h^{k-1})$ and the fact that $h^\ell \in \omega_\mu(i^{\ell+1 \pmod k})$ for all ℓ , it follows that $\{h^0, \dots, h^{k-1}\} \subseteq \Omega(C|\omega_\mu, \mu)$. Thus, coalition C can indirectly exclusion block μ by reallocating their most preferred houses among themselves, which is a contradiction. Therefore, 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 $i \rightarrow h_0 \rightarrow i$ and $\mu(i) = \tau^1(i) = h_0$, let $K^1 = (i, h_0)$. Otherwise, let $K^1 = (i^0, h^0, \dots, i^{k-1}, h^{k-1})$ 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 .

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 $(K^1, K^2, \dots, K^T)$ until no agents remain in Γ^T . (The outside option h_0 is always a member of Γ^T .)

Next, we use the sequence of defined cycles to define an extension $\tilde{\succ}$ of $\succ$. 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 $i^0 \rightarrow h^0 \rightarrow i^1 \rightarrow \dots \rightarrow h^{k-1} \rightarrow i^0$. For each h^ℓ let $i^{\ell+1 \pmod k}$ be the (unique) maximal element under $\tilde{\succ}_{h^\ell}$, i.e. $i^{\ell+1 \pmod k} \tilde{\succ}_{h^\ell} j$ for all $j \neq i^{\ell+1 \pmod k}$. The rest of $\tilde{\succ}_{h^\ell}$ can be defined in any way not violating $\triangleright_{h^\ell}$.

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 $i^0 \rightarrow h^0 \rightarrow i^1 \rightarrow \dots \rightarrow h^{k-1} \rightarrow i^0$. For each h^ℓ define $\tilde{\succ}_{h^\ell}$ as follows. First, identify all agents $j \in I \cap (\bigcup_{\tau < t} K^\tau)$ such that $j \triangleright_{h^\ell} i^{\ell+1 \pmod k}$. Let J be this set. Order these agents in an arbitrary manner not violating $\triangleright_{h^\ell}$. Place $i^{\ell+1 \pmod k}$ in the $\tilde{\succ}_{h^\ell}$ 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 $\triangleright_{h^\ell}$ is not violated. The constructed

³²Note that such i^k always exists. If $\mu^{-1}(h^{k-1})$ is $\triangleright_{h^{k-1}}$ -maximal in I^1 , then $i^k = \mu^{-1}(h^{k-1})$. If $\mu^{-1}(h^{k-1})$ is not $\triangleright_{h^{k-1}}$ -maximal in I^1 , then there exist (possibly multiple) $\triangleright_{h^{k-1}}$ -maximal agents dominating $\mu^{-1}(h^{k-1})$ with respect to $\triangleright_{h^{k-1}}$. The one with the lowest index number is i^k .

³³Otherwise $(i^0, h^0, \dots, i^{k-1}, h^{k-1})$ 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.

extension $\tilde{\triangleright}_{h^\ell}$ should have the following form:

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

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

One can now verify that Algorithm 2 outputs μ when $\tilde{\triangleright}_h$ is the linear extension of $\triangleright_h$ 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. $\square$

Proof of Proposition 4. The exclusion core's internal stability is implied by its definition. To verify external stability, we show that every allocation outside the strong exclusion core can be indirectly exclusion blocked by some coalition with a strong exclusion core allocation.

Fix a relational economy $\langle I, H, \succ, \triangleright \rangle$ with an acyclic priority structure. Let μ be an allocation not in the strong exclusion core. Let $\tilde{\triangleright}$ be an extension of $\triangleright$ such that for each h , $i \tilde{\triangleright}_h \mu^{-1}(h)$ if and only if $i \triangleright_h \mu^{-1}(h)$.³⁴ In other words, $\tilde{\triangleright}_h$ gives $\mu^{-1}(h)$ higher priority for h relative to any j for whom $j \not\triangleright_h \mu^{-1}(h)$. Let σ be the allocation identified by Algorithm 2 given $\tilde{\triangleright}$. Let $\tilde{I}^t$ be the set of agents assigned to a house in step t of the algorithm. Since $\triangleright$ is acyclic, σ is a strong exclusion core allocation. Let $C = \{i | \sigma(i) \succ_i \mu(i)\}$. This set is not empty because $\sigma \neq \mu$ and if $\mu(i) \succeq_i \sigma(i)$ for all i , then σ would not be Pareto efficient, a contradiction. We will show that coalition C can indirectly exclusion block μ with σ given ω_μ .

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 $\mu(j) \notin \Omega(C | \omega_\mu, \mu)$. Among agents satisfying these conditions, let j^0 be among those assigned earliest by the GTTC algorithm, say in step t_0 (i.e. $j^0 \in \tilde{I}^{t_0}$). Since $\mu(j^0) \succ_{j^0} \sigma(j^0) \succeq_{j^0} h_0$, $\mu(j^0) = h^1$ for some $h^1 \in H$ and h^1 must have been assigned before step t_0 , say in step $t_1 < t_0$. Thus, there exists $j^1 \in \tilde{I}^{t_1}$ such that $j^1 \tilde{\triangleright}_{h^1} i$ for all $i \in \bigcup_{t \geq t_1} \tilde{I}^t$. In particular, given the definition of $\tilde{\triangleright}$, $j^1 \tilde{\triangleright}_{h^1} j^0 = \mu^{-1}(h^1)$ 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 Algorithm 2 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.

³⁴By assumption $i \triangleright_h \emptyset$ for every $h \in H$. Thus, $\tilde{\triangleright}_h$ is defined for all $h \in H$.

Thus, we conclude $h^2 = \mu(j^1) = \sigma(j^1)$ and h^2 is assigned at step t_1 of the GTTC algorithm given $\tilde{\succ}$. Since more than one house is assigned in step t_1 , $h^2 \in H$. And so, there exists an agent $j^2 \in \tilde{I}^{t_1}$, such that $j^2 \tilde{\succeq}_{h^2} i$ for all $i \in \bigcup_{t \geq t_1} \tilde{I}^t$. From $j^1 \in \tilde{I}^{t_1}$, it follows that $j^2 \tilde{\succeq}_{h^2} j^1$, and from $j^2 \neq j^1$,³⁵ it follows that $j^2 \tilde{\succ}_{h^2} j^1 = \mu^{-1}(h^2)$. So, by the construction of $\tilde{\succ}$, it follows that $j^2 \triangleright_{h^2} j^1 = \mu^{-1}(h^2)$ and hence $h^2 \in \omega_\mu(j^2)$.

Following the arguments outlined above, it must be the case that $\sigma(j^2) = \mu(j^2)$ to avoid a contradiction.³⁶ Thus we can find an arbitrarily long chain of distinct houses and agents $(h^1, j^1, h^2, j^2, \dots)$ such that they are all cleared (in that order) at step t_1 of the algorithm, $\mu(j^k) = \sigma(j^k) = h^{k+1}$ for all $k = 1, 2, \dots$, and $j^k \tilde{\succeq}_{h^k} i$ for all $i \in \bigcup_{t \geq t_1} \tilde{I}^t$ and for all $k = 1, 2, \dots$. Note that the chain cannot cycle back to h^1 because $h^1 = \mu(j^0)$ and $j^0 \notin \tilde{I}^{t_1}$, which would contradict the fact that $\mu(j^k) = \sigma(j^k)$ for all k . This implies that $\tilde{I}^{t_1}$ contains infinitely many agents, which is impossible. $\square$

Proof of Proposition 5. We prove the proposition for the case of the strong exclusion core. The other cases follow similarly. It is sufficient to show that for all $C \subseteq I$ and any allocation μ , $\omega'_\mu(C) \subseteq \omega_\mu(C)$, where ω_μ (resp. ω'_μ) is the weak conditional endowment system in $\langle I, H, \succ, \triangleright \rangle$ (resp. $\langle I, H, \succ, \triangleright' \rangle$). If $h \in \omega'_\mu(C)$, then $i \succeq'_h \mu^{-1}(h)$ for some $i \in C$. And so, $i \succeq_h \mu^{-1}(h)$, which implies $h \in \omega_\mu(C)$. $\square$

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

(b) Consider a simple economy $\langle I, H, \succ, \omega \rangle$ where every house is privately or publicly owned. If $\triangleright$ represents ω , then $\triangleright$ is acyclic. By Lemma 2, the weak and strong exclusion cores of $\langle I, H, \succ, \triangleright \rangle$ coincide. Noting part (a), it is sufficient to show that the exclusion core of $\langle I, H, \succ, \omega \rangle$ is contained in the strong exclusion core of $\langle I, H, \succ, \triangleright \rangle$.

Let μ be an exclusion core allocation in $\langle I, H, \succ, \omega \rangle$. Suppose coalition C can indirectly exclusion block μ with σ given ω_μ in $\langle I, H, \succ, \triangleright \rangle$. Since μ is Pareto efficient, $\mu(j) \succ_j \sigma(j)$ for some $j \in I$. This implies $j \notin C$, $\mu(j) \in \Omega(C|\omega_\mu, \mu)$ and, therefore, $\mu(j) \neq h_0$. Thus, there exists a sequence of necessarily distinct agents $j, i^1, \dots, i^K$ such that $\mu(j) \in \omega_\mu(i^1)$, $\mu(i^1) \in \omega_\mu(i^2)$, $\dots$, $\mu(i^{K-1}) \in \omega_\mu(i^K)$ and $i^K \in C$.

As $\triangleright$ represents ω , $h \in \omega(i)$ if and only if $i \triangleright_h j$ for all $j \neq i$; otherwise, agents are not $\triangleright$ -comparable. Given μ , $h \in \omega_\mu(i)$ if and only if (i) $h \in \omega(i)$, (ii) $h = \mu(i)$, or (iii) $\mu^{-1}(h) = \emptyset$.

³⁵We know $j^1 \neq j^2$ because as part of the cycle at step t_1 , h^1 points to j^1 , h^2 points to j^2 , and $h_1 \neq h_2$.

³⁶If $\sigma(j^2) \succ_{j^2} \mu(j^2)$, then $j^2 \in C$ and, as $\mu(j^1) = h^2 \in \omega_\mu(j^2)$, it follows that $\mu(j^1) \in \Omega(C|\omega_\mu, \mu)$. We know that $\mu(j^0) = h^1 \in \omega_\mu(j^1)$ and so $\mu(j^0) \in \Omega(C|\omega_\mu, \mu)$, which is a contradiction. If $\mu(j^2) \succ_{j^2} \sigma(j^2)$ instead, then, since $j^0 \in \tilde{I}^{t_0}$, $j^2 \in \tilde{I}^{t_1}$, $t_1 < t_0$ and j^0 was chosen to be the earliest agent j assigned by the GTTC algorithm for whom both $\mu(j) \succ_j \sigma(j)$ and $\mu(j) \notin \Omega(C|\omega_\mu, \mu)$, it follows that $\mu(j^2) \in \Omega(C|\omega_\mu, \mu)$. Since $\mu(j^0) \in \omega_\mu(j^1)$ and $\mu(j^1) \in \omega_\mu(j^2)$, it follows that $\mu(j^0) \in \Omega(C|\omega_\mu, \mu)$, which is again a contradiction.

We have $\mu(j) \in \omega_\mu(i^1)$ but, as $j \neq i^1$ and $\mu^{-1}(\mu(j)) = j$, it follows that $\mu(j) \in \omega(i^1)$. We can extend the same logic to see that we must have $\mu(j) \in \omega(i^1)$, $\mu(i^1) \in \omega(i^2)$, $\dots$, $\mu(i^{K-1}) \in \omega(i^K)$. Therefore, $\mu(j) \in \Omega(C|\omega, \mu)$ and coalition C can indirectly exclusion block μ in $\langle I, H, \succ, \omega \rangle$, which is a contradiction. $\square$

B Generalized Top Trading Cycles: An Example

This example shows the step-by-step operation of Algorithm 2. There are four agents and four houses. The agents' preferences are

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

The priority structure is as follows: $i_1 \triangleright_{h_1} \{i_2, i_3, i_4\}$; $i_4 \triangleright_{h_2} \{i_1, i_2, i_3\}$; $i_1 \triangleright_{h_3} \{i_3, i_4\}$ & $i_2 \triangleright_{h_3} \{i_3, i_4\}$; $i_4 \triangleright_{h_4} \{i_1, i_2, i_3\}$. In words, $i_1 \triangleright_{h_1}$ -dominates everyone and there are no other $\triangleright_{h_1}$ -dominances. The remaining orders are read similarly. A diarchic structure governs h_3 .

First, suppose the extension of $\triangleright$ is

$$\tilde{\triangleright}_{h_1}: i_1, i_2, i_3, i_4 \quad \tilde{\triangleright}_{h_2}: i_4, i_1, i_2, i_3 \quad \tilde{\triangleright}_{h_3}: i_1, i_2, i_3, i_4 \quad \tilde{\triangleright}_{h_4}: i_4, i_1, i_2, i_3. \quad (\text{B.1})$$

Figure B.1 illustrates the operation of Algorithm 2. (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, i_2 is assigned h_4 . Finally, i_3 is assigned h_1 in step 3.

Suppose instead that the extension of $\triangleright$ is

$$\tilde{\triangleright}'_{h_1}: i_1, i_2, i_3, i_4 \quad \tilde{\triangleright}'_{h_2}: i_4, i_1, i_2, i_3 \quad \tilde{\triangleright}'_{h_3}: i_1, i_2, i_3, i_4 \quad \tilde{\triangleright}'_{h_4}: i_4, i_3, i_1, i_2. \quad (\text{B.2})$$

(B.2) is identical to (B.1) except i_3 ranks ahead of i_1 and i_2 in $\tilde{\triangleright}'_{h_4}$. Figure B.2 illustrates the operation of Algorithm 2 given (B.2). 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 above allocations are the only possible assignments identified by Algorithm 2 in this economy. Every other extension will lead to one of these two assignments.




Figure B.1: Operation of Algorithm 2 given the extension (B.1).




Figure B.2: Operation of Algorithm 2 given the extension (B.2).

References

- Abdulkadiroğlu, A., Pathak, P. A., and Roth, A. E. (2009). Strategy-proofness versus efficiency in matching with indifference: 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 indifference: The top trading absorbing sets mechanisms. *Games and Economic Behavior*, 73(1):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(3):481–513.
- Aziz, H. (2018). An impossibility result for housing markets with fractional endowments. Mimeo.
- 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.
- 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(6):2489–2497.
- 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.
- 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(5):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(1):155–171.
- Kesten, O. (2009). Why do popular mechanisms lack efficiency in random environments? *Journal of Economic Theory*, 144(5):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.
- Kreps, D. M. (2013). *Microeconomic Foundations I: Choice and Competitive Markets*. Princeton University Press, Princeton, NJ.
- Ma, J. (1994). Strategy proofness and the strict core in a market with indivisibilities. *International Journal of Game Theory*, 23(1):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.
- 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). Equilibrium 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(2):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 allocation with indifference: A generalization and a unified view. In *Proceedings of the Fourteenth ACM Conference on Electronic Commerce*, EC '13, pages 803–820, New York. ACM.
- Schummer, J. and Vohra, R. (2013). Assignment of arrival slots. *American Economic Journal: Microeconomics*, 5(2):164–185.
- 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.
- 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(2):167–190.
- Tang, Q. and Zhang, Y. (2016). Hierarchical exchange rules and the core in indivisible objects allocation. Mimeo.
- 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.

Endowments, Exclusion, and Exchange

Online Appendix

Ivan Balbuzanov*

Maciej H. Kotowski†

October 15, 2018

A Production Economies and the Proof of Theorem 6

To prove Theorem 6, we construct an exclusion core outcome using a modification of the Generalized Top Trading Cycles algorithm.

Algorithm A.1 (Generalized Top Trading Cycles for a Simple Economy with Production). Given $\langle (I, F), (H_0, \mathcal{H}), \succ, \omega \rangle$, fix an enumeration of $I = \{i_1, \dots, i_n\}$ and for each $h \in H := H_0 \cup \left(\bigcup_{k,j} H_{k,j} \right)$ define a linear order of I , $\tilde{\succ}_h$, as follows. First, rank all agents in C^h in order of their fixed enumeration. Then, rank all remaining agents according to their fixed enumeration. The resulting order should have the following structure:

$$\underbrace{i_k \tilde{\succ}_h \cdots \tilde{\succ}_h i_{k'}}_{C^h} \tilde{\succ}_h \underbrace{i_\ell \tilde{\succ}_h \cdots \tilde{\succ}_h i_{\ell'}}_{I \setminus C^h}$$

Given (A6) and the fixed enumeration of agents, we note that $\tilde{\succ}_h = \tilde{\succ}_{h'}$ for all $h, h' \in \bigcup_j H_{k,j}$.

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 as follows. The set of vertices is $I^t \cup H^t \cup \{h_0\}$. Draw an arc

*Department of Economics, Level 4 111 Barry Street, The University of Melbourne 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

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 $H^t \cup \{h_0\}$. For each $h \in H^t$, draw an arc from h to the $\succsim_h$ -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 $\tilde{H}^t = \emptyset$.
- (b) Otherwise, the constructed graph contains at least one cycle. (Cycles are necessarily disjoint.) Choose any cycle and carry out the implied assignments. That is, if $i \rightarrow h$ in the cycle then set $\mu(i) = h$. Let $\tilde{I}^t \subseteq I^t$ be the set of agents from the identified cycle and let $\mu(\tilde{I}^t)$ be the set of assigned houses. For each $h \in \mu(\tilde{I}^t)$, define the set $\hat{H}_h^t$ as follows.
 - (i) If $h \in H_0$, set $\hat{H}_h^t = \emptyset$.
 - (ii) Else, if $h \in H_{k,j}$ for some k and j , let $\hat{H}_h^t = \bigcup_{\ell \neq j} H_{k,\ell}$. Additionally, define $\rho(f_k) = H_{k,j}$.

Remove the agents $\tilde{I}^t$ and the houses $\tilde{H}^t = \mu(\tilde{I}^t) \cup \left(\bigcup_{h \in \mu(\tilde{I}^t)} \hat{H}_h^t \right)$ from the market.

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

The above process continues until $I^t = \emptyset$. Any remaining houses are left unassigned. If $\rho(\cdot)$ has not been defined for firm f_k before the algorithm's termination, simply set $\rho(f_k)$ to any feasible production plan in $\mathcal{H}_k$.

Remark A.1. The output of Algorithm A.1 is a feasible outcome. To see this, observe that prior to being removed from the market, all houses in $\bigcup_j H_{k,j}$ are pointing to the same agent. If a house $h \in H_{k,j}$ is assigned in step t , all houses in the production plan $H_{k,j}$ remain in the market and may therefore be assigned at some future step of the algorithm. It is impossible for a house in a different production plan belonging to firm f_k to be assigned in step t (otherwise cycles would intersect). These houses, $\hat{H}_h^t = \bigcup_{\ell \neq j} H_{k,\ell}$, are removed from the market at step t . Therefore, they cannot be assigned in any future step of the algorithm. Thus, the algorithm identifies a unique production plan for each firm.

Proof of Theorem 6. The following parallels the proof of Theorem 3. Let (μ, ρ) be an outcome identified by Algorithm A.1. The algorithm constructs μ sequentially by removing sets of agents $(\tilde{I}^1, \tilde{I}^2, \dots)$ and houses $(\tilde{H}^1, \tilde{H}^2, \dots)$. Houses that are removed in step t are either assigned to someone in $\tilde{I}^t$ or belong to a production plan that is not implemented.

To derive a contradiction, suppose coalition C can indirectly exclusion block (μ, ρ) with the feasible outcome (σ, ρ') . Thus, and without loss of generality, $\sigma(i) \succ_i \mu(i)$ if and only if

$i \in C$ and

$$\mu(j) \succ_j \sigma(j) \implies \mu(j) \in \Omega(C|\omega, \mu). \quad (\text{A.1})$$

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 remains in the market. Thus, if house $\sigma(i)$ was not yet removed before 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 $t < t_i$. $\diamond$

Claim 2. *Let $J \subseteq I$ and suppose $h \in \omega(J)$. 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. Fix $h \in \tilde{H}^t$, but assume that for all $i \in J$, $t_i > t$. That is, each agent in J is assigned strictly *after* house h . Since $h \in \omega(J)$, $C^h \subseteq J$. This implies that all agents in C^h must be present in the market at step t . In particular, the top-ranked agent according to $\tilde{\succ}_h$, say j , is present in the market and house h is pointing to him at step t . But if house h is removed from the market at step t , that implies that either h or another house h' with $C^{h'} = C^h$ was assigned at that step. All those houses point to j and therefore $j \in C^h \subseteq J$ was assigned a house at step t —a contradiction. $\diamond$

Claim 3. *Let t_i be the step of Algorithm A.1 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 $j \in (\mu^{-1} \circ \omega)(J)$.*

Proof of Claim 3. Let $j \in (\mu^{-1} \circ \omega)(J)$. Observe that $\mu(j) \in \omega(J)$ 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 $(\mu^{-1} \circ \omega)(J)$. $\diamond$

Henceforth, consider the earliest cycle occurring in Algorithm A.1 that contains an agent j^0 such that $\mu(j^0) \succ_{j^0} \sigma(j^0)$. It has to be the case that $\mu(j^0) \in H$. Suppose this cycle is removed at step t_0 . Let $\tilde{I}^{t_0}$ and $\tilde{H}^{t_0}$ be the sets of agents and houses, respectively, removed from the market at this step.

Claim 4. *If $i \in C$, then $i \in \bigcup_{t > t_0} \tilde{I}^t$, i.e. $C \cap (\bigcup_{t \leq t_0} \tilde{I}^t) = \emptyset$.*

Proof of Claim 4. Suppose there exists $j^1 \in C \cap (\bigcup_{t \leq t_0} \tilde{I}^t)$. If there are multiple such agents, choose one who is assigned during the earliest cycle occurring in Algorithm A.1. Let the cycle

containing j^1 be removed at step $t_1 \leq t_0$. By Claim 1, $\sigma(j^1) \in \tilde{H}^{t_2}$ for some $t_2 < t_1 \leq t_0$. There are two cases:

- (a) Suppose $\sigma(j^1)$ is available in the outcome (μ, ρ) . Then, there exists some $j^2 \in \tilde{I}^{t_2}$ such that $\mu(j^2) = \sigma(j^1)$. In particular, this implies that $\mu(j^2) \neq \sigma(j^2)$. If $\mu(j^2) \succ_{j^2} \sigma(j^2)$, this would contradict the way j^0 was chosen. If $\sigma(j^2) \succ_{j^2} \mu(j^2)$, this would contradict the way j^1 was chosen. Either way, we have arrived at a contradiction.
- (b) Suppose $\sigma(j^1)$ is not available in the outcome (μ, ρ) . Thus, $\sigma(j^1) \in H_{k,\ell} \neq \rho(f_k)$ for some k . This implies there must exist some $j^2 \in \tilde{I}^{t_2}$ such that $\mu(j^2) \in \rho(f_k)$. However, this means that house $\mu(j^2)$ is not available in outcome (σ, ρ') and therefore $\mu(j^2) \neq \sigma(j^2)$. If $\mu(j^2) \succ_{j^2} \sigma(j^2)$, this would contradict the way j^0 was chosen. If $\sigma(j^2) \succ_{j^2} \mu(j^2)$, this would contradict the way j^1 was chosen. Either way, we have arrived at a contradiction.

Cases (a) and (b) exhaust all possibilities. Thus, $C \cap (\bigcup_{t \leq t_0} \tilde{I}^t) = \emptyset$. ◇

Continuing with the same agent j^0 and his assignment $\mu(j^0)$ as above, we now argue that $\mu(j^0) \notin \Omega(C|\omega, \mu)$. This will contradict (A.1) and therefore prove the theorem. Recall that $\Omega(C|\omega, \mu) := \omega(\bigcup_{k=0}^{\infty} C_k)$ where $C_0 = C$ and $C_k = C_{k-1} \cup (\mu^{-1} \circ \omega)(C_{k-1})$. As $C_{k-1} \subseteq C_k$, it is sufficient to show that $\mu(j^0) \notin \omega(C_k)$ for each k .

First, suppose $\mu(j^0) \in \omega(C_0)$. Claim 2 implies that there is some $i' \in C_0 = C$, such that $i' \in \tilde{I}^{t_{i'}}$ and $t_{i'} \leq t_0$. By Claim 4, no members of C are assigned a house at step t_0 , or earlier. Hence, we have arrived at a contradiction.

Continuing by induction, suppose $\mu(j^0) \notin \omega(C_{k'})$ for all $k' < k$. Suppose $\mu(j^0) \in \omega(C_k)$. By definition, $C_k = C_{k-1} \cup (\mu^{-1} \circ \omega)(C_{k-1})$. Again, Claim 2 implies that there is some $i' \in C_k$, such that $i' \in \tilde{I}^{t_{i'}}$ and $t_{i'} \leq t_0$. However, repeated application of Claim 3 implies that the agent in C_k who is assigned a house earliest is necessarily a member of $C_0 = C$. By Claim 4, no members of C are assigned a house at step t_0 , or earlier, of Algorithm A.1—a contradiction. □