Economies with Complex Property Rights: The Role of Exclusion

IVAN BALBUZANOV
University of Melbourne
and
MACIEJ H. KOTOWSKI
University of Notre Dame

We discuss the exclusion core, a solution concept for object-allocation and object-exchange problems. The exclusion core is based on the right of exclusion and is especially useful for the analysis of economies with complicated property arrangements, such as those with shared ownership. The exclusion core coincides with the (strong) core in classic settings, and is closely related to the celebrated Top Trading Cycles algorithm.

General Terms: Economics

Additional Key Words and Phrases: Property Rights, Cooperative Games, Exchange Economies

1. INTRODUCTION

In a recent paper [Balbuzanov and Kotowski 2019a], we introduced the *exclusion core*, a cooperative solution concept for discrete exchange economies and allocation problems. Our objective was to better understand and model economies with complex property relations, such as those seen in the real world. Some goods might be private property, in the intuitive sense. Others might be jointly owned, perhaps by spouses or by a community. Still others may be enmeshed in a byzantine hierarchy of rights and obligations. Examples are common. Co-ownership of real-estate property may take the form of joint tenancy, tenancy in common, or even tenancy by the entirety. Digital goods, in particular, grant some property rights while withholding others. Uses of copyrighted e-books or music may be restricted by Digital Rights Management. Smartphone manufacturers can prevent owners from modifying or repairing their devices in certain ways. Intellectual property rights also come with their own complications, such as patent pools, patent thickets, and patent trolling.[1]

Economists' interest in property rights is of course nothing new. The "Coase theorem" [Coase 1960], the property rights theory of the firm [Grossman and Hart 1986; Hart and Moore 1990], and Ostrom's [1990] analysis of communal-property management are all Nobel-winning contributions. Our point of departure and motivation differ from this prior literature. We start by asking how well traditional solutions to economic models, the core in particular, perform when the property regime is nudged away from the "standard model." To the extent predictions diverge from intuition or fail completely, we then ask how the standard framework

[1] Heller and Salzman [2021] provide a popular account of the complexity of property relations.

Authors' addresses: ivan.balbuzanov@unimelb.edu.au, maciej.h.kotowski@nd.edu

can be reinterpreted to facilitate fruitful analysis on the enriched domain. For our reinterpretation, we draw upon an old and fundamental idea in the legal understanding of property, *the right to exclude*. The United States Supreme Court has called this "one of the most essential sticks in the bundle of rights that are commonly characterized as property"[2] and it appears in Blackstone's *Commentaries on the Laws of England*, a classic eighteenth-century treatise.[3] The right to exclude permits an owner to prevent others from entering or otherwise using his property.

Our analysis traces the logical implications of the right of exclusion by crafting a solution concept, which we call the exclusion core. At the heart of our construction is a reinterpretation of "endowments" in an economy—what agents start out with before any trade takes place—as a distribution of exclusion rights and not as bundles of tradable goods. Our analysis also uncovers some surprising connections with well-known market protocols, such as David Gale's Top Trading Cycles (TTC) algorithm [Shapley and Scarf 1974]. We believe our study sheds light on some conceptual issues in the definition, operation, and design of successful markets. In particular, market design exercises should not take "property rights" for granted or view their definition as implicitly obvious. Property rights help enable successful markets and while some market designers may consider them to be a constraint to accommodate, they may at times be a design variable as well.

In this letter, we aim to introduce our definitions and to outline some of our main results. We limit our discussion to the case of a so-called simple economy, which includes private- and public-ownership economies as special cases. While only reflecting a portion of our research on these topics, we believe this is the easiest venue to convey the intuition behind our analysis. Unless noted otherwise, proofs of all results presented below are in Balbuzanov and Kotowski [2019a].

2. SIMPLE ECONOMIES

In a seminal paper, Shapley and Scarf [1974] introduce the "house-exchange economy." In their model, there is a set of agents, each of whom initially owns one indivisible good, called a house. Each agent has use for at most one house, houses cannot be shared, and there are no other goods. There is no money or other medium of exchange. If the agents' preferences differ, it is natural to imagine them bartering, trading, or swapping houses with the aim of acquiring a more preferable dwelling. Who will (should?) get which house? Perhaps unexpectedly, this bare-bones model has proven remarkably influential and adaptable. Applications such as kidney exchange [Roth et al. 2004], student assignment [Abdulkadiroğlu and Sönmez 2003], and airport landing-slot allocation [Schummer and Vohra 2013] all build upon its foundation.

The starting point for our analysis is a generalization of Shapley and Scarf's setup. A *simple economy* is a tuple $\langle I, H, (\succeq_i)_{i\in I}, \omega\rangle$. $I := \{i_1, \dots, i_n\}$ and $H := \{h_1, \dots, h_m\}$ are finite sets of agents and indivisible objects ("houses"), respectively. $(\succeq_i)_{i\in I}$ is a profile of preferences, one for each agent. Each agent's preference $\succeq_i$ is a complete transitive binary relation defined over the set of houses H and an

[2] *Kaiser Aetna v. United States*, 444 U.S. 164 [1979].

[3] For more on the right of exclusion, see Merrill [1998], Merrill and Smith [2001], Balbuzanov and Kotowski [2019a], and the citations therein.

outside option $h_0 \notin H$ that represents "no consumption." We assume each agent's preference is strict and he is never indifferent between two distinct options. We write $h \succ_i h'$ if i strictly prefers h to h'; $h \succeq_i h'$ means $h \succ_i h'$ or $h = h'$.

Shapley and Scarf [1974] assumed that each agent is initially endowed with one house. We depart from this assumption as follows. An *endowment system* $\omega \colon 2^I \to 2^H$ specifies the houses owned by each coalition.[4] For example, if there are three agents and three houses, an endowment system ω might specify that

$$\begin{array}{llll} \omega(\varnothing) = \varnothing, \ \omega(i_1) = \{h_1, h_2\}, & \omega(\{i_1, i_2\}) = \{h_1, h_2\}, & \omega(\{i_1, i_3\}) = \{h_1, h_2\}, \\ \omega(i_2) = \varnothing, \ \omega(i_3) = \varnothing, & \omega(\{i_2, i_3\}) = \{h_3\}, & \omega(\{i_1, i_2, i_3\}) = \{h_1, h_2, h_3\}. \end{array} \tag{1}$$

In this case, i_1 is endowed with $\{h_1, h_2\}$. Agents i_2 and i_3 own nothing personally, but jointly own h_3. In all results and examples to follow, we maintain the assumption that the endowment system satisfies the requirement of

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

This property guarantees that each house has a well defined set of one or more joint owners. We call C^h the *minimal controlling coalition of* h. The (NC) condition is satisfied by the plain-vanilla private-ownership economy typical in economic analysis; it is even satisfied by the non-standard endowment system defined in (1) above.[5] Though already a generalization of the usual setup, we must stress that (NC) is nevertheless a strong assumption. It precludes, for example, cases where ownership rights are contested, ambiguous, or uncertain. Thus, there is scope for investigating alternatives to (NC) and we have done so elsewhere.[6]

Many well-known models satisfy the assumptions we have laid out so far. If we assume that there are exactly as many houses as agents, each agent owns one house (i.e., $\omega(i_k) = \{h_k\}$ for each k), and all houses are acceptable to all agents (i.e., $h \succ_i h_0$ for all i and all $h \neq h_0$), then we recover Shapley and Scarf's [1974] private-ownership economy. In this case, the aggregate endowment of coalition C is simply the union of the coalition members' endowments, $\omega(C) = \bigcup_{i \in C} \omega(i)$. We refer back to this special case in many examples below. Hylland and Zeckhauser [1979] consider a model that is the polar opposite of Shapley and Scarf's. In their object-assignment problem, all houses belong only to the social endowment, i.e., $\omega(C) = \varnothing$ for all $C \subsetneq I$ and $\omega(I) = H$. The "house-allocation problem with existing tenants" [Abdulkadiroğlu and Sönmez 1999] is a hybrid—each house's minimal controlling

[4] We omit braces when confusion is unlikely, e.g., $\omega(i)$ means $\omega(\{i\})$.

[5] To break (NC) in (1), one can, for example, replace $\omega(i_2) = \varnothing$ with $\omega'(i_2) = \{h_2\}$.

[6] It is worth noting that (NC) implies three weaker properties:

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

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

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

Our experience suggests that (A1)–(A3) are necessary properties that any reasonable endowment system should satisfy, even if (NC) is relaxed. In Balbuzanov and Kotowski [2019b] we investigate an economy satisfying a relaxed version of (NC) called *weak non-contestability*. A more significant departure from (NC) is found in Balbuzanov and Kotowski [2019a] where we study *relational economies*. See Section 6.

coalition is either a singleton (i.e., the house is privately owned) or the grand coalition (i.e., the house is part of the social endowment).

An *allocation* $\mu\colon I \to H \cup \{h_0\}$ is an assignment of agents to houses such that at most one agent is assigned to each $h \in H$. It is the final outcome after any barter, exchange, and trade has taken place. Let $\mu(C) := \bigcup_{i \in C} \mu(i)$ denote the aggregate allocation of coalition $C \subseteq I$.

3. AN EXAMPLE

Which final allocations are likely to arise or persist in a market? Two benchmark answers are provided by the weak and strong cores. Each solution consists of allocations that cannot be improved upon, or "blocked," by any coalition via a reassignment of the goods that it owns. An allocation μ is in the *weak core* if and only if there does not exist a nonempty coalition $C \subseteq I$ and an allocation σ such that (a) $\sigma(i) \succ_i \mu(i)$ for all $i \in C$ and (b) $\sigma(C) \subseteq \omega(C) \cup \{h_0\}$. When a coalition C and allocation σ satisfy conditions (a) and (b), we say that C can *strongly block* μ (via σ). An economy's *strong core* is defined like the weak core except point (a) is weakened to (a′) $\sigma(i) \succeq_i \mu(i)$ for all $i \in C$ with $\sigma(j) \succ_j \mu(j)$ for some $j \in C$, i.e., at least one member of a blocking coalition must be strictly better off. When a coalition C and allocation σ satisfy conditions (a′) and (b), we say that C can *weakly block* μ (via σ).[7] In both definitions, point (b) plays a critical role and is often taken for granted. It translates the model primitive delineating property, the endowment system ω, into an understanding about what coalitions can do with the economy's goods. In this case, a coalition is restricted to autarkic reallocations of any goods in its aggregate endowment.

There are many cases where the strong or weak core provide a compelling benchmark for economic analysis. One such case is the Shapley and Scarf [1974] economy where the (unique) strong core allocation is the consensus selection.[8] Nevertheless, both solutions exhibit weaknesses once we depart from the standard case. The following example is also discussed in Balbuzanov and Kotowski [2019a].

Example 3.1 *The Kingdom.* Consider an economy with three agents—peasants i and j, and their King k—and two indivisible houses—h_1 and h_2. Everyone agrees that h_1 is strictly better than h_2. The King is an absolute monarch and owns both houses: $\omega(k) = \{h_1, h_2\}$ and $\omega(i) = \omega(j) = \varnothing$.

We argue that there are only two allocations that can credibly arise in this small economy. The King owns both houses and likes h_1 more than h_2 so he is sure to take h_1 for himself. With his needs met, he no longer needs h_2 and so it should be occupied by either i or j. The remaining agent, j or i respectively, will settle for the outside option h_0. These two allocations are efficient and respect the Kingdom's ownership rights. To us, they appear as the focal outcomes in this economy.

However, neither the strong nor the weak core selects exactly these two outcomes. The strong core is empty. It is not hard to check that every allocation can be weakly blocked by a coalition involving the King and an agent originally receiving

[7] Admittedly, the terminology can become confusing. If μ is in the strong core, it is also in the weak core. If C can strongly block μ (via σ), it can also weakly block μ (via σ).

[8] Roth and Postlewaite [1977] and Ma [1994] reinforce the case favoring the strong core in Shapley and Scarf's model.

h_0. The weak core is not empty and includes the two focal allocations we identified above. However, it also includes an inefficient allocation that is difficult to justify. Specifically, the weak core includes the allocation in which h_2 is uninhabited and both i and j remain homeless. Without externalities, such as the King having an explicit preference for h_2 to remain empty, this allocation seems incredible and may prove unstable in the long run. A peasants' revolt may ensue.

The core solutions' deficiencies in the Kingdom can be attributed to an uneasy interplay between a coalition's incentive to block an allocation—points (a) or (a′) in the definitions above—and its ability to block an allocation—point (b). The strong core is empty because coalitions have too much blocking power. They can always rely on the participation of any agent who is indifferent between the blocked and blocking allocation. These agents' membership in a blocking coalition enlarges the set of houses that can be reallocated. Justifying the participation of indifferent agents is difficult, however. Sometimes altruism is advanced as a rationale. Joining a blocking coalition is an ambiguous action, however: it may just as well hurt some agents while helping others. Another explanation is the possibility of side payments. Indifferent agents, perhaps, can be "bribed" by those who strictly benefit. If this is possible, however, a pivotal agent might instead demand payments from those hurt by the blocking action. Without further restrictions on the existence and plausibility of side payments, he may even be able to set up a money pump with a suitably chosen pair of allocations that mutually block each other.

The weak core has the opposite problem: blocking is too hard. A coalition is limited to reallocating the goods that it owns, but it can form only if all coalition members are strictly better off. Even if an alternative allocation is a Pareto improvement, i.e., no one is harmed and someone is strictly better off, it is impossible to implement if the (co)owner of the some of the reallocated resources is indifferent.

4. THE EXCLUSION CORE

The Kingdom identifies some of the core definitions' shortcomings. To address those weaknesses, we rethink the balance between agents' incentives to block an allocation and the property rights that they draw upon to do so. Rather than assuming that a blocking coalition autarkically exchanges its own goods, we show that the weaker notion of excluding others from those goods is sufficient to support desirable and plausible outcomes.

Towards defining the exclusion core, we will first introduce a weaker solution concept, the direct exclusion core. We do so to better highlight our interpretation of exclusion rights. The exclusion core proper, defined below, will involve an inductive extension of the ideas in the following definition.

Definition 4.1. An allocation μ is in the *direct exclusion core* if and only if there does not exist a nonempty coalition $C \subseteq I$ and an allocation σ such that

(a) $\sigma(i) \succ_i \mu(i)$ for all $i \in C$, and

(b) $\mu(j) \succ_j \sigma(j) \Rightarrow \mu(j) \in \omega(C)$.

When there is a coalition C and an allocation σ satisfying (a) and (b), then coalition C can *directly exclusion block* μ (via σ).

Compared to the core definitions above, the direct exclusion core's novelty lies in part (b). It states that if an agent is made worse off by a blocking coalition, then that agent must have been evicted, i.e., *excluded*, from a house owned by the coalition. In other words, a coalition can block an assignment whenever each member of the coalition benefits from an alternative and any harm caused is legitimately rooted in the coalition's right to exclude. A direct exclusion core allocation cannot be destabilized by a coalition exercising its exclusion rights. This reasoning differs from the "exchange within a coalition" logic found in the classic formulations of the core in exchange economies.

The direct exclusion core resolves the Kingdom's troubles. The example's direct exclusion core coincides with the two intuitive and focal allocations identified above. The King claims his top choice and the remaining house is occupied. The agent without a house has no way of dislodging his peer (he lacks the exclusion rights to do so) and the King would not come to his aid (he does not gain by doing so). The inefficient weak-core allocation where h_2 is unoccupied is not in the direct exclusion core, either: it can be directly exclusion blocked by either i or j claiming h_2 for himself. In that case, condition (a) of Definition 4.1 is obviously satisfied for the blocking agent, while (b) holds vacuously as no one is harmed.

A natural question is whether the direct exclusion core solution can be strengthened? One tempting modification concerns replacing point (a) of Definition 4.1 with (a′) from the strong core definition. However, this route runs into the same problems identified above. Thus, our proposal for strengthening the direct exclusion core rests on using the interdependencies implied by exchange to increase a coalition's blocking power. The following example illustrates the reasoning.

Example 4.2. Consider an instance of the Shapley and Scarf [1974] economy with five agents and five houses. Suppose agent i_k privately owns house h_k. Table I summarizes the agents' preferences, listing houses in each agent's preferred order.

Table I: Preferences in Example 4.2.

$\succ_{i_1}$	$\succ_{i_2}$	$\succ_{i_3}$	$\succ_{i_4}$	$\succ_{i_5}$
h_2	h_3	h_4	h_2	h_3
h_5	h_1	h_1	h_5	h_1
h_1	h_2	h_3	h_4	h_5
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$

This economy has two direct exclusion core allocations, μ and σ. These are illustrated in Figures 1a and 1b, respectively. In each figure, there is an arc connecting each agent to his assignment at the prevailing allocation, e.g., $i_1 \to \mu(i_1) = h_2$ in Figure 1a and $i_1 \to \sigma(i_1) = h_5$ in Figure 1b. There is also an arc connecting each house to its owner, e.g., $h_1 \to i_1$ in Figures 1a and 1b.

We contend that σ is this economy's most compelling outcome, while μ hides an inherent fragility. At μ, agents i_2 and i_4 have a motive to block since they do not receive their favorite houses. Agent i_2 wants house h_3 and i_4 wants h_2 (Figure 1c). Luckily for i_4, the coalition $C = \{i_2, i_4\}$ has direct exclusion rights to h_2 since $\omega(C) = \{h_2, h_4\}$. Thus, i_4 can legitimately displace i_1, as in Definition 4.1.

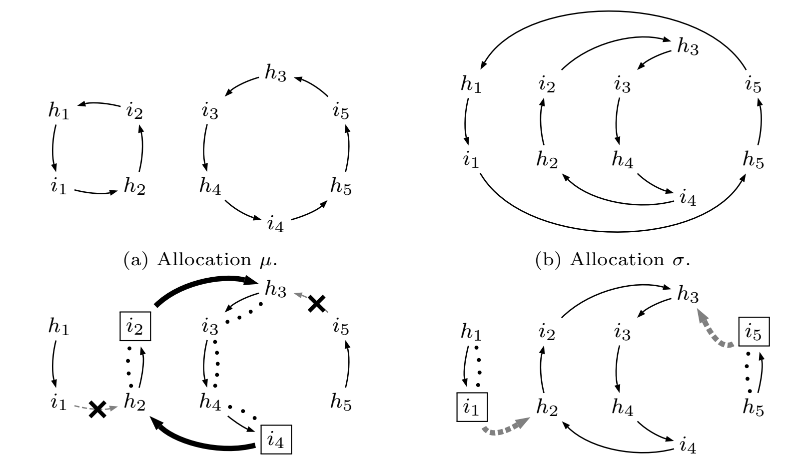

(a) Allocation μ.

(b) Allocation σ.

(c) Agents i_2 and i_4 can indirectly exclusion block μ.

(d) Agents i_1 and i_5 cannot indirectly exclusion block σ.

Fig. 1: Direct exclusion core allocations in Example 4.2.

However, the coalition $C = \{i_2, i_4\}$ cannot directly exclusion block μ since it lacks direct exclusion rights to h_3. However, i_3 (the owner of h_3) is assigned h_4 at μ. This house *is* in the coalition's endowment. Thus, i_3's well-being at μ depends on the coalition's endowment and on trade with the coalition. This gives the coalition leverage over i_3 at μ. Agents i_2 and i_4 may pressure i_3 to exclude i_5 from h_3—thus, making it available for i_2—by threatening to exclude him from h_4. Arguably, i_3 would accept this demand given his interest in a good acquired from the coalition. Thus, the pattern of trade defined by μ seems unstable and may unravel.

The allocation σ is immune to the preceding reasoning. Now, i_1 and i_5 would like to block the allocation to get h_2 and h_3, respectively (Figure 1d). However, the pattern of trade defined by σ fortifies i_2, i_3, and i_4 from any direct or indirect threats of exclusion originating from i_1 or i_5. The assignments of i_2, i_3, and i_4 do not depend on transaction chains that incorporate i_1 or i_5's endowments at their root. Accordingly, this allocation appears more likely to prevail in the long term.

Example 4.2 suggests that the right to exclude carries with it direct and indirect consequences. The latter arises from trade interdependencies. Critical for our analysis is the fact that the logic extends by induction beyond the single step illustrated in the example. This magnifies the exclusion power wielded by a coalition. Starting with Definition 4.1, coalition C has direct exclusion rights to all houses in the set

$$Z_0 = \omega(C).$$

At an allocation μ, it enjoys indirect exclusion rights to all houses in the set

$$Z_1 = Z_0 \cup \omega(C \cup \mu^{-1}(Z_0)).$$

The set $\omega(C \cup \mu^{-1}(Z_0)) \setminus Z_0$ consists of the *additional houses* controlled by C through pressure it potentially exerts on any agent assigned by μ to a house directly controlled by C. Continuing this reasoning, by relaying exclusion threats to anyone assigned to a house in Z_1, the coalition's indirect exclusion rights grow to $Z_2 =$

$Z_1 \cup \omega(C \cup \mu^{-1}(Z_1))$. And so on. In general,

$$Z_k = Z_{k-1} \cup \omega(C \cup \mu^{-1}(Z_{k-1}))$$

and we call

$$\Omega(C|\omega,\mu) = \bigcup_k Z_k$$

the *extended endowment of coalition C at allocation μ*.[9] It encompasses all houses over which the coalition has both direct and indirect exclusion rights and reflects the coalition's de facto market power given the prevailing pattern of trade. Observe that $\Omega(C|\omega,\mu)$ changes with μ.[10] Some allocations may expose agents to outside threats while others provide insulation, a distinction seen in Example 4.2 above.

Replacing $\omega(\cdot)$ with $\Omega(\cdot|\omega,\mu)$ in part (b) of Definition 4.1 leads to a strengthening of the direct exclusion core.

Definition 4.3. An allocation μ is in the *indirect exclusion core*, or simply the *exclusion core*, if and only if there does not exist a nonempty coalition $C \subseteq I$ and an allocation σ such that

(a) $\sigma(i) \succ_i \mu(i)$ for all $i \in C$, and
(b) $\mu(j) \succ_j \sigma(j) \Rightarrow \mu(j) \in \Omega(C|\omega,\mu)$.

When there is a coalition C and an allocation σ satisfying (a) and (b), then coalition C can *indirectly exclusion block* μ (via σ).

PROPOSITION 4.4. *The exclusion core is a subset of the direct exclusion core. Moreover, all exclusion core allocations are Pareto efficient and belong to the weak core.*

The first part of this proposition follows from the fact that $\omega(C) \subseteq \Omega(C|\omega,\mu)$. The second part is equally simple. If σ is a Pareto-improvement over μ, then it would satisfy both conditions in Definition 4.3. Condition (a) is satisfied since $\sigma(i) \succ_i \mu(i)$ for some i while (b) is vacuously true. Thus, μ cannot be in the exclusion core. The final part is less immediate. Suppose allocation μ is not in the economy's weak core. Thus, some coalition C can assign itself the houses in $\omega(C) \cup \{h_0\}$ in a way that leaves everyone in the coalition strictly better off. The same coalition can directly or indirectly exclusion block μ. If $i \in C$ receives a house originally assigned to some other $j \in C$, then no harm is done and this change is admissible. If $i \in C$ receives a house originally assigned to some $j \notin C$, then the disruption is permissible since $\mu(j) \in \omega(C) \subseteq \Omega(C|\omega,\mu)$.

The relationship between the exclusion core and the strong core is more nuanced. Sometimes the strong core is empty while the exclusion core is not. This occurs in the Kingdom example above.[11] There are also examples where the exclusion core

[9]This definition of the extended endowment is more tractable than the one in Balbuzanov and Kotowski [2019a]. The two definitions are equivalent [Balbuzanov and Kotowski 2019b].
[10]To see how sensitive $\Omega(C|\omega,\mu)$ is to the allocation μ, consider the case of the Shapley and Scarf [1974] economy, which we defined in Section 2. If μ is the "no trade" allocation where each agent keeps his endowment, i.e., $\mu(i_k) = \omega(i_k)$ for all k, then $\Omega(i_k|\omega,\mu) = \omega(i_k)$. If instead μ involves *all* agents trading in one cycle, e.g., $\mu(i_1) = h_2, \mu(i_2) = h_3, \ldots, \mu(i_n) = h_1$, then $\Omega(i_k|\omega,\mu) = H$.
[11]In the Kingdom the direct and indirect exclusion cores are the same.

is a strict subset of the strong core [Balbuzanov and Kotowski 2019a]. Equivalence occurs in some special cases that we discuss below.

We have already seen how some familiar solutions may be empty, despite the relatively restrictive setting of a simple economy. There are sound reasons to doubt whether an exclusion core allocation exists as well. Coalitions are unusually powerful with inductively inflated extended endowments. Moreover, extended endowments are functions of the prevailing allocation. Blocking cycles therefore are a real possibility. Surprisingly, perhaps, this does not occur. The right pattern of exchange can entirely neutralize agents' ability to gain from excluding others.

THEOREM 4.5. *The exclusion core of a simple economy is not empty.*

In Balbuzanov and Kotowski [2019a] we prove Theorem 4.5 constructively using a generalization of the TTC algorithm. The TTC algorithm was proposed by David Gale to Shapley and Scarf as a way to identify a price equilibrium in their house-exchange economy. Our algorithm builds upon earlier extensions of this market protocol in applications such as student assignment [Abdulkadiroğlu and Sönmez 2003] or transplant organ allocation [Roth et al. 2004].

To appreciate the link between trading cycles and the exclusion core, it is best to specialize to the Shapley and Scarf [1974] house-exchange economy. In this case, our algorithm reduces to that of the classic TTC algorithm, which identifies this economy's unique strong core allocation [Roth and Postlewaite 1977].

Algorithm 4.6 *Top Trading Cycles.* All agents and houses start off unassigned. In each step of the algorithm, each unassigned agent "points" to his most preferred unassigned house and each unassigned house "points" to its owner. There are finitely many agents and houses so there exists at least one cycle of alternating houses and agents, such as $h \to i \to \cdots \to h' \to i' \to h$. (The cycle may potentially consist of a single agent pointing to the house he owns.) Choosing any one of the resulting cycles, we assign each agent in the cycle to the house he is pointing to. Removing thus assigned agents and houses, we iterate the process with the next step. The algorithm terminates when all agents and houses have been assigned.

PROPOSITION 4.7. *In the Shapley and Scarf [1974] economy, the strong core and the exclusion core coincide. Thus, the unique exclusion core allocation is identified by the TTC algorithm.*

An informal graphical sketch of the proposition's proof is possible. First, without loss of generality, we can restrict attention to allocations where each agent is assigned to a house.[12] Such assignments partition the set of agents and their privately owned houses into disjoint cycles. The agents within a cycle swap houses among themselves. Now consider the particular allocation μ that is illustrated in both panels of Figure 2. It has five cycles $K_1, \ldots, K_5$.[13] As in Figure 1, each house in a cycle is "pointing" to its owner, i.e, $h_k \to i_k$, and each agent is "pointing" to

[12]By assumption, all houses are acceptable to all agents. If $\mu(i) = h_0$ for some i, agent i can weakly block μ since $\omega(i) \succ_i h_0$. Agent i can also exclusion block μ by claiming an unoccupied house. Thus, μ cannot be in the strong core or in the exclusion core.

[13]Agent i_{12} receives his endowment and forms a cycle by himself.

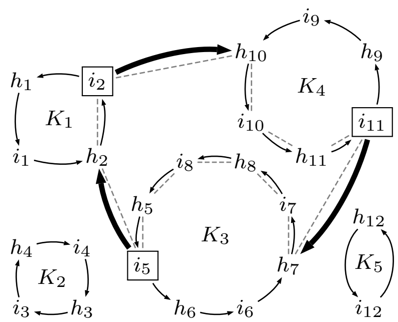

(a) If μ can be exclusion blocked, then it can also be weakly blocked.

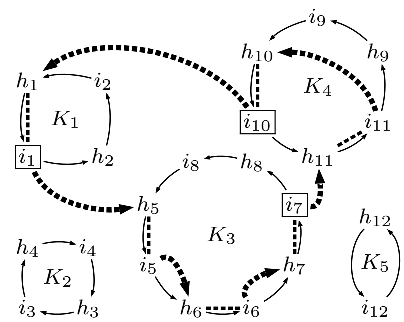

(b) If μ can be weakly blocked, then it can also be exclusion blocked.

Fig. 2: An allocation μ that is neither in the strong nor exclusion core.

his assigned house, i.e, $i_k \to \mu(i_k)$. We will explain that if μ is not in the exclusion core, then it cannot be in the strong core, and vice-versa.

Suppose coalition $C = \{i_2, i_5, i_{11}\}$ (boxed in Figure 2a) can indirectly exclusion block μ via σ. For concreteness, suppose this alternative allocation involves

$$\sigma(i_2) = h_{10} \qquad \sigma(i_5) = h_2 \qquad \sigma(i_{11}) = h_7,$$

as illustrated by the bolded links in Figure 2a.[14] Among these houses, the coalition only owns h_2 outright. It must access the others indirectly. In this case, h_7 is part of the same cycle as i_5 and a path of exchanges links i_7 to i_5's endowment. Likewise, h_{10} belongs to i_{10} who receives h_{11} from i_{11} in cycle K_4. Thus, h_2, h_7 and h_{10} all belong to the coalition's extended endowment as required by indirect exclusion blocking. Examining this situation, however, reveals that the coalition $C' = \{i_2, i_5, i_7, i_8, i_{10}, i_{11}\}$ can weakly block μ. The agents can exchange the houses in their endowment $\omega(C')$ in a cycle so that i_2, i_5, and i_{11} are strictly better off than at μ and i_7, i_8, and i_{10} are no worse off (follow the dashed links in Figure 2a to trace this cycle). Thus, if μ can be exclusion blocked, then it can be weakly blocked as well. Hence, the strong core is contained in the exclusion core.

Now consider the converse. Suppose μ can be weakly blocked by a coalition C via σ'. Without loss of generality we can assume coalition C cyclically exchanges the houses in its endowment $\omega(C)$.[15] In Figure 2b we posit such a cycle involves coalition $C = \{i_1, i_6, i_7, i_{10}, i_{11}\}$. The bolded dashed arrows identify each coalition member's assignment at σ'. Some members of this blocking coalition are strictly better off (e.g., i_1) while others' assignments are unchanged (e.g., i_6).[16] Suppose agent i_k is strictly better off. There are two cases. First, if $\sigma'(i_k) = h_\ell$ is in the same cycle as i_k at μ, then h_ℓ is in i's extended endowment at μ. We can trace a path from $\omega(i_k) = \{h_k\}$ back to the assignment of agent $i_\ell = \omega^{-1}(h_\ell)$. Second, if instead h_ℓ is in a different cycle than i_k, then as clear in Figure 2b there is some $i_m \in C$ who is in that cycle and for whom $\sigma'(i_m) \succ_{i_m} \mu(i_m)$ as well. Accordingly, the agents who strictly improve can together indirectly exclusion block μ. In the case of Figure

[14]The rest of σ does not matter.
[15]If there are multiple cycles, they are disjoint and we may focus on any one of them.
[16]Preferences are strict. Thus, if i is part of a blocking coalition, $\sigma'(i) \neq \mu(i) \Rightarrow \sigma'(i) \succ_i \mu(i)$.

2b, coalition $C' = \{i_1, i_7, i_{10}\}$ (boxed in the figure) can indirectly exclusion block μ via σ'.[17] Thus, if μ can be weakly blocked, then it can be indirectly exclusion blocked as well. Hence, the exclusion core is contained in the strong core.

5. PRIVATE AND PUBLIC OWNERSHIP ECONOMIES

The coincidence of the strong and exclusion cores in the Shapley and Scarf [1974] economy reflects a deeper relationship between these solutions in a private-ownership economy. In a *private-ownership economy*, the minimal controlling coalition of each house is a singleton though a particular agent may own multiple houses, as in the Kingdom.

PROPOSITION 5.1. *If the strong core of a private-ownership economy is not empty, then it coincides with the exclusion core.*

In a *public-ownership economy* all houses belong only to the social endowment. As no particular agent has any *a priori* claim to any particular house, the usual objective in such a market is to ensure a Pareto efficient assignment.

PROPOSITION 5.2. *In a public-ownership economy, the exclusion core equals the set of Pareto efficient allocations.*

The proof of this proposition follows from the observation that exclusion blocking always allows a coalition to block via a Pareto-improving allocation. When no one is harmed, point (b) of Definition 4.3 is moot. A Pareto efficient allocation in a public ownership economy cannot be exclusion blocked. Doing so would necessarily impose harm which is illegitimate since $\Omega(C|\omega, \mu) = \varnothing$ if $C \neq I$. The strong core of a public-ownership economy also equals the Pareto frontier. However, the weak core may include Pareto-inferior outcomes.[18]

It is simple to construct a Pareto-efficient allocation algorithmically using a *serial dictatorship*. Agent i_1 picks his favorite house. Then, agent i_2 picks his favorite among whatever is left. And so on. Changing the order in which agents pick may change the identified assignment. All Pareto efficient allocations can be identified by iterating through all $n!$ orders in which the agents might pick.

As a final case, we can consider a hybrid situation where some houses are privately owned and some are part of the social endowment. This case was introduced by Abdulkadiroğlu and Sönmez [1999] and is called the "house-allocation problem with existing tenants." A subset of agents $\hat{I} = \{i_1, \ldots, i_\ell\}$ are existing tenants and each is assumed to privately own a single distinct house, i.e., $\omega(i_k) = \{h_k\}$ for all $k = 1, \ldots, \ell$. Thus, $\hat{H} = \{h_1, \ldots, h_\ell\}$ are privately owned houses. The remaining agents $I \setminus \hat{I} = \{i_{\ell+1}, \ldots, i_n\}$ have nothing in their personal endowments. And, the remaining houses $H \setminus \hat{H} = \{h_{\ell+1}, \ldots, h_n\}$ belong solely to the social endowment. The minimal controlling coalition of each $h \in H \setminus \hat{H}$ is the grand coalition.

The challenge in the house allocation problem with existing tenants is to assign the houses efficiently while respecting existing tenants' property rights. An existing

[17] Agent i_{11} is also strictly better off. We can include him in C' without changing the argument.

[18] Any allocation where at least one agent, say i, receives his favorite house is in the weak core, even if every other agent receives his least-preferred assignment. Since i cannot be made *strictly* better off, he will never participate in any blocking coalition. But this implies every other coalition has no houses to reallocate making strong blocking impossible.

tenant may wish to upgrade his assignment by trading with another current owner or by opting for an unoccupied house in the social endowment. If he opts for an unoccupied house, his old house can be assigned to someone else. Motivating applications include office or dorm-room (re)assignment. Transplant organ exchanges exhibit some of this problem's key features as well [Roth et al. 2004].

Abdulkadiroğlu and Sönmez [1999] propose an ingenious algorithm to solve this allocation problem. Their "You Request My House—I Get Your Turn" (YRMH-IGYT) mechanism combines the features of the TTC and the serial dictatorship.[19] This mechanism always identifies a Pareto-efficient allocation, ensures each existing tenant gets an assignment at least as good as his original house, and is strategy-proof. This final property means that if agents' preferences must be solicited, then no agent can improve his assignment by falsifying his preference ranking. Since the YRMH-IGYT mechanism incorporates features from a serial dictatorship, its output will generally depend on an (exogenous) ordering of agents.

Proposition 5.3. *The exclusion core of the house allocation problem with existing tenants equals the range (over all orderings of agents) of Abdulkadiroğlu and Sönmez's [1999] YRMH-IGYT mechanism.*

6. OPEN QUESTIONS AND CONCLUSION

The complexity of real-world property rights is undeniable. However, our ability to model such rights and to formally understand their implications is surprisingly limited. In our work, we have sought take a step toward enriching the classic model of an exchange economy by considering less standard property arrangements. To do so, we have built our analysis around the right of exclusion. This is a very limited interpretation of property, even among legal scholars. Nevertheless, our analysis suggests that it can have wide-ranging consequences.

The key step for our analysis is the reinterpretation of endowments in an economic model as distributions of exclusion rights rather than as bundles of goods to trade. This point of view offers great flexibility that, regrettably, the simple economy exposited above may unintentionally conceal. Logically, multiple agents may hold exclusion rights to the same good, the exercise of this right might require collective action, and exclusion rights can be defined for both tangible and intangible goods, such as ideas (e.g., patents) or software (e.g., licensing and copyright).

There are numerous avenues for further research. One extension we have investigated is the possibility of production. In Balbuzanov and Kotowski [2019b] we maintain the basic setup of a simple economy, but we allow for the existence of firms that transform inputs into outputs. An output of one firm may be an input of another, thereby forming supply chains or production networks. A production connection is essentially a trading relationship among firms and it too can transmit exclusion threats, exactly like in the analysis above. We argue that production networks further expand coalitions' extended endowments and their configuration at an exclusion core outcome will need to strike a delicate balance. On one hand, they need to be expansive to assure the production and supply of desirable goods. On

[19] It reduces to the TTC if all houses have existing tenants. It reduces to a serial dictatorship if all houses are part of the social endowment.

the other hand, they must be sufficiently insulating to ensure firms are not beholden to opportunistic threats of hold-up. We identify sufficient conditions ensuring the exclusion core of a production economy is not empty. Our analysis points to two distinct types of criteria. Either there is sufficient integration or concentration of ownership rights within a supply chain or there is sufficient opportunity for multi-sourcing. As a technical contribution, we also propose an extension of the trading cycles algorithm to account for production.

A number of open questions about the exclusion core of simple economies remain. For example, the exclusion core as defined, may be empty if agents are allowed to consume multiple houses. Discrete exchange economies with multi-unit demand are known to be intractable—even their weak core may be empty [Konishi et al. 2001]. Nevertheless, there may be an appropriate reformulation of the exclusion core that preserves its intuition while also delivering useful predictions in these cases.[20]

A challenging problem concerns complex within-coalition relations when goods are jointly owned. Like all core-like definitions, Definition 4.3 requires coalition-unanimity to block an outcome. However, in practice weaker requirements, e.g., majority rule, are common. Admitting majority rule in our model effectively dispenses with the (NC) assumption and risks the emergence of Condorcet cycles. The exclusion core may be empty as different majorities cycle through blocking allocations.

Some features of joint ownership require a formal extension of the model to study adequately. For instance, in Balbuzanov and Kotowski [2019a] we examine a variant of the Kingdom economy from Example 3.1 called the Diarchy. In this case, two equally-powerful Kings co-own both houses. Surprisingly, there is no natural specification for the endowment system $\omega(\cdot)$ such that the exclusion core consists only of the two intuitive outcomes where the Kings get both houses. An extension of our framework to so-called relational economies (see below) lets us resolve this situation by making exclusion rights conditional on particular outcomes. In the Diarchy, a King can always exclude the peasant, but never his co-monarch.[21]

Another interesting extension concerns dynamic, or multi-period, markets. Dynamic variants of the house allocation problem have been studied before (see Kurino 2014, among others). An interesting question in such markets concerns what exactly is traded. Are the exclusion rights to a house exchanged or is only its use temporarily reassigned? The former, like buying or selling, sets agents' (future) outside options; the latter, like renting, affects only (current) welfare. Are some market protocols, like the TTC, better at handling one type of exchange or another? Can exclusion rights be created or destroyed over time?[22] A single-period economy obscures many of the dynamic questions closely tied to property.

Finally, one may ask a more fundamental question: where does an economy's endowment system, and the rights it implies, come from? Is it derived from some

[20]In Balbuzanov and Kotowski [2019b] we accommodate multi-unit consumption by introducing "firms" that produce notional goods that represent agent-specific bundles of goods.

[21]Other possible complications resulting from jointly-owned goods in simple economies have been examined by Sun et al. [2020].

[22]In practice, the answer to this question is "yes." Privatization and nationalization are examples of this process in action. Parsimoniously modeling this phenomenon as an equilibrium process seems challenging.

other primitive describing society?[23] For example, the exercise of property rights is often mired by formal and informal hierarchies. Common examples include a contractual right of first refusal or the challenges of shared ownership with (strong-willed) family members. These situations often lead to qualified exclusion rights or rights that are conditional on particular outcomes. These far more complicated forms of property relations can also be studied in our framework. In Balbuzanov and Kotowski [2019a] we introduce and examine *relational economies* where the endowment system is replaced by a *priority structure* that defines hierarchies among agents. The economy's endowment system, which is a necessary parameter for the exclusion core's definition, is then derived from these priorities. Necessarily, there are many ways in which this can be done and each leads to a different version of the exclusion core. Under appropriate conditions—an acyclicicty restriction on the priority structure—the close connection between the exclusion core and a trading cycles mechanism is maintained.

There is much left for us to explore about the exclusion core solution. However, our hope is that our proposal will encourage broader investigation of the questions that initially led us down this path. The definition, meaning, and interpretation of property rights in economic models has been taken for granted, deemed "obvious," or has often been imprecisely formulated. Translating the complexities of real-world property arrangements and understanding their economic implications is undoubtedly important. We believe the exclusion core solution offers a valuable lens for understanding some of these implications given its connection to a very basic understanding of property rights. We look forward to the development of other solutions and approaches that may improve upon its insights.

REFERENCES

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

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

Balbuzanov, I. and Kotowski, M. H. 2019a. Endowments, exclusion, and exchange. *Econometrica 87,* 5 (September), 1663–1692.

Balbuzanov, I. and Kotowski, M. H. 2019b. The property rights theory of production networks. HKS Faculty Research Working Paper RWP19-033, Harvard Kennedy School. November.

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

Demsetz, H. 1967. Toward a theory of property rights. *American Economic Review: Papers and Proceedings 57,* 2 (May), 347–359.

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 (August), 691–719.

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

Heller, M. A. and Salzman, J. 2021. *Mine! How the Hidden Rules of Ownership Control Our Lives.* Doubleday, New York.

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

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 (February), 1–15.

[23]There are several economic accounts concerning the emergence of property rights. A well-known theory is that of Demsetz [1967].

KURINO, M. 2014. House allocation with overlapping generations. *American Economic Journal: Microeconomics 6,* 1 (February), 258–289.

MA, J. 1994. Strategy-proofness and the strict core in a market with indivisibilities. *International Journal of Game Theory 23,* 1 (March), 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. 2001. What happened to property in law and economics? *Yale Law Journal 111,* 2 (November), 357–398.

OSTROM, E. 1990. *Governing the Commons: The Evolution of Institutions for Collective Action.* Cambridge University Press, Cambridge.

ROTH, A. E. AND POSTLEWAITE, A. 1977. Weak versus strong domination in a market with indivisible goods. *Journal of Mathematical Economics 4,* 2 (August), 131–137.

ROTH, A. E., SÖNMEZ, T., AND ÜNVER, M. U. 2004. Kidney exchange. *Quarterly Journal of Economics 119,* 2 (May), 457–488.

SCHUMMER, J. AND VOHRA, R. 2013. Assignment of arrival slots. *American Economic Journal: Microeconomics 5,* 2 (May), 164–185.

SHAPLEY, L. S. AND SCARF, H. 1974. On cores and indivisibility. *Journal of Mathematical Economics 1,* 1 (March), 23–37.

SUN, X., TANG, Q., AND XIAO, M. 2020. On the core of markets with co-ownerships and indivisibilities. Mimeo. Available at SSRN: https://dx.doi.org/10.2139/ssrn.3540210.