Matching with single-peaked preferences.

Sophie Bade[1]

[1]sophie.bade@rhul.ac.uk, 44 (0) 1784 437520, Royal Holloway, University of London and Max Planck Institut for research on Collective Goods, Bonn. Funding through the ARCHES prize is gratefully acknowledged. I would like Annika Johnson, Michael Mandler, Marcus Pivato, Karol Szwagrzak, Rann Smorodinsky, and Roee Teper for their comments.

Preprint submitted to Elsevier *November 13, 2018*

Matching with single-peaked preferences.

Sophie Bade[1]

Abstract

The crawler is a new efficient, strategyproof, and individually rational mechanism for housing markets with single-peaked preferences. In a housing market each agent is endowed with exactly one house. These houses are ordered - by their size for example - and all agents preferences are single-peaked with respect to that order. The crawler screens agents in order of their houses' sizes, starting with the smallest. The first agent who does not want to move to a larger house is matched with his most preferred house. Agents who currently occupy houses sized between this agent's original and chosen houses "crawl" to the next largest unmatched house. This process is repeated until all agents are matched. The crawler is easier to understand than Gale's top trading cycles and can be extended to allow for indifferences.

Keywords: Matching, Single-Peaked Preferences, Gale's top trading cycles, Obvious Strategyproofness.

1. Introduction

Consider a housing market where each agent i in a set $\{1, \ldots, n\}$ is endowed with a house, also called i. Suppose there is some objective linear order on all houses. Houses could be ordered by their location, so that $i < j$ means that house i lies to the south of house j. Alternatively houses could be ordered by their sizes, their energy efficiency, etc. All agents preferences

[1]sophie.bade@rhul.ac.uk, 44 (0) 1784 437520, Royal Holloway, University of London and Max Planck Institut for research on Collective Goods, Bonn. Funding through the ARCHES prize is gratefully acknowledged. I would like Annika Johnson, Michael Mandler, Marcus Pivato, Karol Szwagrzak, Rann Smorodinsky, and Roee Teper for their comments.

are single-peaked with respect to the objective order on houses. If preferences are single-peaked with respect to the north-south ordering then agents hope to live as close as possible to their preferred latitude. If size is the relevant objective order, then each agent has an ideal house size. Such an agent prefers a house that is a bit smaller (larger) than his ideal house to any other house that is yet smaller (larger). For ease of presentation I assume throughout that preferences are single-peaked with respect to house sizes.

A mechanism maps each profile of all agents' preferences to a matching. A matching, in turn, is a one-to-one function between agents and houses. The crawler, a new matching mechanism for the single-peaked domain, determines matchings by screening all houses in order of their size, starting with the smallest. Once a house whose current occupant i wants to either stay put or move to a smaller house is found, the crawler matches this agent i with his most preferred house. If this house is not the house that agent i occupied at the beginning of this step, then each occupant of a house sized between these two "crawls" to the next largest house[2]. This process is repeated until all agents are matched.

A mechanism is strategyproof if no agent can ever benefit from misrepresenting his preferences. It is efficient if it maps each profile of preferences to a matching for which there does not exist an alternative matching weakly preferred by all and strictly by some. It is individually rational if no agent is ever matched with a house he deems worse than the one he was endowed with. Theorem 1 shows that the crawler is efficient, strategyproof, and individually rational.

Without the assumption of single peakedness, exactly one mechanism satisfies these three criteria: when all linear orders are permitted as preferences, then Gale's top trading cycles is the unique efficient, strategyproof, and individually rational matching mechanism. In Gale's top trading cycles each agent points to the owner of his most preferred house. Any agent in a pointing cycle is matched with the house he points to. The procedure is repeated with all unmatched agents and the restriction of their preferences

[2]Fixing any unmatched house i, house $j > i$ is called the "next largest" house, if j remains unmatched and if any house j' with $i < j' < j$ is matched

to the unmatched houses and the algorithm terminates once a matching is reached.

While Shapley and Scarf's [17] and Roth's [13] results that Gale's top trading cycles is efficient, strategyproof, and individually rational also apply to the domain of single-peaked preferences, Ma's [10] result that Gale's top trading cycles is the only such mechanism, does not. I provide a new proof of Ma's [10] result to show how this result depends on richness of the domain of all linear preferences. Proposition 1 then indeed shows that the crawler differs from Gale's top trading cycles.

On the domain of single-peaked preferences the crawler has an advantage over Gale's top trading cycles. It has an extensive form implementation that is - in a well-defined sense - easier to understand than any extensive form implementation of Gale's top trading cycles. To define mechanisms that are more or less easy to understand, consider a strategy for some agent i in an extensive form mechanism. Arbitrarily fix a history where agent i moves and that can be reached if i plays the given strategy. This strategy is obviously dominant following Li [9] if i (weakly) prefers the worst outcome associated with the continuation of his strategy to the best outcome following a deviation at the current history (and all later histories). To calculate the relevant worst (best) payoff the agent considers the most harmful (favorable) choices by all other agents at all histories following the current one. Li [9] argues that even cognitively impaired agents or agents who suspect the designer of fraud never see a reason to deviate from an obviously dominant strategy. Theorem 3 shows that the crawler can be implemented in obviously dominant strategies. Conversely, I show that even on the restricted domain of single-peaked preferences Gale's top trading cycles cannot be implemented in obviously dominant strategies.[3]

In Section 6 I define a variant of the crawler that can be used on a larger domain of single-peaked preferences where agents may be indifferent

[3]Li [9] already showed in his Proposition 5 that Gale's top trading cycles cannot be implemented in obviously dominant strategies on the domain of all linear preferences. Bade and Gonczarowski [6] provide a characterization of all efficient mechanisms for housing markets that can be implemented in obviously dominant strategies.

between some houses. Theorem 4 shows that this variant inherits the three crucial properties of the crawler: it is efficient, individually rational and implementable in obviously dominant strategies.

While, the assumption of single-peaked preferences has a long pedigree in the social choice literature (see for example Moulin [11]), it is a novel assumption in the matching context. Independently of the present paper Damamme, Beynier, Chevaleyre, and Maudet [7] also study single-peaked preferences in the context of a housing market and find that sequences of individually rational bilateral swaps always reach the Pareto frontier in such housing markets. Without the restriction to single-peaked preferences the Pareto frontier can generally only be reached if larger groups of agents exchange their endowments. So Damamme, Beynier, Chevaleyre, and Maudet [7] and the present paper propose two different criteria according to which Gale's top trading cycles is not the best mechanism on the the domain of single-peaked preferences: some mechanisms require less centralization (Damamme, Beynier, Chevaleyre, and Maudet [7]) others satisfy more stringent incentive properties (Theorem 3).

2. Definitions

There is a set of agents $N\colon = \{1, \ldots, n\}$ and a set of equally many houses $O\colon = \{1, \ldots, n\}$. Each agent $i \in N$ is initially endowed with house i. Each agent i has a transitive and complete preference $\succsim_i$ over all houses. A profile of all agents' preferences is denoted $\succsim$. A house j is $\succsim_i$**-acceptable** if $j \succsim_i i$. If $\succsim_i$ ranks some house j above all houses in some set $N' \subset N$, I write $j \succsim_i N'$.

Houses are ordered by their sizes and $i < j$ means that i is smaller than j. Extending this terminology to agents I say that agent i is smaller than agent j if $i < j$. The preference $\succsim_i$ is **single-peaked** (with respect to the order $<$ on all houses) if there exists an $i^* \in N$ such that $j \succsim_i j'$ holds if either $j' < j \leq i^*$ or $j' > j \geq i^*$. A preference is a **linear** order if it is antisymmetric. Arbitrary domains of agent i's preferences and of preference profiles are denoted Ω_i and $\Omega\colon = \Omega_1 \times \cdots \times \Omega_n$. The domains of all linear preferences, of all single-peaked preferences and of all linear, single-peaked preferences respectively are Ω^l, $\hat{\Omega}$, and $\hat{\Omega}^l = \Omega^l \cap \hat{\Omega}$.[4]

A a **matching** is a one-to-one function $\mu : N \to N$. A match is a pair (i, o) of an agent i and a house o. The set of all matchings is $\mathcal{M}$. The identity $id : N \to N$ (with $id(i) = i$ for all $i \in N$) represents the initial endowment. Each agent i only cares about the house he is matched with: he prefers matching μ to matching μ' if and only if $\mu(i) \succsim_i \mu'(i)$. A matching μ is **efficient** at $\succsim$ if any matching μ' that is strictly better than μ for some agent is strictly worse than μ for a different agent. The same μ is **individually** rational at $\succsim$ if $\mu(i)$ is $\succsim_i$-acceptable for each i.

A social choice function $scf : \Omega \to \mathcal{M}$ maps each profile $\succsim$ in the arbitrary domain Ω to a matching in $\mathcal{M}$. Any social choice function can be viewed as a direct revelation mechanism. In such a mechanism, the agents declare their preferences to the designer who chooses the matching $scf(\succsim)$ given that $\succsim$ is the profile of stated preferences. The mechanism scf is efficient (individually rational) if $scf(\succsim)$ is efficient (individually rational) at $\succsim$ for each $\succsim \in \Omega$. It is **strategyproof** if no agent has an incentive to misrepresent his preferences, so $scf(\succsim)(i) \succsim_i scf(\succsim'_i, \succsim_{-i})(i)$ holds for all i, $\succsim'_i$ and $\succsim$.

3. The Crawler

The crawler $C : \hat{\Omega}^l \to \mathcal{M}$ is defined via a trading algorithm that screens agents and houses in ascending order. The smallest agent who wants to either stay or move to a yet smaller house leaves the mechanism with his most preferred house as his match. All agents who currently occupy houses at least as large as this agent's choice and smaller than the house he vacated "crawl" to the next largest house. The process is repeated until all agents are matched. To differentiate houses from agents, a generic house is now denoted at o_t. At Step k the agents in N^k and the houses in O^k remain unmatched.

Initialize: $N^1 \leftarrow N$, $O^1 \leftarrow N$.

Step k:

[4]The preference represented by $u_i(j) = - \mid j - i \mid$ is in $\hat{\Omega}$ but not in $\hat{\Omega}^l$ since i indifferent between $i+1$ and $i-1$. Since any $\succsim \in \hat{\Omega}$ may have multiple most preferred houses, such preferences are sometimes called single-plateaued.

Indexation: Let $m\colon =\mid N^k \mid$. Index agents and houses such that $N^k\colon = \{i_1, \ldots, i_m\}$ and $O^k\colon = \{o_1, \ldots, o_m\}$ with $i_t < i_{t+1}$ and $o_t < o_{t+1}$ for all $t \in \{1, \ldots, m-1\}$.

Screening: If $o_t \succ_{i_t} o_{t+1}$ holds for some t, let t^* be the minimal such t. If not, let $t^*\colon = m$. Say that o_r is the $\succsim_{i_{t^*}}$-best house in O^k.

Matching: Let $C(\succsim)(i_{t^*}) = o_r$.

Updating: Let $N^{k+1}\colon = N^k \setminus \{i_{t^*}\}$ and $O^{k+1}\colon = O^k \setminus \{o_r\}$. If $N^{k+1} = \emptyset$ terminate, otherwise go to Step $k+1$.

Exactly one agent, the agent i_{t^*} identified in Screening, is matched at each step. The Indexation step of the algorithm implicitly defines current occupancies in the sense that Agent i_t occupies House o_t for all $t \in \{1, \ldots, m\}$. So at each Step k smaller unmatched agents occupy smaller unmatched houses. Say Agent i_{t^*} gets matched with a house smaller than the one he currently occupies, so he gets matched with some o_r with $r < t^*$. Agent i_{t^*-1} then crawls to o_{t^*}, the houses vacated by i_{t^*}. This crawl frees up house o_{t^*-1} and crawling goes on until Agent i_r moves into house o_{r+1}. Any crawling agent prefers his new house to the house he occupied at the beginning of this step.

Example 1 Define a profile of preferences $\succsim$ for agents and houses $\{1, \ldots, 7\}$ with $2 \succsim_2 N$, $6 \succsim_4 N$, $3 \succsim_6 N$, $5 \succsim_7 N$, and where all other agents want to move to the largest possible house, so $7 \succsim_i N$ for $i = 1, 3, 5$. Figure 1 illustrates the crawling process. Each line represents the 7 houses. Houses are denoted by the top labels, agents by the bottom labels and finalized matches by boxes.

In the first line each house i is occupied by agent i. Screening all houses starting with the smallest, Step 1 finds Agent 2 as the smallest agent who does not want to move to a larger house. Since House 2 is the $\succsim_2$-best house, Agent 2 is matched with House 2 in Step 1. In Step 2, Agent 6 is the smallest agent who does not want to move to a larger house. The solid arrow in the second line illustrates that House 3 is the $\succsim_6$-best house. The dashed arrows, in the same line, show how Agents 3, 4, and 5 each crawl to the next largest house. The matches between Agent 2 and House 2 and between Agent 6 and House 3 are illustrated by the boxes around these pairs. In Step 3 Agent 7 is the smallest agent who does not want to move to a larger house. The solid arrow in the fourth line shows that Agent 7 top-ranks House 5. The dashed arrows represent Agent 4 and 5's, respective crawls to the next largest house. Onwards from the fifth line Agent 7 is boxed with House 5. Steps 4 through 7 match the remaining agents with the house they currently occupy, yielding the matches $(4, 6)$, $(5, 7)$, $(3, 4)$, and $(1, 1)$. No agent crawls in these steps. The last line in Figure 1 illustrates $C(\succsim)$.

Theorem 1 *The crawler $C : \hat{\Omega}^l \to \mathcal{M}$ is a well-defined, efficient, strategyproof, and individually rational mechanism.*

Proof Fix a profile $\succsim \in \hat{\Omega}^l$.

At any Step k, Screening finds an agent $i_{t^*} \in N^k$. Since this agent i_{t^*} is the only agent matched at Step k and since he is matched with $C(\succsim)(i_{t^*}) \in O^k$, an as of yet unmatched house, $C(\succsim)$ is a matching and the crawler is well-defined.

To see that $C(\succsim)$ is efficient, note that the agent matched in Step 1, is matched with his most preferred house. The agent matched in Step 2 is then matched with his most preferred house other than the house matched in Step 1. Proceeding inductively we see that the agent matched at any Step k is matched with his most preferred house that remains unmatched at Step k, so that $C(\succsim)$ is efficient.

Theorem 3 shows that C can be implemented in obviously dominant strategies which implies that it is also strategyproof.

Fix an arbitrary agent i. Say $\succsim_i^*$ is such that i is the $\succsim_i^*$-best house. Note that $C(\succsim_i^*, \succsim_{-i}')(i) = i$ holds for any $\succsim_{-i}'$. Since C is strategyproof we have $C(\succsim)(i) \succsim_i C(\succsim_i^*, \succsim_{-i})(i) = i$ and C is individually rational. □

4. Gale's top trading cycles

When there are at least three agents and three houses then the crawler differs from Gale's top trading cycles. Gale's top trading cycles $G : \Omega \to \mathcal{M}$ is defined for any domain Ω of linear preferences. To define Gale's top trading

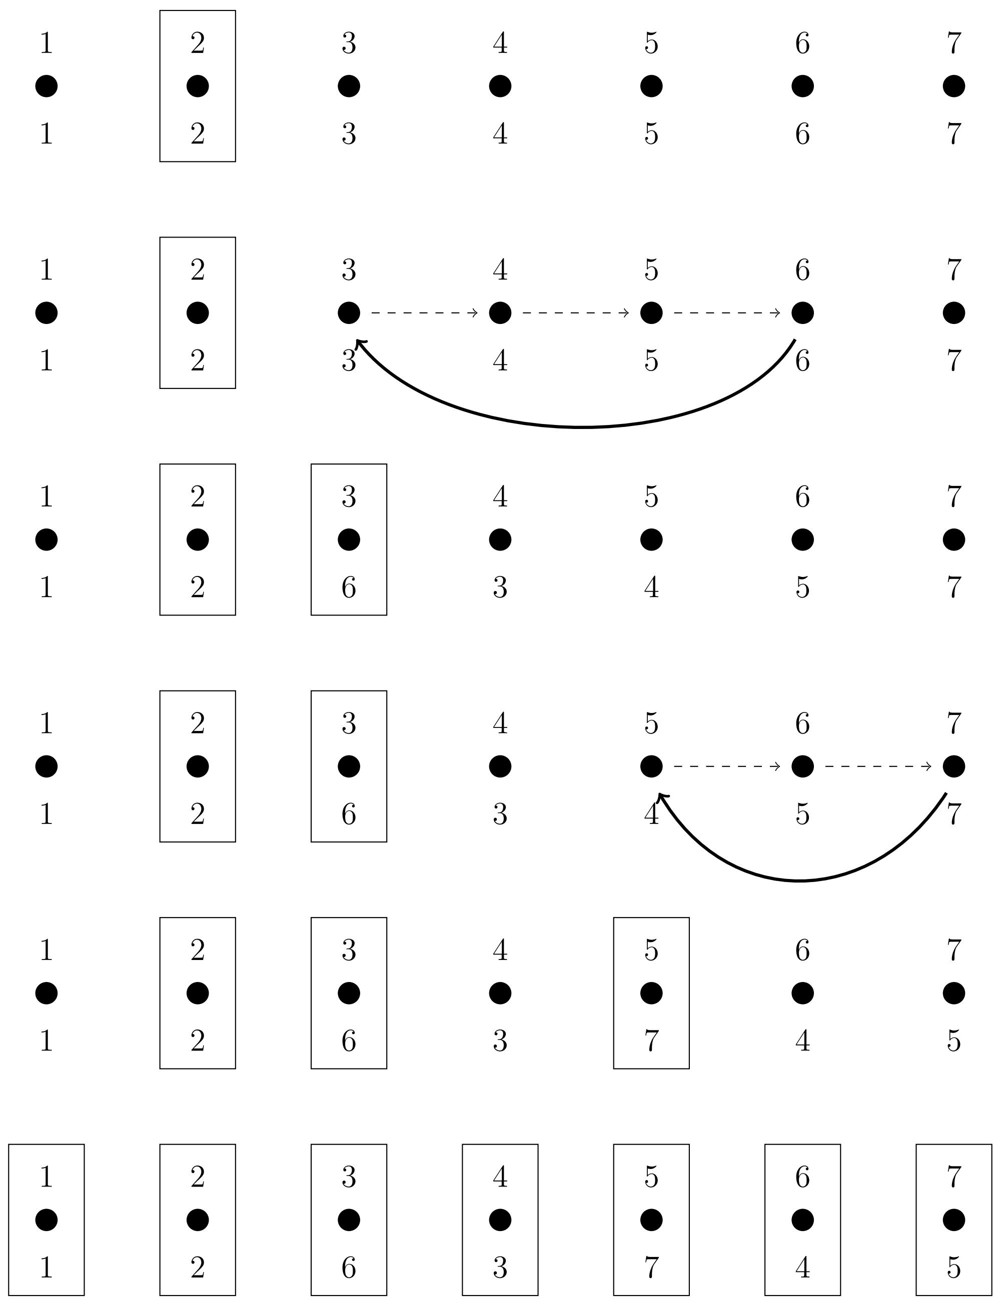

Figure 1: The crawling process

cycles, say $\rho : N_\rho \to N_\rho$ is a **cycle** on $N_\rho \subset N$ if for each $i, j \in N_\rho$ there exists some number m such that $\rho^m(i) = j$, where ρ^m is the m-th composition of ρ. Call the set of agents not in this cycle $\overline{N}_\rho \colon = N \setminus N_\rho$. The following algorithm finds $G(\succsim)$ for any profile $\succsim \in \Omega$.

Initialize: $N^1 \leftarrow N$.

Step k: Let each agent in N^k point to his most preferred house in N^k. Define a cycle $\rho : N_\rho \to N_\rho$ such that each agent $i \in N_\rho$ points to house $\rho(i) \in N_\rho$. Match each agent $i \in N_\rho$ with $\rho(i)$. Let $N^{k+1} \leftarrow N^k \setminus N_\rho$. Terminate if $N^{k+1} = \emptyset$, otherwise go to Step $k+1$.[5]

The parameter N^k denotes the set of unmatched agents at Step k. Since each agent i initially owns house i and since a house gets matched if and only if its owner gets matched, the set N^k is also the set of unmatched houses at Step k. Say all agents in N_ρ get matched in the first step of Gale's top trading cycles. Shapley and Scarf [17] and Roth [13] show that G is efficient, strategyproof, and individually rational on any domain Ω of linear preferences. Ma [10] shows that G is the unique such mechanism on Ω^l the domain of all linear preferences.

Theorem 2 ***[Ma [10]]*** *Fix a mechanism $M : \Omega^l \to \mathcal{M}$ on the domain of linear preferences. If M is efficient, strategyproof, and individually rational then it is Gale's top trading cycles.*

The upcoming proof of Theorem 2 highlights the reason why Theorem 2 does not apply to the single-peaked domain. I show that Theorem 2 applies to any domain where an agent i who sometimes finds a house j acceptable may be picky about that house j. A preference $\succsim_i$ is **picky about** j if j and i are the only two $\succsim_i$-acceptable houses. While Ω^l, the domain of linear preferences, contains all picky preferences, the single-peaked domain $\hat{\Omega}^l$ does not. Over the years Svensson [18], Anno [1] and Sethuraman [16]

[5] While this definition asks for all agents in one cycle to be matched in any given trading round, the standard definition asks for all cycles in any trading round to be matched. This difference is not relevant, since any cycle that forms at some round remains a cycle until it gets matched.

gave substantially more concise proofs of Theorem 2. My proof combines their simplifying ideas: Following Svensson [18] I use induction over the steps of Gale's algorithm. In line with Sethuraman [16] I directly work with efficient, strategyproof, and individually rational mechanisms (as opposed to the core).

Proof Say the cycle ρ forms at Step 1 of $G(\succsim^*)$. Define $\succsim_i'$ such that each $i \in N_\rho$ is picky about $\rho(i)$. To see $M(\succsim^*)(i) = \rho(i)$ for all $i \in N_\rho$ I show

$$M(\succsim'_{N_\rho \setminus S}, \succsim^*_{S \cup \overline{N}_\rho})(i) = \rho(i) \text{ for all } i \in N_\rho \text{ and } S \subset N_\rho.$$

by induction over $l\colon =\mid S \mid$. Since M is efficient and individually rational, $M(\succsim'_{N_\rho}, \succsim^*_{\overline{N}_\rho})(i) = \rho(i)$ holds for all $i \in N_\rho$ and the hypothesis holds for $l = 0$. Assume the hypothesis up to some $l \geq 0$. Fix a set S with $\mid S \mid = l+1$ and an agent $i \in S$. Let $S'\colon = S \setminus \{i\}$. Since M is strategyproof, and since the hypothesis applies to $S'\colon = S \setminus \{i\}$ we have

$$M(\succsim'_{N_\rho \setminus S}, \succsim^*_{S \cup \overline{N}_\rho})(i) \succsim_i^* M(\succsim'_{N_\rho \setminus S'}, \succsim^*_{S' \cup \overline{N}_\rho})(i) = \rho(i).$$

Since ρ forms at Step 1 of $G(\succsim^*)$, $\rho(i)$ is the $\succsim_i$-best house in N, and $M(\succsim'_{N_\rho \setminus S}, \succsim^*_{S \cup \overline{N}_\rho})(i)$ must equal $\rho(i)$ for any $i \in S$. Since M is individually rational and since $\rho : N_\rho \to N_\rho$ is a cycle $M(\succsim'_{N_\rho \setminus S}, \succsim^*_{S \cup \overline{N}_\rho})(i) = \rho(i)$ also holds for each $i \in N_\rho \setminus S$ and the hypothesis holds for all $S \subset N_\rho$.

Since only the preferences $\succsim^*_{N_\rho}$ of the agents in the cycle ρ matter for the above argument, $M(\succsim^*_{N_\rho}, \succsim_{\overline{N}_\rho})(i) = \rho(i)$ holds for all $i \in N_\rho$ and all $\succsim_{\overline{N}_\rho}$. So the mechanism M' with $M'(\succsim_{\overline{N}_\rho})(i)\colon = M(\succsim^*_{N_\rho}, \succsim_{\overline{N}_\rho})(i)$ for all $i \in \overline{N}_\rho$ is well-defined. Since M is efficient, strategyproof and individually rational M' is too. Say that the cycles ρ and ρ' form in Steps 1 and 2 of $G(\succsim^*_{N_\rho}, \succsim_{\overline{N}_\rho})$. By the arguments in the preceding paragraph, $M'(\succsim_{\overline{N}_\rho})(i) = \rho'(i)$ holds for each $i \in N_{\rho'}$. Proceeding inductively we see that $M = G$. $\square$

As most picky preferences are excluded from the linear domain of single-peaked preferences $\hat{\Omega}^l$, there may be multiple efficient, strategyproof and individually rational mechanisms $M : \hat{\Omega}^l \to \mathcal{M}$. The next proposition shows that this is indeed the case: the crawler differs from Gale's top trading cycles.

Proposition 1 *If* $\mid N \mid \geq 3$ *then the crawler* $C : \hat{\Omega}^l \to \mathcal{M}$ *differs from Gale's top trading cycles* $G : \hat{\Omega}^l \to \mathcal{M}$.

Proof Define $\succsim \in \hat{\Omega}^l$ such that the ideal houses of Agents 1, 2, and 3 respectively are 3, 1, and 1. If there are any agents $i > 3$ these agents are endowed with their ideal house. Step 1 of the crawler $C(\succsim)$ and of Gale's top trading cycles $G(\succsim)$ match two different agents (namely Agents 2 and 3) with House 1. $\square$

The alignment of interests implied by the restriction to single-peaked preferences allows for a larger set of efficient, strategyproof and individually rational mechanisms. Clearly the crawler and Gale's top trading cycles are not the only such mechanisms: the crawler has a dual mechanism that screens agents and houses in descending order. Mutatis mutandis all of the arguments in the present paper also apply to this dual mechanism.

5. Obvious Dominance

Some implementations of strategyproof mechanisms are easier to understand than others. To make this intuitive idea precise, Li [9] defines the concept of "obvious dominance" which distinguishes strategies that are merely dominant from strategies that can easily be recognized as such. Fix a strategy for an agent in some extensive form game. This strategy is obviously dominant if the following condition holds at any history where the agent gets to choose and that can be reached if he follows the given strategy. The agent's minimal utility if he continues the strategy given any possible follow-up choices of all other agents must be at least as high as his maximal utility given any collective deviation starting at the present history. Obvious dominance distinguishes between different extensive form mechanisms that implement the same strategyproof social choice function.[6] Li [9] provides experimental evidence that obviously dominant implementations of strategyproof mechanisms are indeed easier to understand than alternative implementations that do not satisfy this criterion.

[6]Li's [9] lead example to motivate obvious dominance shows that ascending clock auctions implement second price auction in obviously dominant strategies whereas the corresponding direct revelation mechanism does not.

Li [9] shows that Gale's top trading cycles with at least 3 agents cannot be implemented in obviously dominant strategies. The crawler is simpler to understand: there is an extensive form mechanism that implements the crawler in obviously dominant strategies.

An **extensive form mechanism** M is an extensive game form where a rooted tree represents the set of **histories** H. A history is **terminal** if it is not a subhistory of any other history. The set of all terminal histories is Z. The set of possible actions after the nonterminal history h is $A(h) := \{a \mid (h,a) \in H\}$. The set of players is N. The player function P maps any nonterminal history $h \in H \setminus Z$ to a player $P(h) \in N$ who gets to choose from all actions $A(h)$ at h.[7] Each terminal history $h \in Z$ is associated with a matching $\mu \in \mathcal{M}$. A **behavior** B_i for player i is a vector of actions that specifies a choice $a \in A(h)$ for each history h with $P(h) = i$. The path $Path(B)$ of a behavior profile $B := B_1 \times \cdots \times B_n$ is the set of all histories that are reached if all agents follow B. So $\emptyset \in Path(B)$ and $(h,a) \in Path(B)$ if $h \in Path(B)$ and $B_{P(h)}(h) = a$. The outcome $M(B)$ is associated with the unique terminal history $h \in Path(B)$. A **strategy** S_i for player i maps each $\succsim_i \in \Omega_i$ to a behavior $S_i(\succsim_i)$. A strategy profile $S = S_1 \times \cdots \times S_n$ consists of strategies for all agents.

A strategy S_i is **obviously dominant** (Li [9]) for agent i if for every $\succsim_i \in \Omega_i$, behavior profiles B and B', and history h, with $h \in Path(S_i(\succsim_i), B_{-i})$, $h \in Path(B')$, $P(h) = i$, $S_i(\succsim_i)(h) \neq B'_i(h)$ we have

$$M(S_i(\succsim_i), B_{-i})(i) \succsim_i M(B')(i).$$

A social choice function (or direct revelation mechanism) $scf : \Omega \to \mathcal{M}$ is implementable in obviously dominant strategies if there exists an extensive form mechanism M and a profile of obviously dominant strategies S such that $M(S(\succsim)) = scf(\succsim)$ for all $\succsim \in \Omega$.

To see that the crawler is implementable in obviously dominant strategies define an extensive form mechanism where each choice has the same format:

[7] Bade and Gonczarowski [6] show that simultaneous moves can be ignored without loss of generality when considering extensive form mechanisms that implement a social choice function in obviously dominant strategies.

an agent may either pick a house no larger than his current house or he may pass the choice to the next largest agent. At the start of the mechanism, Agent 1 may either choose the smallest house or pass. If he passes then Agent 2 may either choose one of the two smallest houses or pass. As long as agents pass, the next agent may either pass or be matched with a house no larger than the one he currently occupies. The occupant of the largest house cannot pass; he must pick a house as his match. Once an agent chooses to be matched, repeat the above questions, starting with the smallest remaining agent.

Theorem 3 *The crawler $C : \hat{\Omega}^l \to \mathcal{M}$ can be implemented in obviously dominant strategies.*

Proof To define an extensive form mechanism M, label each nonterminal history h with a vector (N_h, O_h, t_h). At h, N_h and O_h are the sets of unmatched houses and agents. These sets are indexed as outlined in the definition of the crawler and t_h designates the agent choosing at h. Let $(N_\emptyset, O_\emptyset, t_\emptyset) = (N, N, 1)$. The player function P maps each history h to $P(h) = i_{t_h}$. The choice set of agent i_{t_h} at h is $A(h) := \{o_1, \ldots, o_{t_h}, p\}$ if $t_h <\mid N_h \mid$ and $A(h) := O_h = \{o_1, \ldots, o_{t_h}\}$ otherwise. Now consider (h, a) for all possible actions $a \in A(h)$. The history (h, a) is terminal if and only if $A_h = \{o_r\}$. Otherwise, the label of (h, a) is updated so that $(N_{(h,o_r)}, O_{(h,o_r)}, t_{(h,o_r)}) := (N_h \backslash \{P(h)\}, O_h \backslash \{o_r\}, 1)$ and $(N_{(h,p)}, O_{(h,p)}, t_{(h,p)}) := (N_h, O_h, t_h + 1)$.

Define a strategy profile S such that $S_{P(h)}(\succsim_{P(h)})(h) = p$ if house o_{t_h+1} is strictly $\succsim_{P(h)}$-preferred to o_{t_h}, otherwise let $S_{P(h)}(\succsim_{P(h)})(h)$ be the $\succsim_{P(h)}$-best house in O_h.

Each string of pass choices corresponds to a step of the crawling algorithm: Since agents are asked in order of their index, any such string goes on until the smallest agent who does not want to move to a larger house. According to S, such an agent is then matched with his most preferred unmatched house. So the behavior $S(\succsim)$ induces the choices made in all steps of the crawling algorithm at $\succsim$. In sum we obtain $M(S(\succsim)) = C(\succsim)$ for each $\succsim \in \hat{\Omega}^l$.

To see that S is obviously dominant fix an arbitrary history h. Say that o_r is the $\succsim_{i_{t_h}}$-best house in O_h. If $r \leq t_h$ then $P(h) = i_{t_h}$ is matched with

o_r if he follows S. Since o_r is the $\succsim_{i_{t_h}}$-best house in O_h agent $P(h) = i_{t_h}$ cannot be made strictly better off by any collective deviation starting at h. If $r > t_h$ and if agent i_{t_h} follows S, then he gets matched with a house in $\{o_{t_h}, \ldots, o_r\}$. If he deviates at h, his match is in $\{o_1, \ldots, o_{t_h}\}$. Since o_r, the $\succsim_{i_{t_h}}$-best house in O_h, is larger than o_{t_h} and since $\succsim_{i_{t_h}}$ is single-peaked, the agent weakly prefers any house in $\{o_{t_h}, \ldots, o_r\}$ to any house in $\{o_1, \ldots, o_{t_h}\}$. Since h was chosen arbitrarily, S is obviously dominant. $\square$

Any social choice function that can be implemented in obviously dominant strategies is strategyproof. Theorem 3 therefore implies that the crawler $C : \hat{\Omega}^l \to \mathcal{M}$ is strategyproof, completing the proof of Theorem 3.

The combination of Ma's [10] uniqueness result (Theorem 2) with Li's [9] proof that Gale's top trading cycles with more than two agents is not implementable in obviously dominant strategies, implies that no efficient and individually rational mechanism for housing markets with more than two agents can be implemented in obviously dominant strategies. Bade and Gonczarowski [6] study obvious dominance in a variety of settings and come to the conclusion that only very few efficient social choice functions can be implemented in obviously dominant strategies. They show in particular that median voting with single-peaked preferences is not implementable in obviously dominant strategies. Arribillaga, Masso, and Neme [3] come to a similar conclusion on the dearth of obviously strategyproof voting mechanisms.

Theorem 3 then simultaneously applies the domain restrictions of the impossibility results by Li [9] (housing market) and Bade and Gonczarowski [6] as well as Arribillaga, Masso and Neme [3] (single peakedness) to obtain a possibility result for the single-peaked housing markets. The crawler is not only efficient and individually rational, it can be implemented in obviously dominant strategies. This possibility result differs in two dimensions from the preceding impossibility results. It concerns a novel mechanism (the crawler) as well as a novel domain of preferences (single-peaked housing markets). So one may now wonder whether the restriction on the domain is sufficient for Gale's top trading cycles to be implementable obviously dominant strategies. In the appendix I give a negative answer. I show that $G : \hat{\Omega}^l \to \mathcal{M}$ can be implemented in obviously dominant strategies if and only if there are less than four agents.

6. Indifferences

While Gale's top trading cycles is the unique efficient, strategyproof, and individually rational mechanism for the linear domain Ω^l, there exists a plethora of different such mechanisms on the grand domain Ω that permits indifferences. With single-peaked preferences there are - of course - even more such mechanisms, given that Ma's [10] uniqueness result (Theorem 2) does not apply.

Saban and Sethuraman [15] lay out a protocol for the construction of efficient, strategyproof, and individually rational matching algorithms that reduce to Gale's top trading cycles on the domain of linear preferences. Firstly, the rules governing trading cycles need to be amended such that they do not depend on unique most preferred houses. Trading cycles are, secondly not used to define matches but to update endowments for the continuation of the algorithm. Saban and Sethuraman [15], thirdly, require occasional matching steps that match some agents with their current endowments. The fourth requirement is that agents stay unmatched as long as their participation in the algorithm may increase the efficiency of the outcome. The fifth requirement is that the algorithm terminates. Saban and Sethuraman [15] show that the mechanisms defined by Jaramillo and Manjunath [8] as well as by Alcalde-Unzu and Molis [2] fit their requirements.[8] Here I extend Saban and Sethuraman's [15] protocol to the crawler to obtain the circle-crawler, an efficient and individually rational matching mechanism that reduces to the crawler on the domain of linear preferences and that can be implemented in obviously dominant strategies.

Following Saban and Sethuraman's [15] first requirement I amend the crawler's rules for the identification of agents and houses so that these rules are well-defined whether or not an agent has a unique most preferred house.

[8] With an eye towards the reduction of running times, Plaxton [12] proposes a different protocol for the generation of efficient, strategyproof, and individually rational matching algorithms that reduce to Gale's top trading cycles on the domain of linear preferences. Plaxton's [12] set of algorithms is contained in Aziz and De Keijzer's [4]. While the mechanisms defined by Aziz and De Keijzer's [4] need not be strategyproof, they outperform the mechanisms of Jaramillo and Manjunath [8] as well as by Alcalde-Unzu and Molis [2] in terms of running times.

While each step of the crawler and the circle-crawler identifies the smallest agent i who strictly prefers his current house to the next largest (remaining) house, they use different rules to identify a house: Where the crawler refers to agent i's best house the circle-crawler refers to this agent's *smallest* best house. The crawler would then match agent i with his (smallest) best house. However, in accordance with Saban and Sethuraman's [15] second requirement, moves from house to house do not immediately yield matches and the circle-crawler designates agent i to become the current occupant of his smallest best house. In line with Saban and Sethuraman's [15] third requirement the circle-crawler periodically matches some agents with the houses they currently occupy. Such matches occur whenever the algorithm revisits a current occupancy. This condition to match agents not only guarantees that the circle-crawling algorithm terminates (Saban and Sethuraman's [15] fifth requirement), it is also responsible for the efficiency of the circle-crawler (Saban and Sethuraman's [15] fourth requirement).

In the circle crawling algorithm agents move to smaller houses without immediately getting matched. The principle that smaller agents occupy smaller houses does therefore not implicitly define current occupancies. **Current occupancies** are instead explicitly defined as one-to-one functions $\nu : N' \to N$ for $N' \subset N$, with the understanding that agent i currently occupies $\nu(i)$. To define the circle-crawler $C^{\sim} : \hat{\Omega} \to \mathcal{M}$ use the following algorithm to calculate $C^{\sim}(\succsim)$ for each $\succsim \in \hat{\Omega}$:

Initialize$^{\sim}$: $N^1 \leftarrow N$ and $\nu[1]$ the identity on N.

Step k:

Indexation$^{\sim}$: Let $m :=| N^k |$. Index houses such that $\nu[k](N^k) := \{o_1, \ldots, o_m\}$ with $o_t < o_{t+1}$ for all $t \in \{1, \ldots, m-1\}$. Index agents $N^k := \{i_1, \ldots, i_m\}$ such that $\nu[k](i_t) = o_t$.

Screening$^{\sim}$: If $o_t \succ_{i_t} o_{t+1}$ holds for some t, let t^* be the minimal such t. If not, let $t^* := m$. Choose agent i_{t^*}. Say o_r is the smallest $\succsim_{i_{t^*}}$-best house in $\nu[k](N^k)$.

Circle-Crawling$^{\sim}$: Define an occupancy $\tilde{\nu} : N^k \to N$ with

$$\tilde{\nu}(i_{t^*}) :\, = o_r,$$
$$\tilde{\nu}(i_t) :\, = o_{t+1} \text{ if } r \leq t < t^*$$
$$\tilde{\nu}(i_t) :\, = o_t \text{ otherwise .}$$

If $\nu[k'] \neq \tilde{\nu}$ for all $k' \leq k$ let $\tilde{N} :\, = \emptyset$, otherwise define $\tilde{N}$ as the set of all agents $i \in N^k$ who were chosen in Screening$^{\sim}$ at some Step $k' < k$.

Matching$^{\sim}$: Let $C^{\sim}(\succsim)(i) :\, = \tilde{\nu}(i)$ for each agent $i \in \tilde{N}$.

Updating$^{\sim}$: Let $N^{k+1} :\, = N^k \setminus \tilde{N}$. If $N^{k+1} = \emptyset$ terminate, otherwise define $\nu[k+1]$ as the restriction of $\tilde{\nu}$ to N^{k+1} and go to Step $k+1$.

Example 2 illustrates the circle-crawler.

Example 2 Consider $N :\, = \{1, 2, 3\}$ together with the profile $\succsim$ where $\succsim_1$, $\succsim_2$, and $\succsim_3$ respectively top rank the sets $\{1, 2, 3\}$, $\{1, 2\}$, and $\{1\}$. In the unique efficient matching at $\succsim$, Agents 1 and 3 swap houses, while Agent 2 keeps his endowment. In Step 1 of the circle crawler, Screening$^{\sim}$ identifies Agent 2 as the first agent who strictly prefers the house he currently occupies (House 2) to the next largest remaining house (House 3). So Circle-Crawling$^{\sim}$ yields the current occupancy $\tilde{\nu}(1) = 2$, $\tilde{\nu}(2) = 1$, and $\tilde{\nu}(3) = 3$. Since $\nu[1] \neq \tilde{\nu}$, no agent is matched at Step 1 and $\nu[2]$ is set to $\tilde{\nu}$. At Step 2, Screening$^{\sim}$ identifies Agent 3 and Circle-Crawling$^{\sim}$ yields the current occupancy $\tilde{\nu}(3) = 1$, $\tilde{\nu}(2) = 2$, and $\tilde{\nu}(1) = 3$. Since $\nu[1] \neq \tilde{\nu} \neq \tilde{\nu}[2]$, no agent is matched at Step 2 and $\nu[3]$ is set to $\tilde{\nu}$. At Step 3, Screening$^{\sim}$ once again identifies Agent 3 as the first agent who does not strictly prefer the next largest house and we obtain $\tilde{\nu} = \nu[3]$. Since the preceding steps identified Agents 2 and 3 in Screening$^{\sim}$, we have $\tilde{N} = \{2, 3\}$ and Matching$^{\sim}$ requires to match Agents 2 and 3 with the houses they currently occupy (respectively House 1 and 2). Only Agent 1 is left at Step 4, where he gets matched with House 3. The circle-crawler, in sum, finds the unique efficient matching at $\succsim$.

Theorem 4 *The circle-crawler $C^{\sim} : \hat{\Omega} \to \mathcal{M}$ is an efficient and individually rational matching mechanism that can be implemented in obviously dominant strategies. For any $\succsim \in \hat{\Omega}^l$ we have $C(\succsim) = C^{\sim}(\succsim)$.*

Proof To see that $C^{\sim}$ is well-defined and efficient fix a profile $\succsim \in \hat{\Omega}$.

To see that the circle-crawling algorithm terminates, suppose not. So suppose there exists a Step k^{∞} with $N^{k^{\infty}} = N^k \neq \emptyset$ for all $k > k^{\infty}$. Since finitely many agents and houses remain unmatched there must exist two steps k' and k'' with $k^{\infty} \leq k' < k''$ such that the current occupancy $\tilde{\nu}$ determined at Step k'' equals $\nu[k']$. So at least one agent must, by Matching$^{\sim}$, get matched in Step k'', a contradiction. Since any Step k either matches no one or bijectively matches uniquely defined sets of agents and houses, and since new steps are initiated as long as some agents remain unmatched, the circle-crawling algorithm results in a matching $C^{\sim}(\succsim)$.

To see that $C^{\sim}(\succsim)$ is efficient fix an arbitrary agent i. Say i gets matched with $C^{\sim}(\succsim)(i)$ at Step k. It suffices to see that agent i (weakly) prefers his match $C^{\sim}(\succsim)(i)$ to any house that remains unmatched at Step k and strictly prefers his match $C^{\sim}(\succsim)(i)$ to any house that remains unmatched at Step $k+1$. Say Circle-Crawling$^{\sim}$ in Step k finds the current occupancy $\tilde{\nu}$ and the set of agents $\tilde{N} \ni i$, which is then matched in Matching$^{\sim}$. Let Step k' be such that $\nu[k'] = \tilde{\nu}$. Let Step k'' be the earliest step such that no agent is matched at the Steps $\{k'', k''+1, \ldots, k-1\}$. Say $O\colon = \nu[k](N^k)$ is the set of unmatched houses at Step k'' and O^* the set of $\succsim_i$-best houses in O. Say o_{min} and o_{max} are the minimal and maximal elements of O^*. To see that $C^{\sim}(\succsim)$ is efficient at $\succsim$ it then suffices to show that $C^{\sim}(\succsim)(i) \in O^* \subset C^{\sim}(\succsim)(\tilde{N})$.

Since $i \in \tilde{N}$, Screening$^{\sim}$ at some Step $k^* \in \{k'', \ldots, k\}$ must choose Agent i, where he moves from o_{max} to o_{min}. Screening$^{\sim}$ and Circle-Crawling$^{\sim}$ imply that any agent i° who remains unmatched at some Step k° weakly prefers $\nu[k^{\circ}+1](i^{\circ})$ to $\nu[k^{\circ}](i^{\circ})$. So if $k^* \neq k$, we have $o_{min} = \nu[k^*+1](i) \precsim_i \nu[k^*+2](i) \precsim_i \cdots \precsim_i \nu[k](i) \precsim_i \tilde{\nu}(i) = C^{\sim}(\succsim)(i)$. Since $o_{min} \in O^*$ and since $\succsim_i$ is transitive we have $C^{\sim}(\succsim)(i) \in O^*$.

If O^* is a singleton the first claim implies the second. So suppose there exists a $j \in O^* \setminus \{C^{\sim}(\succsim)(i)\}$. Circle-Crawling$^{\sim}$ implies that any agent who moves from house o_t at some step either moves to a smaller house or to the next largest house o_{t+1}. This together with Agent i's move from o_{max} to o_{min} at Step k^* (as defined above) implies that Agent i occupies each house in O^*, particularly House j, at some step between k'' and k. Since

$C^{\sim}(\succsim)(i) = \tilde{\nu}(i) \neq j$, House j is not occupied by the same agent at all steps between k'' and k. The Agent i' with $\tilde{\nu}(i') \colon = j$ therefore moves houses at some steps between k'' and k. Since $\nu[k'](i') = \tilde{\nu}(i')$, Agent i' must at some step before k move to a smaller house. At that step Screening$^{\sim}$ must choose Agent i', implying $i' \in \tilde{N}$. We in sum have $j = C^{\sim}(\succsim)(i')$ for some $i' \in \tilde{N}$ and therefore $O^* \subset C^{\sim}(\succsim)(\tilde{N})$.

To see that the circle-crawler $C^{\sim}$ is implementable in obviously dominant strategies define an extensive form mechanism $M^{\sim}$. Label each nonterminal history h in $M^{\sim}$ with a vector $(N_h, \nu_h, t_h, \overline{o}_h)$. At h, N_h is the set of unmatched agents and ν_h the current occupancy. The sets N_h and $\nu_h(N_h)$ of unmatched agents and houses are indexed as outlined in the definition of the circle-crawler. Agent i_{t_h} is the player $P(h)$ at h. Finally the function $\overline{o}_h : N \to N$ sets upper bounds on the agents' matches, so that no terminal history following on h matches agent i with a house $o > \overline{o}_h(i)$.

Let $N_{\emptyset} = N$, $\nu_{\emptyset}(i) = i$ for all $i \in N$, $t_{\emptyset} = 1$ and $\overline{o}_h(i) = \max N$ for all $i \in N$. Let $A(h) \colon = \{o_1, \ldots, o_{t_h}, p\}$ if $o_{t_h} < \overline{o}_h(i_{t_h})$ and $A(h) \colon = \{o_1, \ldots, o_{t_h}\}$ if $o_{t_h} = \overline{o}_h(i_{t_h})$. Now consider (h, a) for all possible actions $a \in A(h)$. If $a = p$ update labels such that $(N_{(h,p)}, \nu_{(h,p)}, t_{(h,p)}, \overline{o}_{(h,p)}) \colon = (N_h, \nu_h, t_h + 1, \overline{o}_h)$. If $a = o_r$ calculate $\tilde{\nu}$ according to the rules set out in Circle-Crawling$^{\sim}$. If $\tilde{\nu} = \nu_{h'}$ for a sub-history $h' \subset h$, let $\tilde{N}$ be the set of all agents $i \in N_h$ with $P(h'') = i$ and $(h'', p) \not\subset h$ for some $h'' \subset h$, otherwise let $\tilde{N} = \emptyset$. Match each agent $i \in \tilde{N}$ with $\tilde{\nu}(i)$. If $\tilde{N} = N_h$ then (h, o_r) is a terminal history. If $\tilde{N} \neq N_h$ update $\overline{o}_{(h,o_r)}(i_{t_h}) \colon = o_{t_h}$ and $\overline{o}_{(h,o_r)}(i) \colon = \overline{o}_h(i)$ for all $i \neq i_{t_h}$. Let $N_{(h,o_r)} \colon = N_h \setminus \tilde{N}$, $\nu_{(h,o_r)}$ the restriction of $\tilde{\nu}$ to $N_{(h,o_r)}$, and $t_{(h,o_r)} \colon = 1$.

After a pass choice by $P(h)$ the occupant of o_{t_h} nothing much happens: the occupant of the next largest unmatched house gets to choose, all else stays the same. Conversely, the choice of a house at h triggers extensive updating. We firstly have to calculate a new current occupancy. If that new current occupancy coincides with a preceding one, then each unmatched agent who chose a house at subhistory of h gets matched with the house he currently occupies. The mechanism terminates if all agents are matched. If not, the upper bound of agent $i_{t_h} = P(h)$ is updated to the house he occupied at the start of the step. The new current occupancy is the restriction of the occupancy determined at the present history to the set of unmatched agents.

The occupant of the smallest unmatched house gets to choose next.

To define a strategy profile S fix an arbitrary history h and say that $o_{\min}$ and $o_{\max}$ are the minimal and maximal most preferred unmatched house of agent $i_{t_h} = P(h)$. If either $p \notin A(h)$ or $o_{\max} \leq o_{t_h}$ then $S(\succsim)(h)$ is agent $P(h)$'s minimal most preferred house in $A(h)$. Otherwise (if $p \in A(h)$ and $o_{\max} > o_{t_h}$) let $S(\succsim)(h) = p$. To see that S is obviously dominant consider the fixed history h. **Case 1:** $p \notin A(h)$. For all terminal histories following on h agent i_{t_h} gets matched with a house in $A(h)$. Since i_{t_h} gets matched with a $\succsim_{i_{t_h}}$-best house in $A(h)$ if he follows S at h, this choice is obviously dominant. **Case 2:** $p \in A(h)$ and $o_{\max} \leq o_{t_h}$. Since Agent i_{t_h} gets matched with a $\succsim_{i_{t_h}}$-best house in the set of all unmatched houses $\nu_h(N_h)$ if he follows S at h, this choice is obviously dominant. **Case 3:** $p \in A(h)$ and $o_{\max} > o_{t_h}$. If agent i_{t_h} follows S at h he chooses p and gets matched with a house in $\big\{ \min\{o_{\min}, o_{t_h}\}, \ldots, o_{\max}\big\}$. If he deviates to choose some o_r at h his maximal potential match $\overline{o}_{(h,o_r)}(i_{t_h})$ is updated to o_{t_h} and his match is in $\big\{o_1, \ldots, o_{t_h}\big\}$. Since i_{t_h}'s largest most preferred house is strictly larger than o_{t_h} and since $\succsim_{i_{t_h}}$ is single-peaked, the agent weakly prefers any house in $\big\{ \min\{o_{\min}, o_{t_h}\}, \ldots, o_{\max}\big\}$ to any house in $\big\{o_1, \ldots, o_{t_h}\big\}$, and S also prescribes in this final case an obviously dominant choice.

To see that $M^{\sim}(S(\succsim)) = C^{\sim}(\succsim)$ holds for all $\succsim \in \hat{\Omega}$ fix an arbitrary $\succsim$. Consider the first sequence of pass choices given $M^{\sim}(S(\succsim))$ which terminates with the choice of a house by some agent who is either the occupant of the largest house ($\max N$) or strictly prefers the house he occupies to the next largest house. This agent's maximal potential house is updated to the house he occupied at the start of the current step. Given that the agent follows S this house is at least as large the agent's maximal most preferred unmatched house. The ensuing circle crawl in $M^{\sim}(S(\succsim))$ exactly corresponds to the first circle crawl in $C^{\sim}(\succsim)$. So both $M^{\sim}(S(\succsim))$ and $C^{\sim}(\succsim)$ find the same current occupancy. Both $M^{\sim}(S(\succsim))$ and $C^{\sim}(\succsim)$ then apply the same rules whether and whom to match. Finally both $M^{\sim}(S(\succsim))$ and $C^{\sim}(\succsim)$ restrict the current occupancy to the set of all unmatched agents to obtain a new current occupancy for the start of the next sequence of pass choices.

The next sequence of pass choices differs from the preceding one only insofar as that the maximal potential house for some agent may have changed.

However, the strategy profile S is such that no upper bound $\overline{o}_h(i)$ of any agent i is ever updated to a house smaller than agent i's maximal most preferred unmatched house, implying that on the path of S, S yields the same choices whether there are upper bounds on houses or not. So the second sequence of passes in $M^{\sim}(S(\succsim))$ terminates with the choice of the agent chosen in Step 2 of $C^{\sim}(\succsim)$. From then on $M^{\sim}(S(\succsim))$ and $C^{\sim}(\succsim)$ again follow the same rules to update and to match. Proceeding inductively we see that $M^{\sim}(S(\succsim)) = C^{\sim}(\succsim)$.

The proof that C is individually rational (see Theorem 1) applies unchanged to the case of $C^{\sim}$.

To see that $C(\succsim) = C^{\sim}(\succsim)$ holds for each $\succsim \in \hat{\Omega}^l$ fix an arbitrary such $\succsim$. In Step 1 we obtain identical indexations of all agents N and houses N. Screening and Screening$^{\sim}$ choose the same agent $i = i_{t^*}$. They also choose the same house o_r since the smallest $\succsim_{i_{t^*}}$-best house in N is the unique $\succsim_{i_{t^*}}$-best house. According to the crawler, Agent i_{t^*} is then matched with o_r. If $r = t^*$, then the current occupancy $\tilde{\nu}$ determined in Step 1 equals $\nu[1]$ and Step 1 of the circle-crawler also matches i_{t^*} with o_r. If not, then each Agent i_t with $r \leq t < t^*$ crawls to the next largest house. Step 2 of the circle-crawler then once again chooses the same agent i (now agent i_r) and house o_r in Screening$^{\sim}$. Since agent i (after Step 1) occupies his best house o_r, he stays put in the current round of Circle-Crawling$^{\sim}$ and therefore gets matched with o_r at Step 2. No matter whether the circle-crawler matches Agent i at Step 1 or 2, the current occupancy is such that smaller agents occupy smaller houses. So the next step of the circle-crawler starts out the same as does Step 2 of the crawler. Proceeding inductively we see $C^{\sim}(\succsim) = C(\succsim)$.$\square$

The presence of updated maximal potential houses $\overline{o}_h$ complicates the proof that $M^{\sim}(S(\succsim))$ equals $C^{\sim}(\succsim)$ for any profile of preferences $\succsim$. In fact, $M'(S(\cdot)) = C^{\sim}(\cdot)$ also holds for the mechanism M' which is identical to $M^{\sim}$ except that maximal potential houses are not updated in M'. Dropping all references to maximal potential houses, the proof that $M^{\sim}(S(\cdot))$ equals $C^{\sim}(\cdot)$ directly applies to show $M'(S(\cdot))$ equals $C^{\sim}(\cdot)$.

Updated maximal potential houses are however crucial for the obvious strategyproofness of $M^{\sim}$. To see that S is not obviously dominant in the alternative mechanism M' consider a problem with just three agents and assume that Agent 2 top ranks the largest house. Consider the history (p) where Agent 1 has chosen to pass. At (p) the strategy S prescribes for Agent 2 to choose p. If Agent 3 chooses House 3 at (p, p) Agent 2 does not get House 3. If Agent 2 deviates from S and chooses House 1 at (p) he may be matched House 3. This indeed happens if Agents 1 and 3 behave as if they are indifferent between all houses, whereas Agent 2, following on (p, o_1), behaves as if he most prefers House 3. What stands out is that Agent 2's choice of o_1 at (p) together with his later behavior is inconsistent with single peaked preferences. $M^{\sim}$ forces all behavior in line with the assumption of single peaked preferences. According to $M^{\sim}$ Agent 2 cannot behave in the way outlined above: by choosing o_1 at p, Agent 2 bars himself from getting matched with any house larger than House 2.[9]

The algorithm of the circle-crawler typically goes on for many more steps than the algorithm of the crawler. Quicker mechanisms for $\hat{\Omega}$ proceed in two steps: first use a fixed rule to break all indifferences in $\succsim$, then use the resulting profile in the crawler. But are such quick mechanisms efficient, strategyproof, and individually rational? A **tie-breaker** $TB : \hat{\Omega} \to \hat{\Omega}^l$ transforms any $\succsim \in \hat{\Omega}$ to a linear order $\succsim^{TB} \in \hat{\Omega}^l$ such that $j \succ_i j'$ implies $j \succ_i^{TB} j'$ for all $i, j, j' \in N$.[10] The mechanism $C^{TB} : \hat{\Omega} \to \mathcal{M}$ is then defined such that $C^{TB}(\succsim) := C(\succsim^{TB})$ for all $\succsim \in \hat{\Omega}$. So C^{TB} maps each profile $\succsim \in \hat{\Omega}$ to the outcome of the crawler at $\succsim^{TB}$, the profile obtained by applying the tie-breaker TB to $\succsim$. I next show that any such mechanism C^{TB} is strategyproof and individually rational. However, no such mechanism is efficient.

Proposition 2 *Fix any tie-breaker $TB : \hat{\Omega} \to \hat{\Omega}^l$. Then the mechanism C^{TB} is strategyproof and individually rational. If $n \geq 3$, then C^{TB} is not efficient.*

[9] I am grateful to a referee for pointing out that S is not even dominant in M' and that the introduction of variable upper bounds $\overline{o}_h$ not only renders S dominant in $M^{\sim}$ but obviously dominant.

[10] Note that the tie-breaker TB may use different standards for different agents. Say $N = \{1, 2, 3\}$. Then we may have $1 \succ_1^{TB} 2 \succ_1^{TB} 3$ and $3 \succ_2^{TB} 2 \succ_2^{TB} 1$ even though 1, 2, and 3, are indifferent according to both $\succsim_1$ and $\succsim_2$.

Proof Fix an agent i, a profile $\succsim \in \hat{\Omega}$ and a preference $\overline{\succsim}_i \in \hat{\Omega}_i$. Since C is individually rational we have $C(\succsim^{TB})(i) \succsim_i^{TB} i$. The definitions of TB and C^{TB} then imply that $C^{TB}(\succsim)(i) = C(\succsim^{TB})(i) \succsim_i i$ and C^{TB} is individually rational. Next $C(\succsim^{TB})(i) \succsim_i^{TB} C(\overline{\succsim}_i^{TB}, \succsim_{-i}^{TB})(i)$ holds since C is strategyproof. By the definition of C^{TB} we then have $C^{TB}(\succsim)(i) \succsim_i^{TB} C^{TB}(\overline{\succsim}_i, \succsim_{-i})(i)$. Finally the definition of TB yields $C^{TB}(\succsim)(i) \succsim_i C^{TB}(\overline{\succsim}_i, \succsim_{-i})(i)$ and C^{TB} is strategyproof.

To see that C^{TB} is not efficient with $n \geq 3$, fix a profile of preferences $\succsim \in \hat{\Omega}$ where $\succsim_2$ top ranks the houses $\{1, 2, 3\}$ and finds all other houses unacceptable. Agents 1 and 3 have linear preferences. If there is any agent $i > 3$, then i is the only $\succsim_i$-acceptable house. Since any individually rational matching at $\succsim$ matches any agent $i > 3$ with house i we can ignore any such agent $i > 3$. Efficient and individually rational matchings at $\succsim$ can therefore be denoted as a three-component vectors with the understanding that the i-th component represents agent i's match. The vector $(1, 3, 2)$, for example, denotes the matching where Agents 2 and 3 swap houses, while all other agents keep their houses. Given that Agents 1 and 3 strictly rank all houses we only need to know the tie-breaking rule for Agent 2.

Case 1: $\succsim_2^{TB}$ top ranks House 1. Let $\succsim_1$ and $\succsim_3$ respectively top rank Houses 2 and 1, so that $(2, 3, 1)$ is the unique efficient and individually rational matching at $\succsim$. Since Step 1 of $C(\succsim^{TB})$ matches Agent 2 with House 1, $C(\succsim^{TB})$ is not efficient at $\succsim$.

Case 2: $\succsim_2^{TB}$ top ranks House 2. Let $\succsim_1$ and $\succsim_3$ respectively top rank Houses 2 and 3, so that $(2, 1, 3)$ is the unique efficient and individually rational matching at $\succsim$. Since Step 1 of $C(\succsim^{TB})$ matches Agent 2 with House 2, $C(\succsim^{TB})$ is not efficient at $\succsim$.

Case 3: $\succsim_2^{TB}$ top ranks House 3. Let $\succsim_1$ and $\succsim_3$ respectively top rank Houses 3 and 2, so that $(3, 1, 2)$ is the unique efficient and individually rational matching at $\succsim$. Since Steps 1 and 2 of $C(\succsim^{TB})$ match Agents 3 and 2 with Houses 2 and 3, $C(\succsim^{TB})$ is not efficient at $\succsim$. $\square$

7. Conclusion

The crawler is a new efficient, strategyproof and individually rational matching mechanism on the domain of single-peaked preferences. Differently from Gale's top trading cycles the crawler can, on the domain of single-peaked preferences be implemented in obviously dominant strategies following Li [9].

The crawler is similar to a shell-swapping mechanism used by hermit crabs.[11] Hermit crabs need shells for protection but do not grow their own. They instead salvage empty shells of other animals. To be useful such shells may neither be too large nor too small: A hermit crab that cannot completely retract into his shell is more likely to be eaten, the same crab looses agility if its shell is too large. To accommodate their growth hermit crabs then periodically require new shells, which they may acquire via "synchronous vacancy chains". Such a vacancy chain starts with a crab who wants to move to a larger shell but happens upon a vacant shell that is too large for its own needs. This crab may wait next to the vacant shell. More crabs, who also find the vacant shell too large, may appear on the scene and wait. As they wait the crabs line up in order of their size, each holding tight to the next largest shell. Once a crab, for whom the vacant shell is a good fit, arrives, this crab occupies the vacant shell. The largest crab waiting, that is the crab at the helm of the line, moves into the shell cast away by the newcomer. A cascade of moves ensues: each crab moves into the shell it held onto while waiting. Three things stand out: hermit crabs firstly appear to have single peaked preferences over shells/houses of different sizes. Hermit crabs must, secondly, be considered as cognitively limited, so that Li's [9] argument in favor of obviously strategyproof mechanisms applies. Finally, synchronous vacancy chains much more closely resemble the crawler than Gale's top trading cycles.

While I have shown that the crawler is implementable in obviously dominant strategies, I have not provided a characterization of all efficient and individually rational mechanisms for the domain of single-peaked preferences

[11] I would like to thank Yannai Gonczarowski for pointing out the hermit crab housing problem to me. The description of synchronous vacancy chains follows Rotjan, Chabot and Lewis [14].

$\hat{\Omega}^l$ that are implementable in obviously dominant strategies. Clearly the dual crawler that screens houses in descending order is also implementable in obviously dominant strategies. With just 3 agents even Gale's top trading cycles is implementable in obviously dominant strategies. I conjecture that any social choice function on the matching domain with single-peaked preferences that is implementable in obviously dominant strategies is a combination of these three mechanisms.

When more than one mechanism appears normatively reasonable, then they can be ranked by their computational complexity. In that vein Plaxton [12] proposes a set of efficient, strategyproof, and individually rational mechanisms for (standard) housing markets that are less computationally complex than Jaramillo and Manjunath's [8] and Alcalde-Unzu and Molis' [2] mechanisms. Without Ma's [10] uniqueness result (Theorem 2) we could use the criterion of computational complexity to distinguish between the crawler, Gale's top trading cycles and possibly yet more efficient, strategyproof, and individually rational mechanisms. Since the crawler matches exactly one agent per step while Gale's top trading cycles matches at least one agent per step, the crawler always needs at least as many steps as Gale's top trading cycles to determine a matching. However, the steps of the crawler are less computationally complex than the steps of Gale's top trading cycles: While each agent under Gale's top trading cycles may point to any house, the crawling algorithm restricts agents to point either backwards or to the next largest house. Comparing Gale's top trading cycles and the crawler, I conjecture that the computational advantage of each step of the crawler outweighs the fact that only one agent is matched per step. I furthermore conjecture that this computational advantage extends to the domain $\hat{\Omega}$ that permits indifferences. On this domain the circle-crawler $C^{\sim}$ should be even less computationally complex than the mechanisms defined by Plaxton [12].

While no individually rational and efficient mechanism on the domain of linear preferences Ω^l is implementable in obviously dominant strategies, I have shown that the crawler satisfies these three desiderata on the subdomain of single-peaked preferences. In Bade [5] I use single-peaked preferences to escape from a different impossibility result. In that paper I define shift exchange problems as housing markets with an infinite stream of overlapping generations of agents. No efficient, strategyproof, and individually rational mechanism matches at least one agent in finite time. However if we impose that all agents preferences are single-peaked a version of the crawler is not only efficient, strategyproof, and individually rational, but also matches each agent within a finite time window around his (original) shift.

8. Appendix

Gale's top trading cycles on $\hat{\Omega}^l$, the domain of linear single-peaked preferences, is implementable in obviously dominant strategies if and only if there are at most three agents.[12] The upcoming proof of the claim that Gale's top trading cycles on $\hat{\Omega}^l$ with at least four agents is not implementable in obviously dominant strategies crucially relies on Bade and Gonczarowski's [6] gradual revelation principle. In their Theorem 1, Bade and Gonczarowsi [6] show that a social choice function is implementable in obviously dominant strategies if and only if it is implementable by an obviously incentive compatible gradual revelation mechanism. A gradual revelation mechanism is an extensive form mechanism where each action corresponds to a set of preferences. There are no simultaneous moves, singleton choice sets, or directly consecutive moves for the same agent. An agent's strategy is truthful if he, wherever possible, chooses an action that corresponds to the set of preferences containing his true preference.

Formally, the player function maps each non-terminal history h to a single player $P(h)$ who chooses from $A(h)$ with $\mid A(h) \mid > 1$. For any h and $a \in A(h)$ we have $P(h) \neq P(h,a)$. Each history h is associated with a set of preferences $\Omega_i(h) \subset \Omega_i$, with $\Omega(h) \colon = \Omega_1(h) \times \cdots \times \Omega_n(h)$. The mechanism starts with $\Omega(\emptyset) = \Omega$. If $P(h) = i$, then $\{\Omega_i(h,a)\}_{a \in A(h)}$ partitions $\Omega_i(h)$, if not then $\Omega_i(h) = \Omega_i(h,a)$ holds for each $a \in A(h)$. A strategy S_i in a gradual revelation mechanism is truthful if at each history h agent $i \colon = P(h)$ with preference $\succsim_i \in \Omega_i(h)$ chooses the action $a \in A(h)$ with $\succsim_i \in \Omega_i(h,a)$.[13] If

[12] Li [9] already showed that Gale's top trading cycles on the domain Ω^l is not implementable in obviously dominant strategies.

[13] Since this definition does not determine the agents choice at a history h with $\succsim_i \notin \Omega_i(h)$ agents may have multiple truthful strategies.

all agents follow a truthful strategy then $\Omega(h)$ describes everything that the agents revealed about their preferences up to history h. Finally the gradual revelation mechanism M is obviously incentive compatible if truthtelling is obviously dominant for each agent.

Theorem 5 *Gale's top trading cycles $G : \hat{\Omega}^l \to \mathcal{M}$ can be implemented in obviously dominant strategies if and only if there are at most three agents.*

Proof Suppose M was an obviously incentive compatible mechanism that implements G for $N = \{1, \ldots, 4\}$. Let Ω^* be the set of all preference profiles for which agents 1 and 2 want to move to larger houses while agents 3 and 4 want to move to smaller houses. So $\Omega_i^* : = \{\succsim_i \in \hat{\Omega}^l \mid j \succsim_i H \text{ implies } j > i\}$ for $i = 1, 2$ and $\Omega_i^* : = \{\succsim_i \in \hat{\Omega}^l \mid j \succsim_i H \text{ implies } j < i\}$ for $i = 3, 4$.

Claim (*) For any history h with $\Omega^* \subset \Omega(h)$, there exists an action $a \in A(h)$ such that $\Omega^* \subset \Omega(h, a)$: fixing any history on the path of a truthtelling strategy for any $\succsim \in \Omega^*$, where no agent has revealed more than the direction in which he wants to move, the agent moving at the present history will not reveal any more than the direction in which he wants to move.

To see Claim (*) fix a history h with $\Omega^* \subset \Omega(h)$, say that $i = P(h) \in \{1, 2\}$ and let j be the other agent in $\{1, 2\}$. Suppose there existed two actions $a, a' \in A(h)$ with $\succsim_i \in \Omega_i(h, a)$ and $\succsim_i' \in \Omega_i(h, a')$ for some $\succsim_i, \succsim_i' \in \Omega_i^*$ with $i + 1 \succsim_i' \{1, 2, 3, 4\}$. For $\succsim \in \Omega^*$ with $4 \succsim_j \{1, 2, 3, 4\}$, $j \succsim_3 \{1, 2, 3, 4\}$, and $3 \succsim_4 \{1, 2, 3, 4\}$ we have $G(\succsim)(i) = i$. On the other hand, $G(\succsim')(i) = i + 1$ holds for $\succsim' \in \Omega^*$ with $j + 1 \succsim_j' \{1, 2, 3, 4\}$ and $1 \succsim_3' \{1, 2, 3, 4\}$. Since $\succsim_i \in \Omega_i^*$ and $i \in \{1, 2\}$, $G(\succsim') = i + 1$ is strictly $\succsim_i$-preferred to $G(\succsim)(i) = i$, and the action a is not an obviously dominant choice for agent i with preference $\succsim_i$ at h. There must be an action $a \in A(h)$ with $\Omega_i^* \subset \Omega_i(h, a)$. Since $\Omega_{i'}(h, a) = \Omega_{i'}(h)$ holds for all $i' \neq P(h)$ we obtain $\Omega^* \subset \Omega(h, a)$. Claim (*) then holds since the same arguments apply mutatis mutandis to the case of $P(h) \in \{3, 4\}$.

Now fix any $\succsim^* \in \Omega^*$ and say h^* is the (unique) terminal history reached if all agents follow a truthtelling strategy at $\succsim^*$. Since no agent has revealed anything before the game starts, Ω^* is a subset of $\hat{\Omega}^l = \Omega(\emptyset)$. The inductive application of claim (*) to all histories h with $\succsim^* \in \Omega(h)$ implies $\Omega^* \subset \Omega(h)$ for any such history, in particular $\Omega^* \subset \Omega(h^*)$. So if each agent i with preference $\succsim_i^*$ follows a truthful strategy, then each agent i at most reveals the direction in which he would like to move. But G is not constant on Ω^*: G, for example, maps the two profiles $\succsim, \succsim' \in \Omega^*$ constructed in the preceding paragraph to two different matchings. So the terminal history h^* cannot be mapped to a unique outcome - a contradiction.

Li [9] already showed that the direct revelation mechanism implements Gale't top trading cycles with two agents in obviously dominant strategies. So it only remains to be seen that Gale's top trading cycles with three agents $\{1, 2, 3\}$ with single peaked preferences is implementable in obviously dominant strategies. Consider the extensive form mechanism M^3 where Agent 2 first reveals his preferences. Agent 2 gets the middle house (House 2) if he most likes it. Agents 1 and 3 are then asked whether they would like to swap houses. If both agree, they swap houses, otherwise each agent is matched with the house he was endowed with. If Agent 2 top ranks House 1 (House 3), then Agent 1 (3) reveals his preference. If Agent 1 (3) ranks the smallest (the largest) house at the top, each agent is matched with the house he was endowed with. If Agent 1 (3) top ranks the middle house, Agents 1 (3) and 2 swap houses, while Agent 3 (1) is matched with the largest house. In the final case where Agent 1 (3) prefers larger (smaller) houses Agent 3's (1's) preference is elicited. Once all preferences are known the outcome is established using Gale's top trading cycles.

To see that truthtelling is obviously dominant in M^3 first consider Agent 2. If he most likes his own house, he gets matched with it if he truthfully reveals his preference. If he prefers the smallest house the he gets the smallest or the middle house if he truthfully reveals his preference, whereas he gets with the middle or the largest house if he deviates. The same holds mutatis mutandis for the case that he most prefers the largest house. So the truthful revelation is obviously dominant for Agent 2.

Now consider Agent 1 as a second player. If Agent 2 ranks the middle house at the top, it is obviously dominant for Agent 1 to reveal his preference in the the simple swapping mechanism with the smallest and largest houses. If Agent 2 most likes the smallest house, Agent 1 gets matched with his most preferred house if he truthfully reveals his preference and if his most preferred house is either the smallest or the middle house. In the remaining case where Agent 1 most prefers the largest house, Agent 1 either gets matched with that house or the middle house if he truthfully reveals his preference, if he deviates he either gets matched with the middle or the smallest house. So the truthful revelation of his preference is obviously dominant for Agent 1 if he is the second player to reveal his preference.

If Agent 1 is the third player to reveal his preference, then it is already known that Agent 2 wants to move to a larger house and that Agent 3 wants to move to the smallest possible house. Given that Gale's top trading cycles is used to determine all agents' matches the announcements by Agents 2 and 3 imply that Agent 1 gets matched with the house he declares to be his favorite choice. So truthtelling is also in this final case obviously dominant. Mutatis mutandis we see that truthtelling is also obviously dominant for Agent 3. □

References

[1] Anno, H.: "A short proof for the characterization of the core in housing markets", *Economics Letters*, (2015), pp 66-67.

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

[3] Arribillaga R.P, Masso, J. and Neme, A.: "Not All Majority-based Voting Mechanisms Are Obviously Strategy-proof" *mimeo* (2017) Universitat Autonoma de Barcelona.

[4] Aziz, H. and de Keijzer, B.: "Housing markets with indifferences: A tale of two mechanisms", *In Proceedings of AAAI'12*.

[5] Bade, S.: "Shift Exchange Problems", *mimeo*, (2017), Royal Holloway.

[6] Bade, S. and Gonczarowski, Y,: "Gibbard-Satterthwaite Success Stories and Obvious Strategyproofness" *mimeo*, (2016), The Hebrew University, discussion paper 704.

[7] Damamme, A., Beynier, A., Chevaleyre, Y., and Maudet, N.: "The Power of Swap Deals in Distributed Resource Allocation", (2015), *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems*, pp. 625-633.

[8] Jaramillo, P. and Manjunath, V.: "The dfference indifference makes in strategy- proof allocation of objects", (2012), *Journal of Economic Theory*, pp. 1913 - 1946.

[9] Li, S.: "Obvioulsy Strategyproof Mechanisms", (2017), *American Economic Review*, 107, pp. 3257-87.

[10] Ma, J.: "Strategyproofness and the Strict Core in a Market with Indivisibilites", *International Journal of Game Theory*, 23, (1995), pp. 75-83.

[11] Moulin, H.: *Axioms of Cooperative Decision Making*, Cambridge University Press, 1991.

[12] Plaxton, C. G.: "A simple family of top trading cycles mechanisms for housing markets with indifferences", *Proceedings of the 24th International Conference on Game Theory*, Stony Brook, (2013).

[13] Roth, A.: "Incentive Compatibility in a Market with Indivisibilities" *Economics Letters*, 9, (1982), pp. 127-132.

[14] Rotjan, R., J. Chabot and S. Lewis: "Social context of shell acquisition in coenobita clypheatus hermit crabs" *Behavioral Ecology*, 21, (2010), pp. 639–646.

[15] Saban D. and Sethuraman, J.: "House allocation with indifferences: A generalization and a unified view", *Proc. EC '13* ACM, (2013) pp. 803-820. 30

[16] Sethuraman, J.: "An alternative proof of a characterization of the TTC mechanism" *Operations Research Letters* (2016) 107-108.

[17] Shapley, L. and H. Scarf: "On Cores and Indivisibility", *Journal of Mathematical Economics*, 1, (1974), pp. 23 - 37.

[18] Svensson, L.-G. "Strategy-proof allocation of indivisible goods", *Social Choice and Welfare*, 16, (1991), pp. 557-567.