Random Serial Dictatorship: The One and Only.

Sophie Bade*

January 13, 2016

Abstract

Fix a Pareto optimal, strategy proof and non-bossy deterministic matching mechanism and define a random matching mechanism by assigning agents to the roles in the mechanism via a uniform lottery. Given a profile of preferences, the lottery over outcomes that arises under the random matching mechanism is identical to the lottery that arises under random serial dictatorship, where the order of dictators is uniformly distributed. This result extends the celebrated equivalence between the core from random endowments and random serial dictatorship to the grand set of all Pareto optimal, strategy proof and non-bossy matching mechanisms.

1 Introduction

A matching problem consists of a finite set of agents, a finite set of indivisible objects, henceforth called houses, and a profile of all agents' preferences over all houses. A matching is a maximal set of agent-house pairs. Mechanisms map the sets of preference profiles to the set of matchings, with serial dictatorship and Gale's top trading cycles (GTTC) the two most prominent

*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.

examples. Serial dictatorship matches one agent - the first dictator - to his most preferred house in the grand set of all houses. A next agent - the second dictator - is matched with his most preferred house in remainder, and so forth. A GTTC starts with an initial matching called the endowment and requires that there are equally many agents and houses. In a first round each agent points to his most preferred house and each house points to its owner according to the endowment. At least one pointing cycle forms and each agent in such a cycle is matched with the house he points to and exits. The same procedure is then repeated until all agents are matched.

Serial dictatorship and GTTC satisfy three central properties. They are *Pareto optimal*: they map each profile of preferences to a matching that is Pareto optimal at that profile. They are *strategy proof*: there is no profile of preferences and agent such that this agent could benefit from reporting a false preference, keeping the reports of all other agents fixed. They are *non-bossy*: there is no profile of preferences, agent and deviation for this agent, such that the agent obtains the same house under the original profile and the deviation (keeping all other agents' preferences fixed) while the matches of some other agents change. Call any mechanism that satisfies these three properties *good*.

No good mechanism treats equals equally: when two different agents submit the same preference they end up with different houses. Randomization fixes this flaw. Consider drawing the sequence in which agents become dictators in a serial dictatorship from a uniform distribution on all such sequences. The resulting random matching mechanism is known as *random serial dictatorship* or *random priority*. Any two agents who submit the same preferences in random serial dictatorship face the same lottery over houses. The same method can be used to *symmetrize* any good mechanism. Instead of assigning one agent, say Anton, to assume the role of agent 1 in a mechanism, and assigning Betty to the role of agent 2, and so forth, the symmetrization of the mechanism uses a uniform lottery over all possible such assignments.

The set of all good mechanisms is large. So this method would seem to generate many different random matching mechanisms, viewed as mappings from profiles of preferences to lotteries over matchings. This is not the case: Theorem 1 shows that the symmetrization of any good mechanism coincides with random serial dictatorship.

The first observations of this sort relate to GTTC: Abdulkadiroglu and Sönmez [1] and Knuth [8] independently proved the identity of random serial dictatorship and the symmetrization of GTTC, known as the *core from random endowments*. Both their proofs start by fixing an arbitrary profile of preferences. They then construct a bijection between the set of all sequences of agents as dictators and the set of initial endowments such that the outcome of the serial dictatorship with a given sequence equals the outcome of GTTC with the image of this given sequence. The bijection ensures that (at the fixed profile of preferences) the number of sequences with which serial dictatorship yields some fixed matching equals the number of initial endowments with which GTTC yields the same matching. So the probability of that matching under random serial dictatorship (the proportion of sequences with which serial dictatorship yields the matching) equals the probability of the matching under the core from random endowments (the proportion of initial endowments with which GTTC yields the matching). This result has been extended to increasingly larger sets of good mechanisms by Pathak and Sethuraman [11], Caroll [5], and Lee and Sethuraman [9].

There are three major differences between Theorem 1 and the preceding results. First, I show the equivalence for *all* good mechanisms. Second, Theorem 1 holds whether there are equally many agents and houses or not; in fact my proof does not distinguish between these cases. These two differences are made possible by the third innovation: a new simple strategy of proof.

This strategy relies on the construction of a sequence of good mechanisms $M^0, M^1, \cdots, M^K$ from an arbitrary good mechanism M^0 to a serial dictatorship M^K, such that any two consecutive mechanisms M^k, M^{k+1} have identical symmetrizations. I follow the bijective strategy pioneered by Abdulkadiroglu and Sönmez [1] and Knuth [8] to establish the identity of these symmetrizations. To make this step as straightforward as possible the difference between any two consecutive mechanism is kept at a minimum.

The proof crucially relies on Theorem 2 which characterizes the set of good mechanisms as the set of trading and braiding mechanisms. Like GTTC, trading and braiding mechanisms use rounds of trade in pointing cycles to determine matchings. In any trading round, houses point to their owners and owners point to their most preferred houses. Matchings are obtained through the consecutive elimination of trading cycles. Trading and braiding mechanisms generalize three aspects of GTTC. One agent might own multiple houses, a feature introduced by Papai [10]. There is a second form of control called brokerage, introduced by Pycia and Unver [12]. At any normal trading round there is at most one brokered house. A brokered house points to its broker (just as an owned house points to its owner). The broker points to the house he most likes among the owned houses. Finally, if there are only three houses left a trading and braiding mechanism might terminate in a *braid*. Braids, in turn are good mechanisms that aim to maximally avoid some specific matches between these three houses and three agents.

This paper is the first to characterize the set of all good mechanisms. While Pycia and Unver [12] lays out the same goal it arrives at an incorrect characterization: the set of Pycia and Unver [12] trading cycles mechanisms represents a a strict subset of the set of all good mechanisms. If one excludes braids from the set of good mechanisms one obtains Pycia and Unver's [12] set of trading cycles mechanisms. In response to the present characterization Pycia and Unver [13] presents a revised set of trading cycles mechanisms that accommodates braids (now called "three broker* mechanisms").

My characterization in Theorem 2 and its proof owe a large debt to Pycia and Unver [12]. I build on Pycia and Unver's [12] ingenious idea to use trading processes with two forms of control, ownership and brokerage, to frame good mechanisms. However, I arrive at a different conclusion. While the Pycia and Unver [12] trading cycles mechanisms allow for at most one broker in a round of trade, I show that there can be rounds with three brokers. In this case the mechanism terminates as a braid. All parts of the proof pertaining to braids originate with the current treatment.

My paper is concerned with a fully symmetric treatment of all agents. It does not speak to existing results on random matching mechanisms that treat different agents differently. Ekici [6], for example, considers two different matching mechanisms that both respect initial (private) allocations. He shows that a uniform randomization of any additional "social" endowment of houses in the two mechanisms yields the same random matching mechanism. Similarly Carroll [5] considers the case of matching mechanisms in which agents are partitioned into groups and provides an equivalence result for the case in which agents are treated symmetrically only within groups. Lee and Sethuraman's [9] equivalence results also cover the case where agents are partitioned into groups that are treated asymmetrically.

2 A sketch of the proof: GTTC

To apply my strategy of proof to GTTC let M^0 be the GTTC where each agent i is endowed with house h_i. Construct a sequence of mechanisms $M^1, \ldots, M^K$ such that each mechanism is derived from its predecessor by consolidating the ownership of exactly two agents. The sequence terminates with a serial dictatorship M^K when all ownership has been maximally consolidated. M^1 is identical to GTTC except that agent 1 owns h_1 and h_2. Once agent 1 exits, agent 2 inherits the unmatched house in $\{h_1, h_2\}$.

To see that the symmetrizations of M^0 (GTTC) and M^1 are identical, fix a profile of preferences. Fix an assignment of agents to roles (initial endowment). To keep things simple, let this endowment be the original one where i owns h_i. Suppose that some cycle at the start of M^0 involves agent 1 but not agent 2. In this case, agent 1 is part of the same cycle at the start of M^1; the only difference is that agent 1 additionally owns house h_2 under M^1. But for the formation of cycles under the given preferences this difference does not matter. Once the cycle involving 1 is matched, M^0 and M^1 continue identically (given that 2 inherits house h_2 under M^1). So M^0 and M^1 yield the same outcome under the given profile of preferences and assignment of agents to roles.

Now assume that in the first round of M^0 there is a single pointing cycle. Say this cycle, $2 \to h \to i \to h^* \cdots \to h_2 \to 2$, involves 2 but not 1 who prefers h^* to all other houses. Every agent in the cycle is matched to the house he points to, $i \neq 1$, in particular, is matched with h^*. Under M^1, 1 owns h_2, the cycle $1 \to h^* \cdots \to h_2 \to 1$ forms in the first round, and 1 is matched to h^*. If we switch the roles of agents 1 and 2 in M^1 2 owns h_1 and h_2 at the start of M^1. The cycle $2 \to h \to i \to h^* \cdots \to h_2 \to 2$ which forms in the first round of M^0 with the original assignment also forms in the first round of M^1 with the new assignment. The fact that 2 additionally owns

h_1 is irrelevant. Once the agents and houses in this cycle exit, M^1 with the new assignment is identical to M^0 with the original assignment (given that 1 inherits house h_1).

The preceding two paragraphs illustrate the bijections used to prove the identity of symmetrizations. In the first case where M^0 and M^1 map the given profile of preferences to the same outcome, the bijection maps the original assignment of roles onto itself. In the second case where M^0 and M^1 map the profile of preferences to the different outcomes, the new assignment is derived from the original by swapping the roles of agents 1 and 2. It turns out that the above reasoning also applies when neither 1 nor 2 is matched in the first round and when both take part in the same cycle under M^0. So there exists a bijection between the assignments of roles such that M^0 with the original assignment and M^1 with the image of that assignment map the given profile of preferences to the same matching. The existence of such a bijection implies that the symmetrizations of M^0 and M^1 are identical. The remaining mechanisms in the sequence $M^2, \cdots, M^K$ are constructed via the further consolidation of ownership. At the start of M^2 agent 1 owns h_1, h_2 and h_3, at the start of M^3 he also owns h_4. The process of consolidation terminates when M^0 has been transformed into a serial dictatorship M^K.

To apply the consolidation strategy to any good mechanism M^0 the above arguments have to be extended to the case where agents may own multiple houses. A somewhat different approach is needed to absorb brokers and to replace braids with serial dictatorships. One feature that all these cases share with the simple case of M^0, M^1 is that the bijections between assignments of roles which show that the symmetrization of two consecutive mechanisms M^k and M^{k+1} are identical, switch the roles of at most three agents.

3 Definitions

A housing problem consists of a set of agents $N := \{1, \cdots, n\}$, a finite set of houses H and a profile of preferences $R = (R_i)_{i=1}^n$. The option to stay homeless $\emptyset$ is always available: $\emptyset \in H$. Preferences R_i are linear orders[1] on

[1]So hR_ih' and $h'R_ih$ together imply $h = h'$.

H and each agent prefers any house to homelessness, so $hR_i\emptyset$ holds for all $i \in N, h \in H$. The notation hR_iH' means that agent i prefers h to each house in H'. The set of all profiles R is denoted by $\mathcal{R}$. The **restriction** of R to some $N' \subset N$ and $H' \subset H$ is a profile of preferences $\overline{R}$ (defined for N' and H') with $h\overline{R}_ig \Leftrightarrow hR_ig$ for all $h, g \in H'$ and $i \in N'$.

Submatchings match subsets of agents to at most one house each. A **submatching** is a function $\nu : N \to H$ such that $\nu(i) = \nu(j)$ and $i \neq j$ imply $\nu(i) = \emptyset$. The sets of agents and houses matched under ν are $N_\nu \colon = N \setminus \nu^{-1}(\emptyset)$ and $H_\nu \colon = \nu(N_\nu)$. When $\nu(i) \neq \emptyset$ then i is matched to house $\nu(i)$; $\overline{N}_\nu \colon = N \setminus N_\nu$ and $\overline{H}_\nu \colon = H \setminus H_\nu$ are the sets of agents and houses that remain unmatched under ν. Any submatching ν can also be interpreted as a set of agent-house pairs $\{(i,h) : \nu(i) = h \neq \emptyset\}$. For two submatchings ν and ν' with $N_\nu \cap N_{\nu'} = \emptyset = H_\nu \cap H_{\nu'}$ the submatching $\nu \cup \nu' : N_\nu \cup N_{\nu'} \to H_\nu \cup H_{\nu'}$ is defined by $(\nu \cup \nu')(i) = \nu(i)$ if $i \in N_\nu$ and $(\nu \cup \nu')(i) = \nu'(i)$ otherwise. A submatching ν is considered **maximal** (**minimal**) in a set of submatchings if there exists no ν' in the set such that $\nu \subsetneq \nu'$ ($\nu' \subsetneq \nu$). Any maximal submatching (in the set of all submatchings) is a **matching**, so μ is a matching if and only if $H_\mu = H$ or $N_\mu = N$ (or both) hold. The sets of all matchings and of all lotteries over matchings are denoted $\mathcal{M}$ and $\Delta\mathcal{M}$ respectively.

A (deterministic) **mechanism** is a function $M : \mathcal{R} \to \mathcal{M}$ where i is matched with $M(R)(i)$ under M at R. A mechanism M is **Pareto optimal** if for no R there exists a matching $\mu \neq M(R)$ such that $\mu(i)R_iM(R)(i)$ for all i.[2] A mechanism M is **strategy proof** if $M(R)(i)R_iM(R'_i, R_{-i})(i)$ holds for all triples R, R'_i, i: declaring one's true preference is a weakly dominant strategy. A mechanism M is **non-bossy** if $M(R)(i) = M(R'_i, R_{-i})(i)$ implies $M(R) = M(R'_i, R_{-i})$ for all triples R, R'_i, i, so an agent can only change someone else's match if he also changes his own match. A mechanism M is **good** if it is Pareto optimal, strategy proof and non-bossy.

Let P be the set of all permutations $p : N \to N$. The permutation involving only agents j and j' is denoted (j, j'), so $(1,2)(2) = 1$, $(1,2)(1) = 2$, and $(1,2)(i) = i$ for $i \neq 1, 2$. Abusing notation p or p^{-1} also stands for the

[2]Since all R_i are linear some i^* must strictly prefer $\mu(i^*)$ to $M(R)(i^*)$ for μ to differ from $M(R)$ and for $\mu(i)R_iM(R)(i)$ to hold for all i.

restriction of p or p^{-1} for which some given composition is well-defined. So for any submatching ν, $\nu \circ p$ denotes the more precise and cumbersome $\nu \circ \overline{p}$ where $\overline{p}$ is the restriction of p that has N_ν as its image. Similarly $\nu \circ p$ with $p = (1, 2)$, $1 \notin N_\nu$, and $2 \in N_\nu$ is such that $\nu(p(i)) = \nu(i)$ for all $i \in N_\nu \setminus \{2\}$ and $\nu(p(1)) = \nu(2)$.

For any mechanism M and any permutation p define the **permuted mechanism** $p \odot M : \mathcal{R} \to \mathcal{M}$ via $(p \odot M)(R)(i) = M(R_{p(1)}, \cdots, R_{p(n)})(p^{-1}(i))$. The permutation p assigns each agent in N to a "role" in the mechanism, such that agent $p(i)$ under $p \odot M$ assumes the role that agent i plays under M.[3] If $S : \mathcal{R} \to \mathcal{M}$ is the serial dictatorship with agent i as the ith dictator, then $p(i)$ is the i-th dictator under $p \odot S$. To calculate $(p \odot S)(R)$ we need to substitute $p(i)$'s preference for agent i's preference to obtain the new profile of preferences $(R_{p(1)}, \cdots, R_{p(n)})$. Under $S(R_{p(1)}, \cdots, R_{p(n)})$ agent 1 is matched $p(1)$'s most preferred house. Under $(p \odot S)(R)$ this house is matched with $p(1)$: $(p \odot S)(R)(p(1)) = S(R_{p(1)}, \cdots, R_{p(n)})(p^{-1}(p(1))) = S(R_{p(1)}, \cdots, R_{p(n)})(1)$.

A **(random matching) mechanism** is a function that maps the set of preference profiles $\mathcal{R}$ to the set of all lotteries over matchings $\Delta\mathcal{M}$: The **symmetrization** of a mechanism $M : \mathcal{R} \to \mathcal{M}$ is a random matching mechanism $\Delta M : \mathcal{R} \to \Delta\mathcal{M}$ that calculates the probability of matching μ at the profile R as the probability of a permutation p with $\mu = (p \odot M)(R)$ under the uniform distribution on P. So we have

$$\Delta M(R)(\mu) :\, = \frac{|\{p : (p \odot M)(R) = \mu\}|}{n!}.$$

Abdulkadiroglu and Sönmez [1] call ΔM a **random serial dictatorship** if M is a serial dictatorship and the **core from random endowments** if M is GTTC.

Definition 1 *Two (deterministic) mechanisms M and M' are* ***s-equivalent***[4] *if $\Delta M = \Delta M'$.*

[3] The symbol $\odot$ is chosen as a reminder that $p \odot M$ arises out of a non-standard composition of the permutation p and the mapping M: $\odot$ is similar to but different from $\circ$, the standard operator for compositions.

[4] The letter "s" is a reminder that symmetrizations are the base of s-equivalence.

4 The Result

Theorem 1 *Any good mechanism is s-equivalent to serial dictatorship.*

For the proof fix an arbitrary good mechanism M^0 and construct a sequence of of mechanisms $M^1, \cdots, M^K$ such that $\Delta M^k = \Delta M^{k+1}$ holds for $0 \leq k < K$ and M^K is a serial dictatorship. The s-equivalence of any two adjacent M^k, M^{k+1} implies the s-equivalence of M^0 and M^K. The minimal difference between M^k and M^{k+1} is crucial. It simplifies the task to find a bijection $f : P \to P$ with $(p \odot M^k)(R) = (f(p) \odot M^{k+1})(R)$ for all $p \in P$. Such a bijection exists if and only if $|\{p : (p \odot M^k)(R) = \mu\}| = |\{p : (p \odot M^{k+1})(R) = \mu\}|$ holds for all R and all matchings μ. The equality $|\{p : (p \odot M^k)(R) = \mu\}| = |\{p : (p \odot M^{k+1})(R) = \mu\}|$ is, in turn, equivalent to $\Delta M^k(R)(\mu) = \Delta M^{k+1}(R)(\mu)$, implying that a bijection f exists if and only if M^k and M^{k+1} are s-equivalent.

5 Trading and braiding mechanisms

My construction of marginally different good mechanisms M^k, M^{k+1} relies on the characterization of the set of all good mechanisms as trading and braiding mechanisms. Just like GTTC, trading and braiding mechanisms use sequential trading rounds to determine matchings. In such trading rounds owned houses point to their owners, and owners point to their most preferred houses. Agents in cycles are matched to the houses they point to. If a matching results the mechanism terminates, if not a new round ensues. There are three differences between GTTC and trading and braiding mechanisms. Following Papai's [10] hierarchical exchange mechanisms trading and braiding mechanisms permit the ownership of multiple houses. Following Pycia and Unver [12] trading cycles mechanisms there is a second form of control in trading and braiding mechanisms: houses might be "brokered". At any given trading round there is at most one broker and he brokers exactly one house. This house points to the broker and the broker points to his most preferred house among the owned ones. Finally when there are only three houses left a trading and braiding mechanism might terminate as a braid. Braids are good mechanisms that match three houses to three agents with the aim to obtain a matching that maximally differs from some fixed "avoidance matching".

The **braid** $B^\omega : \mathcal{R} \to \mathcal{M}$ is a mechanism for a problem with exactly three houses and at least as many agents. It is fully defined through the **avoidance matching** ω. Outcomes $B^\omega(R)$ are chosen to avoid matching i to $\omega(i)$ while keeping the set of matched agents equal to the set of agents matched under ω. For any R let $\overline{PO}(R)$ be the set of Pareto optima μ with $N_\omega = N_\mu$ and let $Mini(R) := argmin_{\mu \in \overline{PO}(R)} \mid \{i : \mu(i) = \omega(i)\} \mid$ be the subset of all $\mu \in PO(R)$ that minimally coincide with ω. If $Mini(R)$ is a singleton let $Mini(R) = \{B^\omega(R)\}$. If not, at least two agents in N_ω must rank some house $\omega(i)$ at the top. If only one agent $j \neq i$ ranks $\omega(i)$ at the top then $B(R)$ is the unique element in $Mini(R)$ that matches j to $\omega(i)$. If both agents $j \neq i$ rank $\omega(i)$ at the top, then $B^\omega(R)$ is i's preferred matching in $Mini(R)$.

To concretely illustrate braids let $H = \{e, f, g\}$ and $N = \{1, 2, 3\}$. Given $\mid H \mid = \mid N \mid = 3$ it is convenient to denote matchings as vectors with the understanding that the i-th component represents agent i's match. Moreover, $N_\omega = N_\mu$ is satisfied by any matching μ and can be ignored. Let $\omega := (e, f, g)$. A matching μ is **maximally avoidant** if $\mu(i) \neq \omega(i)$ for all $i \in N$. There are exactly two such matchings: $\omega' := (g, e, f)$, and $\omega'' := (f, g, e)$. If all three agents rank e at the top and f at the bottom under R then $Mini(R) = \{\omega', \omega''\}$. At least two agents rank house $e = \omega(1)$ at the top. Since 2 and 3 both rank $\omega(1)$ at the top, agent 1's preference of $g = \omega'(1)$ over $f = \omega''(1)$ implies $B(R) = \omega'$. Under (R_2', R_{-2}) where R_2' ranks g at the top and e at the bottom there are four Pareto optima: $\omega, (e, g, f), (g, f, e)$ and ω''. So $Mini(R) = \{\omega''\}$ and $B(R)$ equals ω''.

To formally define trading and braiding mechanisms I use Pycia and Unver's [12] notational system for matching mechanisms that determine outcomes via trading processes. In this parsimonious system mechanisms are defined via sets of control rights functions. A **control rights function** at some submatching ν $c_\nu : \overline{H}_\nu \to \overline{N}_\nu \times \{o, b\}$ assigns control rights over any unmatched house to some unmatched agent and specifies a type of control. If $c_\nu(h) = (i, x)$, then agent i **controls** house h at ν. If $x = o$, then i **owns** h; if $x = b$ he **brokers** h. Control rights functions satisfy the following three criteria:

(C1) If more than one house is brokered, then there are exactly three houses and they are brokered by three different agents.
(C2) If exactly one house is brokered then there are at least two owners.
(C3) No broker owns a house.

A **control rights structure** c maps a set of submatchings ν to control rights functions c_ν. For now assume that c is specifies all c_ν for all submatchings ν. The following algorithm uses c to map any profile of preferences R to a matching.[5]

Initialize with $r = 1$, $\nu_1 = \emptyset$

Round r: only consider the remaining houses and agents $\overline{H}_{\nu_r}$ and $\overline{N}_{\nu_r}$.

Braiding: If more than one house is brokered under c_{ν_r}, terminate the process with $M(R) = \nu_r \cup B^\omega(\overline{R})$ where the avoidance matching ω is defined via $c_{\nu_r}(\omega(i)) = (i, b)$ and $\overline{R}$ is the restriction of R to $\overline{H}_{\nu_r}$ and $\overline{N}_{\nu_r}$. If not, go on to the next step.

Pointing: Each house points to the agent who controls it, so $h \in \overline{H}_{\nu_r}$ points to $i \in \overline{N}_{\nu_r}$ with $c_{\nu_r}(h) = (i, \cdot)$. Each owner points to his most preferred house, so owner $i \in \overline{N}_{\nu_r}$ points to house $h \in \overline{H}_{\nu_r}$ if $hR_i\overline{H}_{\nu_r}$. Each broker points to his most preferred owned house, so broker $i_b \in \overline{N}_{\nu_r}$ with $c_{\nu_r}(h_b) = (i_b, b)$ points to house $h \in \overline{H}_{\nu_r} \setminus \{h_b\}$ if $hR_{i_b}\overline{H}_{\nu_r} \setminus \{h_b\}$.

Cycles: Select at least one cycle. Define ν such that $\nu(i) := h$ if i points to h in one of the selected cycles.

Continuation: Let $\nu_{r+1} := \nu_r \cup \nu$. If ν_{r+1} is a matching terminate the process with $M(R) = \nu_{r+1}$. If not, continue with round $r + 1$.

[5]Without any further conditions, the algorithm may map some R to multiple matchings. Consider the problem $H = \{e, g\}$, $N = \{1, 2, 3\}$, $c_\emptyset(e) = (1, o)$, $c_\emptyset(g) = (2, o)$, $c_{\{(1,e)\}}(g) = c_{\{(2,g)\}}(e) = (3, o)$ and R such that 1 and 2 respectively rank e and g highest. In the first round two cycles form. Removing both at once the matching $\{(1, e), (2, g)\}$ results. Removing only $\{(1, e)\}$, the cycle involving g and 3 forms in the next round yielding the matching $\{(1, e), (3, g)\}$. Conditions (C4), (C5), and (C6) below ensure that a mechanism is defined via the application of the algorithm to c. Section 4 of the online appendix proves this claim.

A submatching ν is **reachable under** c at R if one can choose to match cycles in the above algorithm (for the given c and R) such that some round r starts with $\nu = \nu_r$. A submatching ν is c**-relevant** if it is reachable under c at some R.[6] A submatching ν is a **direct** c**-successor** of some c-relevant ν° if there exists a profile of preferences R such that ν° is reachable under c at R and ν arises out of matching a single cycle at ν°. A **trading and braiding mechanism** is a control rights structure c that maps any c-relevant submatching ν to a control rights function c_ν and satisfies requirements (C4), (C5), and (C6).

Fix a c-relevant ν° and a direct c-successor ν to ν°.

(C4) If $i \notin N_\nu$ owns h at ν° then i owns h at ν.
(C5) If at least two owners at ν° remain unmatched at ν and if $i_b \notin N_\nu$ brokers h_b at ν° then i_b brokers h_b at ν.
(C6) If i owns h at ν° and ν and if $i_b \notin N_\nu$ brokers h_b at ν° but not at ν, then i owns h_b at ν and i_b owns h at $\nu \cup \{(i, h_b)\}$.

While every good mechanism can be uniquely represented as a trading and braiding mechanism, the proof of Theorem 1 also makes use of a slightly larger set of representations. A control rights structure c defines a **lax (trading and braiding) mechanism** if it satisfies (C2)' instead of (C2) keeping all else equal.

(C2)' If exactly one house is brokered then there is at least one owner.

The above algorithm, together with the (lax) trading and braiding mechanism c, is used to map any fixed profile of preferences R to an outcome $c(R)$. If one strengthens (C1) to require that that at most one house is brokered, one obtains the set of Pycia and Unver [12] trading cycles mechanisms. There are also some representational differences. A Pycia and Unver [12] trading

[6] Consider a control rights structure c with three agents $\{1, 2, 3\}$ and 4 houses $\{e, f, g, h\}$, where agent 1 starts out owning house e and agent 2 starts out owning the remainder. Let R be such that 1 and 2 respectively rank e, and f at the top. The submatchings $\{(1, e)\}$, $\{(2, f)\}$ and $\{(1, e), (2, f)\}$ are reachable under c at R; $\{(2, g)\}$ is c-relevant since 2 could appropriate house g, but not it is not reachable under c at R given that 2 prefers f to g; $\{(3, h)\}$ is not c-relevant since 3 does not own h at $\emptyset$.

cycles mechanism satisfies (C2)' instead of (C2) and defines control rights functions for every submatching. The latter difference implies a notational difference between (C4), (C5) and (C6), on the one hand, and the Pycia and Unver [12] requirements that link c_{ν° and c_ν for different ν° and ν, on the other hand. The stronger (C2) together with the definition of c only on relevant submatchings allows for the uniqueness statement in the following theorem.

Theorem 2 *Any lax trading and braiding mechanism is good. Any good mechanism has a unique representation as a trading and braiding mechanism.*

Since any trading and braiding mechanism is a lax mechanism, Theorem 2 implies that trading and braiding mechanisms are good. Since any braid B^ω is a trading and braiding mechanism, Theorem 2 comprises the statement that braids are good. Theorem 2 moreover implies that any lax mechanism can be represented as a trading and braiding mechanism. The next proposition explains how to construct such alternative representations.

Proposition 1 *Fix any lax mechanism $\bar{c}$. For any $\bar{c}$-relevant ν that satisfies (C2) let $c_\nu = \bar{c}_\nu$. For any $\bar{c}$-relevant ν with $\bar{c}_\nu(h_b) = (i_b, b)$ and $\bar{c}_\nu(h) = (i^*, o)$ for some $i_b, i^* \in \overline{N}_\nu$, $h_b \in \overline{H}_\nu$ and all $h \in \overline{H}_\nu \setminus \{h_b\}$ let $c_\nu(h) = (i^*, o)$ for all $h \in \overline{H}_\nu$ and $c_{\nu \cup \{(i^*, h_b)\}}(h) = (i_b, o)$ for all $h \in \overline{H}_\nu \setminus \{h_b\}$. Then c is a trading and braiding mechanism and $c(R) = \bar{c}(R)$ holds for all $R \in \mathcal{R}$.*

The proofs of Theorem 2 and of Proposition 1 are in the online appendix. The proof of Theorem 2 starts by showing that braids are good. The remainder of Theorem 2 and of Proposition 1 are proved by induction over the number of agents n. With only one agent both obviously hold true. The next step is to show that lax mechanisms are well-defined: the order of the elimination of trading cycles does not matter. This flexibility together with the inductive hypothesis that any submechanism with fewer than n agents is good shortens the proof that any lax mechanism is good to a page. Pycia and Unver [12] broke the path for converse direction of the proof (any good mechanism can be represented as a trading and braiding mechanism). While my proof builds on the groundwork laid in Pycia and Unver [12], it ultimately deviates to show that more than one house might be brokered at some round of the mechanism and that any such round is a braid.

6 Tools for Trading and Braiding Mechanisms

A c-relevant submatching ν is c-**dictatorial** if a single agent (the **dictator at** ν) owns all houses $\overline{H}_\nu$ according to c_ν. A c-relevant submatching that is not dictatorial is c-**nondictatorial**. A (lax) trading and braiding mechanism c is a **path dependent serial dictatorship** if any c-relevant submatching is c-dictatorial. A path dependent serial dictatorship c is a serial dictatorship if the dictator at any c-relevant ν depends only on the number of agents matched under ν.

Any c-relevant submatching ν^* defines a **submechanism** $c[\nu^*]$ that maps restrictions $\overline{R}$ (of $R \in \mathcal{R}$ to $\overline{N}_{\nu^*}$ and $\overline{H}_{\nu^*}$) to submatchings $\overline{\nu}$ with the feature that $\nu^* \cup \overline{\nu}$ is a matching in the original problem. The control rights structure $c[\nu^*]$ is such that $\nu = \nu^* \cup \nu'$ is c-relevant if and only if ν' is $c[\nu^*]$-relevant. For any such pair ν, ν' we have $c[\nu^*]_{\nu'} = c_\nu$. If c is a (lax) trading and braiding mechanism then $c[\nu^*]$ also defines a (lax) trading and braiding mechanism. Fixing R such that ν^* is reachable under c at R, the definition of the trading-cycles process implies $c(R) = \nu^* \cup c[\nu^*](\overline{R})$.

Consider a c-relevant $\nu^* \neq \emptyset$ such that any c-relevant $\nu \subsetneq \nu^*$ is c-dictatorial. Then the set of all c-relevant submatchings $\nu \subsetneq \nu^*$ can be represented as $\{\nu_l\}_{l=1}^{L}$, with $\nu_1 = \emptyset$ and $\nu_{l+1} = \nu_l \cup \{(i_l, \nu^*(i_l))\}$ with i_l the dictator at ν_l for all $1 \leq l < L =\mid N_{\nu^*} \mid$. Moreover $\nu^* = \nu_L \cup \{(i_L, \nu^*(i_L))\}$. To see that such a set $\{\nu_l\}_{l=1}^{L}$ exists observe that $\nu_1 = \emptyset \subsetneq \nu^*$ is c-relevant. By the assumption that any c-relevant $\nu \subsetneq \nu^*$ is c-dictatorial $\nu_1 = \emptyset$ is c-dictatorial. If $\nu^* = \{(i_1, \nu^*(i_1))\}$ we are done. If not $\nu_2 = \{(i_1, \nu^*(i_1))\}$ is the unique direct c-successor to $\emptyset = \nu_1$ with $\nu_2 \subsetneq \nu^*$. Since $\nu_2 \subsetneq \nu^*$, ν_2 is c-dictatorial. Mutatis mutandis, the application of the preceding arguments to all consecutive ν_l establishes the claim.

A submatching ν is $(p \odot c)$-relevant if and only if $\nu \circ p$ is c-relevant. If ν is $(p \odot c)$-relevant then $(p \odot c)_\nu(h) = (p(i), o)$ holds if $c_{\nu \circ p}(h) = (i, o)$ and $(p \odot c)_\nu(h) = (p(i), b)$ holds if $c_{\nu \circ p}(h) = (i, b)$. So if agent i controls house h at $\nu \circ p$ under c, then agent $p(i)$ controls h at ν under $p \odot c$; the type of control stays the same.

Consider two mechanisms c, c' that only differ on c-relevant submatchings ν with $\nu^* \subset \nu$. For any fixed ν with $\nu^* \not\subset \nu$, we have $c_\nu = c'_\nu$ and ν is c-relevant if and only if it is c'-relevant. If $\nu^* \not\subset c(R)$ then c and c' prescribe the same control rights function for any reachable ν under c at R and we obtain $c(R) = c'(R)$. In the context of permuted mechanisms $(p \odot c')(R) = (p \odot c)(R)$ holds for any p for which $\nu^* \circ p^{-1} \not\subset (p \odot c)(R)$.

7 The proof

Fix any trading and braiding mechanism c^0. In Section 7.1 I construct a sequence from c^0 to a serial dictatorship c^K. Step α determines whether Step β, γ or δ should be used to transform c^k into c^{k+1}: β consolidates ownership, γ replaces a braid with a serial dictatorship, and δ reorders dictators in a path dependent serial dictatorship. In Section 7.2 I show that the sequence is well-defined and terminates indeed with a serial dictatorship. In Section 7.3 I show that any two mechanisms c^k, c^{k+1} are s-equivalent.

7.1 Construction of a sequence $c^0, c^1, \cdots, c^K$

Go to Step $(\alpha, 0)$.

Step (α, k): If c^k is a serial dictatorship end with $k = K$. If not, go to Step (δ, k) if c^k is a path dependent serial dictatorship. If neither case applies fix a minimal c^k-nondictatorial submatching ν^*. If at most one house is brokered at ν^* under c^k go to Step (β, k) if not go Step (γ, k).

Step (β, k): Let c^{k+1} be the trading and braiding mechanism that represents the lax mechanism $\overline{c}$ which is defined as follows. For any c^k-relevant ν with $\nu^* \not\subset \nu$ let $c^k_\nu = \overline{c}_\nu$. Assume w.l.o.g. that 1 and 2 own houses under $c^k_{\nu^*}$. Say that a c^k-relevant ν is of type 0 if $1, 2 \notin N_\nu$. For any type 0 ν let $\overline{c}_\nu(h) = (1, o)$ if $c^k_\nu(h) = (2, o)$ and $\overline{c}_\nu(h) = c^k_\nu(h)$ for all other $h \in \overline{H}_\nu$. Let ν with $1 \in N_\nu$ be a direct $\overline{c}$-successor to a type 0 submatching ν°. If ν is c^k-relevant then ν is of type 1. If not then ν is of type 2. Any direct $\overline{c}$-successor of a type $t \neq 0$ submatching is of type t. For any type 1 submatching ν let

$c^k_\nu = \bar{c}_\nu$, for any type 2 submatching let $\bar{c}_\nu(h) = (2, o)$ if $c^k_{\nu\circ(1,2)}(h) = (1, o)$ and $\bar{c}_\nu(h) = c^k_{\nu\circ(1,2)}(h)$ otherwise. Go to Step $(\alpha, k+1)$.

Step (γ, k): Let $c^k_\nu = c^{k+1}_\nu$ for any c^k-relevant $\nu \neq \nu^*$. Let $c^{k+1}[\nu^*]$ be a serial dictatorship. Go to Step $(\alpha, k+1)$.

Step (δ, k): Fix three c^k-relevant submatchings $\nu^* \subset \nu' \subset \nu''$, such that $c^k[\nu^*]$ is a serial dictatorship with i the dictator at ν' and $j < i$ the dictator at ν''. Say $i = 2$ and $j = 1$. Let $c^k_\nu = c^{k+1}_\nu$ for any c^k-relevant ν with $\nu^* \not\subset \nu$ and $c^{k+1}[\nu^*] = (1,2) \odot c^k[\nu^*]$. Go to Step $(\alpha, k+1)$.

7.2 All transformations from c^k to c^{k+1} are welldefined

Claim 1: If c^{k+1} is constructed via β (γ or δ) then it is a trading and braiding mechanism if $\bar{c}[\nu^*]$ is a lax mechanism (if $c^{k+1}[\nu^*]$ is a trading and braiding mechanism).

If $\nu^* = \emptyset$ the claim trivially holds; so assume that $\nu^* \neq \emptyset$. First consider the case that β is used to construct c^{k+1} and that $\bar{c}[\nu^*]$ is a lax mechanism. Thanks to Theorem 2 it is sufficient to show that $\bar{c}$ is a lax mechanism. For any c^k-relevant ν with $\nu^* \not\subset \nu$ $\bar{c}_\nu = c^k_\nu$ satisfies (C1), (C2), and (C3) as c^k is a trading and braiding mechanism. By the same reason (C4), (C5), and (C6) are satisfied at ν°, ν if ν a direct $\bar{c}$-successor to ν° and $\nu^* \not\subset \nu$. Any ν with $\nu^* \not\subset \nu$ is $\bar{c}$-relevant if and only if it is c^k relevant. So $\bar{c}_\nu$ is for any $\bar{c}$-relevant ν with $\nu \not\subset \nu^*$ uniquely defined as c^k_ν.

Since ν^* is a minimal non-dictatorial submatching all c^k-relevant $\nu \subsetneq \nu^*$ are c^k-dictatorial and can (by the arguments in Section 6) be represented as $\{\nu_l\}_{l=1}^L$ with $\nu_1 = \emptyset$, $\nu_{l+1} = \nu_l \cup \{(i_l, \nu^*(i_l))\}$ for all $1 \leq l < L =\mid N_{\nu^*} \mid$, and $\nu_L \cup \{(i_L, \nu^*(i_L))\} = \nu^*$. There is exactly one pair of a $\bar{c}$-relevant ν° and a direct $\bar{c}$-successor ν with $\nu^\circ \subsetneq \nu^* \subset \nu$: $\nu^\circ = \nu_L$ and $\nu = \nu^*$. Since ν_L is $\bar{c}$-dictatorial (C4), (C5), and (C6) do not impose any restrictions on $\bar{c}_{\nu^*}$.

For any $\bar{c}$-relevant $\nu \supset \nu^*$ $\bar{c}_\nu$ is defined as $\bar{c}[\nu^*]_{\nu'}$ where $\nu = \nu^* \cup \nu'$. By the above arguments there is exactly one path of cycle removal to reach ν^* under $\bar{c}$. So $\nu \supset \nu^*$ is $\bar{c}$-relevant if and only if $\nu \setminus \nu^*$ is $\bar{c}[\nu^*]$-relevant, implying that for any $\bar{c}$-relevant $\nu \supset \nu^*$ there exists a $\bar{c}[\nu^*]$-relevant ν' such that $\nu = \nu^* \cup \nu'$. The uniqueness of the path of cycle removal implies that

ν cannot be reached on any other path, and $\overline{c}_\nu = c[\nu^*]_{\nu\setminus\nu^*}$ is well-defined for any $\overline{c}$-relevant ν. Moreover, since $\overline{c}[\nu^*]$ is - by assumption - a lax mechanism (C1), (C2)', (C3) as well as (C4), (C5), and (C6) hold for the submatchings in $\overline{c}[\nu^*]$ and $\overline{c}$ is a lax mechanism.

The above arguments have to be modified slightly to show that c^{k+1} is a trading and braiding mechanism if $c^{k+1}[\nu^*]$ is a trading and braiding mechanism when Step γ or δ is used to construct c^{k+1}. If γ is used then any c^k-relevant $\nu \subsetneq \nu^*$ is c^k-dictatorial by the same reason as above. If δ is used the same conclusion holds as c^k is a path dependent serial dictatorship. Steps γ and δ directly construct a trading and braiding mechanism c^{k+1} and $\overline{c}$ has to be replaced by c^{k+1} wherever it appears.

Claim 2: $\overline{c}[\nu^*]$, constructed via Step β is a lax mechanism; $c^{k+1}[\nu^*]$, constructed via Step γ or δ is a trading and braiding mechanism.

If c^{k+1} is constructed via Step γ or δ then $c^{k+1}[\nu^*]$ is a serial dictatorship and we are done. So assume that $\overline{c}$ is constructed via Step β. Fix any type 0 ν° together with a direct $\overline{c}$-successor ν. Under $\overline{c}_{\nu^\circ}$ there is one less owner than under $c^k_{\nu^\circ}$. Since c^k satisfies (C2) there is at least one owner under $\overline{c}_{\nu^\circ}$, as required by (C2)'. Since $\overline{c}_{\nu^\circ}(h) = (i, b)$ only holds if $c^k_{\nu^\circ}(h) = (i, b)$ and since c^k satisfies (C1) and (C3), $\overline{c}_{\nu^\circ}$ satisfies (C1) and (C3).

If ν is of type 0 then ν° and ν are both c^k-relevant. If $\overline{c}_{\nu^\circ}(h) = (1, o)$ then $c^k_{\nu^\circ}(h) = (i, o) = c^k_\nu(h)$ holds for $i \in \{1, 2\}$ as c^k satisfies (C4). The definition of $\overline{c}$ then implies $\overline{c}_\nu(h) = (1, o)$. So 1's ownership continues from ν° to ν under $\overline{c}$ as well as under c^k. On the other hand, 2, who is an owner at ν^* under c^k, is by (C4) an owner at ν° and at ν under c^k. Since at least two agents (1 and 2) own houses at ν° and ν under c^k which satisfies (C4) and (C5), $\overline{c}_{\nu^\circ}(h) = c^k_{\nu^\circ}(h) = c^k_\nu(h) = \overline{c}_\nu(h)$ holds for any h with $\overline{c}_{\nu^\circ}(h) = (j, \cdot)$ for some $1 \neq j \in \overline{N}_\nu$. So (C5) and (C6) are satisfied by $\overline{c}$ at ν°, ν as any agent $i_b \in N_\nu$ who brokers at ν° under $\overline{c}$ continues to do so at ν. Since the ownership of of agent 1 continues from ν° to ν under $\overline{c}$ (as shown above) $\overline{c}$ also satisfies (C4) at ν° and ν.

If ν is of type 1, let $\overline{c}_{\nu^\circ}(h) = (i_b, b)$, $i_b \notin N_\nu$ and $\overline{c}_\nu(h) \neq (i_b, b)$, implying $c^k_{\nu^\circ}(h) = (i_b, b)$ and $c^k_\nu(h) \neq (i_b, b)$. Since 2 is neither matched at ν° nor at ν and since 2 is an owner under c^k at ν°, 2 must by (C4) be an owner under c^k at ν. Since c^k satisfies (C5) no other agent is an owner under c^k at both these matchings. The derivation of $\overline{c}$ from c^k implies that no agent is an owner at ν° and ν under $\overline{c}$. So $\overline{c}$ satisfies (C5) and (C6) at ν°, ν. If $\overline{c}_{\nu^\circ}(h) = (i, o)$ and $i \notin N_\nu$ then $i \neq 1, 2$ as 1 is matched under ν (which is of type 1) and as 2 is not an owner under $\overline{c}$ at ν° (which is of type 0). The definition of $\overline{c}$ and the fact that c^k satisfies (C4) then imply $(i, o) = \overline{c}_{\nu^\circ}(h) = c^k_{\nu^\circ}(h) = c^k_\nu(h) = \overline{c}_\nu(h)$. If ν is of type 2 agents 1 and 2 have to be switched in all of the arguments that refer to c^k (including the replacement of ν by $\nu \circ (1, 2)$ whenever ν is of type 2) to show that $\overline{c}$ satisfies (C4), (C5) and (C6) at ν°, ν.

If ν is of type 1 or 2 then $\overline{c}[\nu]$ respectively equals $c^k[\nu]$ and $(1, 2) \odot c^k(\nu \circ (1, 2))$. In either case $\overline{c}[\nu]$ is a trading and braiding mechanism and (C1)-(C6) are satisfied in all remaining cases. Since any $\overline{c}$-relevant ν is of type 0, 1, or 2, and since no submatching ν is of two different types, Step β uniquely specifies $\overline{c}_\nu$ for every $\overline{c}$-relevant ν. In sum $\overline{c}[\nu^*]$ is a lax mechanism.[7]

In either case $\overline{c}[\nu]$ is a trading and braiding mechanism and (C1)-(C6) are satisfied in all remaining cases. Since any $\overline{c}$-relevant ν is of type 0, 1, or 2, and since no submatching ν is of two different types, Step β uniquely specifies $\overline{c}_\nu$ for every $\overline{c}$-relevant ν. In sum $\overline{c}[\nu^*]$ is a lax mechanism.

Claim 3: The sequence $c^1, \ldots, c^K$ ends with a serial dictatorship c^K

Fix a c^{k+1}-relevant ν such that neither ν nor $\nu \circ (1, 2)$ is c^k-relevant. For

[7]To see that no submatching is of type 1 as well as of type 2, suppose some ν was of both types, implying that ν is reached on two different paths for two different profiles of preferences. Say that under $\overline{c}$ with R° ν is reached via a cycle in which 1 exchanges a house that he owns under c^k (type 1), whereas with R^* ν is reached via a cycle in which 1 exchanges a house that 2 owns under c^k (type 2). Now define a profile of preferences R with $\nu(i) R_i H$ for all $i \in N_\nu$. Since trading cycles may be eliminated in any which order (as is shown in Section 4 of the online appendix), ν is reachable under $\overline{c}$ with R via the paths under which it is reached with R^* and R°. Let ν' be the first submatching at which 1 is part of a cycle under $\overline{c}$ with R. Section 4 of the online appendix shows that for a trading cycles not to persist from round to round, a broker must lose control over the house he brokers. This case is not relevant here: to reach ν first a sequence of dictators has to choose in accordance with $\nu^* \subset \nu$. At ν^* under c^k 1 and 2 are both owners. Since they remain owners under c^k until ν is reached no broker may (by C5) loose control under c^k until ν is reached. The derivation of $\overline{c}$ from c^k implies that the same holds for $\overline{c}$. In sum, agent 1 is matched via the cycle that first forms at ν' under $\overline{c}$ with either R, R°, or R^*. A contradiction ensues since this cycle involves a house owned by either 1 or 2 under c^k and the submatching ν is either of type 1 or type 2.

such a ν to exist c^{k+1} must have been derived via β or γ (if c^{k+1} is constructed via a reordering of dictators as prescribed by δ then ν is c^{k+1}-relevant if and only if either ν or $\nu \circ (1,2)$ is c^k-relevant). If γ was used, then c^{k+1} arises out of c^k via the replacement of a braid with a serial dictatorship and ν is c^{k+1}-dictatorial. If β was used then c^{k+1} is derived from $\overline{c}$ via Proposition 1 and ν equals $\nu' \cup \{(i, h_b)\}$ for a house h_b that is brokered at ν' under c^k but owned by i at ν' under c^{k+1}. Such a ν is c^{k+1}-dictatorial. So the number of c^k-nondictatorial submatchings does not increase in k. If c^k-nondictatorial submatching exists, α prescribes to follow β or γ. Since β reduces the number of owners at at least one c^k-relevant submatching and since γ replaces a braid with a serial dictatorship, the process of transformations eventually reaches a trading and braiding mechanism $c^{k'}$ such that any $c^{k'}$-relevant submatching is $c^{k'}$-dictatorial. But such a $c^{k'}$ is a path dependent serial dictatorship. Finally Step δ is iteratively reorders agents as dictators such that an earlier dictator i swaps roles with a later dictator j if $i > j$. Such reordering occurs until a serial dictatorship c^K is attained.

7.3 $\Delta c^k = \Delta c^{k+1}$ holds for all k

The claim holds if for any R there exists a one-to-one $f : P \to P$ with $(p \odot c^k)(R) = (f(p) \odot c^{k+1})(R)$ for all $p \in P$. Fix an arbitrary R, let P^0 be the set of all $p \in P$ with $\nu^* \circ p^{-1} \not\subset (p \odot c^k)(R)$. For any $p \in P^0$ let $f(p) \colon = p$. Since $c^k_\nu = c^{k+1}_\nu$ holds for any c^k- and c^{k+1}-relevant ν with $\nu^* \not\subset \nu$, $(p \odot c^k)(R) = (f(p) \odot c^{k+1})(R)$ holds for any $p \in P^0$. Since the restriction of f to P^0 is one-to-one and since $f(P^0) = P^0$, f is one-to-one if also its restriction to $\overline{P} \colon = P \setminus P^0$ is one-to-one and if $f(\overline{P}) \subset \overline{P}$.

To see that $f(p)$ is an element of $\overline{P}$ holds if $p(i) = f(p)(i)$ holds for all $i \in N_{\nu^*}$ and if $p \in \overline{P}$, note that, no matter whether β, γ, or δ was used to define c^{k+1}, $\nu \subsetneq \nu^*$ is c^k-relevant if and only if it is c^{k+1}-relevant. By the arguments in Section 6 the set of all $\nu \subsetneq \nu^*$ can be represented as $\{\nu_l\}_{l=1}^L$ with $\nu_1 = \emptyset$, $\nu_{l+1} = \nu_l \cup \{(i_l, \nu^*(i_l))\}$ for all $1 \leq l < L =\mid N_{\nu^*} \mid$ and i_l the dictator at ν_l. Given $p(i_1) = f(p)(i_1)$, the role of i_1, the dictator at $\emptyset$ under c^k and c^{k+1} is assumed by $p(i_1)$ under $p \odot c^k$ and by $f(p)(i_1) = p(i_1)$ under $f(p) \odot c^{k+1}$. The trading cycles processes of $p \odot c^k$ and $f(p) \odot c^{k+1}$ at R start out identically: $p(i_1)$ appropriates house $\nu^*(i_1)$. Since $\nu^* \circ p^{-1} \subset (p \odot c^k)(R)$ (as is required for $p \in \overline{P}$), we can proceed inductively to obtain $\nu^* \circ p^{-1} \subset (f(p) \odot c^{k+1})(R)$ and $f(\overline{P}) \subset \overline{P}$. To establish $\Delta c^k = \Delta c^{k+1}$ it is in sum sufficient to construct a one-to-one $f : \overline{P} \to P$ with $f(p)(i) = p(i)$ for all $i \in N_{\nu^*}$ and $(p \odot c^k)(R) = (f(p) \odot c^{k+1})(R)$ for all $p \in \overline{P}$.

7.3.1 When Step (β, k) is used to construct c^{k+1}

Fix any $p \in \overline{P}$ and let $\overline{\nu}$ be the maximal reachable submatching under $p \odot c^k$ at R with $p(1), p(2) \notin N_{\overline{\nu}}$.[8] If $p(1)$ is part of a cycle under $p \circ c^k$ at $\overline{\nu}$ let $p \in \overline{P}^0$ and $f(p) \colon = p$ if not let $p \in \overline{\overline{P}}$ and $f(p) \colon = p \circ (1, 2)$. Since 1 and 2 are not matched at ν^*, $f(p(i))$ equals $p(i)$ for all $i \in N_{\nu^*}$. Since $\overline{c}$ and c^{k+1} represent the same mechanism $(p \odot c^k)(R) = (f(p) \odot c^{k+1})(R)$ holds if and only if $(p \odot c^k)(R) = (f(p) \odot \overline{c})(R)$. In the next two paragraphs I show that $(p \odot c^k)(R) = (f(p) \odot \overline{c})(R)$ holds for any $p \in \overline{P}$.

First, fix some $p \in \overline{P}^0$ and say that ν is the (unique) direct $(p \odot c^k)$-successor of $\overline{\nu}$ with $p(1) \in N_\nu$. If $p(2) \in N_\nu$ represent the unique cycle at $\overline{\nu}$ under $p \odot c^k$ and R as $p(1) \to \nu(1) \to \cdots \to h_2 \to p(2) \to \nu(2) \to \cdots \to h_1 \to p(1)$. Since $p(1)$ owns h_2 under $p \odot \overline{c}$ at $\overline{\nu}$, the cycle $p(1) \to \nu(1) \to \cdots \to h_2 \to p(1)$ forms at $\overline{\nu}$ under $p \odot \overline{c}$ and R. Since $p(1)$ trades h_2 in this cycle and since $(p \odot c^k)_{\overline{\nu}}(h_2) = (p(2), o)$, agent $p(2)$ inherits h_1 under $p \odot \overline{c}$ once $p(1)$ is matched. The cycle $p(2) \to \nu(2) \to \cdots \to h_1 \to p(2)$ forms next under $p \odot \overline{c}$ at R. So ν is reachable under $p \odot \overline{c}$ at R. If $p(2) \notin N_\nu$ then $p(1)$ is part of the same cycle at $\overline{\nu}$ under $p \odot c^k$ and at $\overline{\nu}$ under $p \odot \overline{c}$

[8] To see that for any lax mechanism c, R and $i, j \in N$ there exists a unique maximal reachable submatching ν that leaves i and j unmatched, let $\hat{\nu}$ and $\tilde{\nu}$ be any two reachable submatchings that leave i and j unmatched. Let $\hat{\nu} = \bigcup_{l=0}^{\hat{L}} \hat{\nu}_l$ and $\tilde{\nu} = \bigcup_{l=0}^{\tilde{L}} \tilde{\nu}_l$ where $\hat{\nu}_0 = \tilde{\nu}_0 = \emptyset$ and $\hat{\nu}_l$ and $\tilde{\nu}_l$ arise out of matching a single cycle in round $l \geq 1$. Such sets of submatchings exist, since cycles may be removed in any order (Section 4 of the online appendix). Since $\nu_0 = \emptyset$, $\hat{\nu} \cup \tilde{\nu}_0$ is reachable under c at R. Let $\hat{\nu} \cup (\bigcup_{l=0}^{L} \tilde{\nu}_l) \neq \hat{\nu} \cup \tilde{\nu}$ be reachable for some $L < \tilde{L}$. If $\hat{\nu} \cup (\bigcup_{l=0}^{L} \tilde{\nu}_l) = \hat{\nu} \cup (\bigcup_{l=0}^{L+1} \tilde{\nu}_l)$, then $\hat{\nu} \cup (\bigcup_{l=0}^{L+1} \tilde{\nu}_l)$ is obviously reachable under c at R. If not then $\hat{\nu} \cup (\bigcup_{l=0}^{L+1} \tilde{\nu}_l)$ is reachable under c at R since the cycle yielding $\tilde{\nu}_{L+1}$ forms at $\bigcup_{l=0}^{L} \tilde{\nu}_l$ and may be removed next (if this cycle does not persist, then a broker looses control in the set and two consecutive cycles are needed to form the submatching $\tilde{\nu}_{L+1}$). By induction $\hat{\nu} \cup \tilde{\nu}$ is reachable under c at R. The unique maximal reachable submatching ν with $i, j \notin N_\nu$ is the union of all submatchings ν' with $i, j \notin N_{\nu'}$.

and ν is reachable under $p \odot \overline{c}$ and R. In either case the definition of $\overline{c}$ implies that the submechanisms of $p \odot c^k$ and $p \odot \overline{c}$ following ν are identical: $(p \odot c^k)[\nu] = (p \odot \overline{c})[\nu]$ and therefore $(p \odot c^k)(R) = (p \odot \overline{c})(R) = (f(p) \odot \overline{c})(R)$.

Now fix some $p \in \overline{\overline{P}}$: $p(1)$ does not take part in any cycle at $\overline{\nu}$ under $p \odot \overline{c}$ and R. Instead $p(2)$ takes part in the unique such cycle $p(2) \to \nu(2) \to \cdots \to h_2 \to p(2)$.[9] Let ν be the direct c^k-successor of $\overline{\nu}$ that arises out of matching this cycle. Under $f(p) \odot \overline{c}$ at $\overline{\nu}$ agent $p(2)$ owns all houses owned by agents $p(1)$ and $p(2)$ under $p \odot c^k$ at $\overline{\nu}$ and $p(2) \to \nu(2) \to \cdots \to h_2 \to p(2)$ also forms under $f(p) \odot \overline{c}$ at $\overline{\nu}$. Once this cycle is matched $f(p) \odot \overline{c}$ continues as if $p(1)$ had always played the role of 1 in c^k: Since $(1,2)$ in the definition of $\overline{c}$ is inverted by $(1,2)$ in the transformation of p under f, we have $(f(p) \odot \overline{c})[\nu] = (p \odot c^k)[\nu]$ and $(p \odot c^k)(R) = (f(p) \odot \overline{c})(R)$.

Restricted to $\overline{P}^0$ and to $\overline{\overline{P}}$ f is one-to-one. Since $f(\overline{P}^0) = \overline{P}^0$, f is one-to-one if $f(\overline{\overline{P}}) \subset \overline{\overline{P}}$. Fix any $p \in \overline{\overline{P}}$ and let $p(2) \to \nu(2) \to \cdots \to h_2 \to p(2)$ be the unique cycle at $\overline{\nu}$ under $p \odot c^k$ and R. The uniqueness of this cycle and $p \in \overline{\overline{P}}$ imply that $p(1)$ is part of some chain $p(1) \to h^* \to \cdots \to h'$ under $p \odot c^k$ and R at $\overline{\nu}$. This chain terminates with a house h' in the cycle $p(2) \to \nu(2) \to \cdots \to h_2 \to p(2)$. Under $f(p) \odot c^k$ agent $p(1)$ owns house h_2 at $\overline{\nu}$ so the cycle $p(1) \to h^* \to \cdots \to h' \to i' \to \cdots \to h_2 \to p(1)$ forms at $\overline{\nu}$ under $f(p) \odot c^k$ and R. In that case $p(2)$ is part of the chain $p(2) \to \nu(2) \to \cdots \to h''$ that terminates with house h'' in the cycle $p(1) \to h^* \to \cdots \to h' \to i' \to \cdots \to h_2 \to p(1)$. There cannot be another cycle at $\overline{\nu}$ under $f(p) \odot c^k$ and R since this cycle would only involve agents $i \neq p(1), p(2)$ and would therefore also form under $p \odot c^k$ at R - contradicting the maximality of $\overline{\nu}$ in the set of all reachable submatching under $p \odot c^k$ at R which involve neither $p(1)$ nor $p(2)$. So there is only one cycle at $\overline{\nu}$ under $f(p) \odot c^k$ and R and this cycle involves $f(p(2)) = p(1)$ implying that $f(p) \in \overline{\overline{P}}$ as required.

[9] The existence and uniqueness of this cycle follows from $\overline{\nu}$ being maximal in the set of submatchings ν with $p(1), p(2) \notin N_\nu$ that are reachable under $p \odot c^k$ at R.

7.3.2 When the transformation is constructed using Step (γ, k)

The construction of f on $\overline{P}$ relies on Proposition 2 on random matching mechanisms with just three agents and three houses. Some more concepts are needed. For any random matching mechanism $\mathfrak{M} : \mathcal{R} \to \Delta\mathcal{M}$ let $\mathfrak{M}(R)[i, h]$ be the probability that agent i is matched with house h when the agents announce the profile of preferences R. A random matching mechanism $\mathfrak{M} : \mathcal{R} \to \Delta\mathcal{M}$ is **ex post Pareto optimal** if any μ in the support of $\mathfrak{M}(R)$ is Pareto optimal at R. The mechanism $\mathfrak{M}$ is **ordinally strategy** proof if $\sum_{hR_ih^*} \mathfrak{M}(R)[i, h] \geq \sum_{hR_ih^*} \mathfrak{M}(R'_i, R_{-i})[i, h]$ holds for all R, R'_i, i, h^*. So under an ordinally strategy proof mechanism no agent can misrepresent his preferences to increase his probability to get a house he prefers to some fixed h^*. A mechanism $\mathfrak{M}$ satisfies **equal treatment of equals** if any two agents i and j who announce the same preferences face the same distribution over matches, $R_i = R_j \Rightarrow \mathfrak{M}(R)[i, h] = \mathfrak{M}(R)[j, h]$ for all $h \in H$. The symmetrization of any good mechanism is ex post Pareto optimal, ordinally strategy proof and satisfies equal treatment of equals.[10]

Proposition 2 *Let $H = \{a, b, c\}$ and $N = \{1, 2, 3\}$. Let $\mathfrak{M} : \mathcal{R} \to \Delta\mathcal{M}$ be ex post Pareto optimal, ordinally strategy proof and satisfy equal treatment of equals. Then $\mathfrak{M}$ is a random serial dictatorship.*

The proof of Proposition 2 is in the online appendix. There I fix an arbitrary $\mathfrak{M}$. For each R I derive 9 linearly independent linear equations from ex post Pareto optimality, strategy proofness and equal treatment of equals to uniquely determine the 9 values $\mathfrak{M}(R)[i, h]$. Since random serial dictatorship satisfies the three properties it must equal $\mathfrak{M}$.

To see that $\Delta c^k = \Delta c^{k+1}$ holds when γ is used to derive c^{k+1}, assume w.l.o.g that $\{1, 2, 3\} = N_\omega$ where ω is the avoidance matching that defines the braid $B^\omega = c^k[\nu^*]$ and that the three dictators in the serial dictatorship $c^{k+1}[\nu^*]$ are 1, 2 and 3. Partition $\overline{P}$ into sets $P_1, \ldots, P_M$ such that p, q belong to the same P_m if and only if $p(i) = q(i)$ for all $i \in N_{\nu^*}$, so all permutations in some subset P_m map the same agents to the roles in N_{ν^*}. For any m let

[10] A random matching mechanism is non-bossy if no agent can alter someone else's lottery over matches without altering his own. For a proof that non-bossiness is robust to randomization when the base mechanism is strategy proof, see Bade [3].

$\mathcal{I}_m \colon = \{I \subset \overline{N}_{\nu^* \circ p^{-1}} :\mid I \mid = 3, p \in P_m\}$ be the set of sets consisting of three unmatched agents under c^k at $\nu^* \circ p^{-1}$ for any $p \in P_m$. Since $c^k_\nu = c^{k+1}_\nu$ holds for all c^k-relevant $\nu \neq \nu^*$, we would obtain the same set $\mathcal{I}_m$ if we were to use c^{k+1} to define it. For any $m \in \{1, \ldots, M\}$ and $I \in \mathcal{I}_m$ let $P^I_m \subset P_m$ be such that $p(\{1,2,3\}) = q(\{1,2,3\}) = I$ for all $p, q \in P^I_m$. By Proposition 2 and the definition of P^I_m there exists a bijection $f^I_m : P^I_m \to P^I_m$ with $(p \odot c^k)(R) = (f(p) \odot c^{k+1})(R)$ for all $p \in P^I_m$. Since $\{P^I_m\}_{I \in \mathcal{I}_m, 1 \leq m \leq M}$ partitions $\overline{P}$, the function $f : \overline{P} \to \overline{P}$ defined by $f(p) \colon = f^I_m(p)$ for any $p \in P^I_m$ is one-to-one with $(p \odot c^k)(R) = (f(p) \odot c^{k+1})(R)$ for all $p \in \overline{P}$. Moreover, $p(i) = f(p)(i)$ holds for all $i \in N_{\nu^*}$ as $f(P_m)$ equals P_m for any $m \in \{1, \ldots, M\}$.

7.3.3 When the transformation is constructed using Step (δ, k)

Define $f(p) = p \circ (1, 2)$ for any $p \in \overline{P}$ and note that $(f(p) \odot c^{k+1})[\nu^*] = (p \odot c^k)[\nu^*]$ holds for any $p \in \overline{P}$, since $(1, 2)$ in the definition of c^{k+1} is inverted by $(1, 2)$ in the transformation of p with f. Moreover restricted to $\overline{P}$ f is one-to-one and $p(i) = f(p)(i)$ holds for all $i \in N_{\nu^*}$ since $1, 2 \notin N_{\nu^*}$.

8 Conclusion

Two approaches had so far been used to establish the equivalence between symmetrizations of different good mechanisms: Abdulkadiroglu and Sönmez [1] as well as Knuth [8] constructed bijections to show the s-equivalence of GTTC and serial dictatorship. Carroll [5] constructed more complex bijections to show the s-equivalence of serial dictatorship and any top trading cycles mechanism. Pathak and Sethuraman [11] and Lee and Sethuraman [9] used an inductive strategy over the number of agents in a mechanism to prove that any hierarchical exchange mechanism following Papai [10] with equally many houses and agents is s-equivalent to serial dictatorship.

Could one use either one of these strategies to extend the s-equivalence result to differently many agents and houses? The case with more agents than houses is easily accommodated: Fix a mechanism M with $\mid N \mid$ agents and houses that is s-equivalent to serial dictatorship and create $\mid N \mid - \mid H \mid$

dummy houses. For any profile of preferences R on the original set of houses H define an auxiliary profile of preferences R' such that R is the restriction of R' to the original set of houses H and such that any agent ranks all houses in H above all dummy-houses. Derive $\Delta M(R) = \Delta S(R)$ by equating the probability that an agent obtains a dummy house under $\Delta M(R')$ or $\Delta S(R')$ with the probability that the agent obtains no house under $\Delta M(R)$ and $\Delta S(R)$.

To apply the same trick when there are more houses than agents, the dummy agents would have to be endowed with "dummy preferences". However, when a dummy agent is matched with a house that some real agent prefers to his match, the exclusion of the dummy-house match leads to a Pareto inferior matching and the trick does not work. Carroll's [5] and Lee and Sethuraman's [9] results on partial symmetrizations that treat agents in some sets symmetrically while maintaining their relative place with respect to other sets of agents could be used to cover the case of more houses than agents.

Could we use one of the existing strategies of proof for the case of a good mechanism that is not a hierarchical exchange mechanism? The task of directly constructing a bijection to prove Theorem 1 seems out of the question. Carroll's work [5] probably hits the limit in this dimension. The inductive strategy or Pathak and Sethuraman [11] and Lee and Sethuraman [9] relies on mechanisms being representable as trading mechanisms in which each agent points to their most preferred house. Given that brokers may not do so and given that braids are not representable as such trading mechanisms, this strategy of proof does not extend to the grand set of mechanisms.

Since my strategy of proof relies on sequences of marginally different good mechanisms, it can only be used on a sufficiently rich set of mechanisms. Historically speaking, it would have been difficult to apply the present strategy of proof in 1996 or 1998, when Knuth [8] and Abdulkadiroglu and Sönmez [1] respectively showed the s-equivalence of GTTC and serial dictatorship, given that Papai's [10] hierarchial exchange mechanisms came out in 2000.

The inductive consolidation of ownership works best with hierarchical exchange mechanisms. To see this reconsider this paper without brokers or braids. All that remains in Section 5 on the characterization of mechanisms, are control rights structures c that map any c-relevant submatching to a control rights function in which all houses are owned. In this setup (C1), (C2), (C3), (C5) and (C6) are trivially satisfied, only (C4) matters. Without brokers the definition of lax mechanisms is obsolete. Given that there are no braids Step α assigns any c^k to one of *two* possible transformations: c^{k+1} is either constructed through the consolidation of ownership in Step β or through a reordering of dictators in Step δ. To see that the sequence is well-defined we only need to check that (C4) remains valid in the transformations. The case of hierarchical exchange mechanisms could be dealt with in fewer than half the pages necessary to cover the grand set of good mechanisms. Differently from the predecessors in the literature this proof that any hierarchical exchange mechanism is s-equivalent to serial dictatorship, does not involve any combinatorial arguments.

To deal with braids I showed Proposition 2: random serial dictatorship is the unique ex post Pareto optimal and ordinally strategy proof random matching mechanism for three agents and three houses that satisfies equal treatment of equals. This result yields a more general conjecture than Theorem 1. Could random serial dictatorship be the unique ex post Pareto optimal and ordinally strategy proof random matching mechanism that satisfies equal treatment of equals? Unfortunately, the method used in my proof of Proposition 2 becomes cumbersome with more houses and agents. For each $R \in \mathcal{R}$, the probability distribution $\mathfrak{M}(R)$ would have to be identified via $\mid H \mid \times \mid N \mid$ linearly independent linear equations on the $\mid H \mid \times \mid N \mid$ probabilities $\mathfrak{M}(R)[i, h]$. If i is not matched to h for any Pareto optimum at R then we for example obtain the (linear) equation $\mathfrak{M}(R)[i, h] = 0$. The problem is that Saban and Sethuraman [14] have shown that finding all house-agent pairs that do not form in any Pareto optimum at some R is an NP-complete problem. So a proof of the more general conjecture requires a different attack.

While this question remains open, Erdil [7] provides an interesting contrast with the case of house matching problems in which agents can opt to stay unmatched. He shows that in such problems the conjecture does not hold. Instead random serial dictatorship is ex-ante Pareto dominated by other strategyproof, non-bossy and fair mechanisms.

Some papers, such as Bogomolnaia and Moulin [4] have presented possible tradeoffs between Pareto optimality and strategy proofness while maintaining equal treatment of equals and non-bossiness. In this context, random serial dictatorship is typically used as the benchmark of a mechanism that is best in terms of its incentive properties (ordinally strategy proof) and worst in terms of its welfare properties (only ex post Pareto optimal). This paper strengthens the case for using random serial dictatorship as the benchmark. While initially one could have criticized the choice of a particular good mechanism as the base of the symmetrization, I have shown that this choice does not matter: the symmtrization of any good mechanism leads to random serial dictatorship.

References

[1] Abdulkadiroglu , A. and T. Sönmez: "Random Serial Dictatorship and the Core from Random Endowments in House Allocation Problems", *Econometrica*, 66, (1998), pp. 689 - 701.

[2] Bade, S. "Pareto optimal, strategy proof and non-bossy matching mechanisms" *mimeo*, Royal Holloway, 2014.

[3] Bade, S. "Fairness and group-strategyproofness clash" *mimeo*, Royal Holloway, 2015.

[4] Bogomolnaia, A. and H. Moulin: "A New Solution to the Random Assignment Problem," *Journal of Economic Theory*, 100, (2001), 295-328.

[5] Carroll, G.: "A General Equivalence Theorem for Allocation of Indivisible Objects," *Journal of Mathematical Economics*, 51, (2014), 163-177.

[6] Ekici, O.: "A General Fair and Efficient Discrete Resource Allocation: A Market Approach," *mimeo* Ozyegin University, 2012.

[7] Erdil, A.: "Strategy-proof Stochastic Assignment," *Journal of Economic Theory*, 151 (2014) 146-162.

[8] Knuth, D.: "An exact analysis of stable allocation", *Journal of Algorithms*, 20, (1996), pp. 431-442.

[9] Lee, T. and J. Sethuraman: "Equivalence results in the allocation of indivisible objects: A unified view", *mimeo*, Columbia University, 2011.

[10] Papai, S.: "Strategyproof Assignment by Hierarchical Exchange", *Econometrica*, 68, (2000), pp. 1403 - 1433.

[11] Pathak, P and J. Sethuraman.: "Lotteries in student assignment: An equivalence result", *Theoretical Economics*, 6, (2011), pp. 1-17.

[12] Pycia, M. and U. Ünver: "Incentive Compatible Allocation and Exchange of Discrete Resources," *mimeo* UCLA, February 2014.

[13] Pycia, M. and U. Ünver: "Incentive Compatible Allocation and Exchange of Discrete Resources," *mimeo* UCLA, May 2015.

[14] Saban, D. and J. Sethuraman: "The Complexity of Computing the Random Priority Allocation Matrix" *mimeo* Columbia University, December 2013; extended abstract in WINE 2013.

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