

The Computational Complexity of Random Serial Dictatorship

Haris Aziz^{a,*}, Felix Brandt^b, Markus Brill^b

^a*NICTA and UNSW, 223 Anzac Parade, Sydney, NSW 2033, Australia,
Phone: +61 2 8306 0490*

^b*Institut für Informatik, Technische Universität München, 85748 Garching, Germany
Phone: +49 89 289 17512*

Abstract

In social choice settings with linear preferences, *random dictatorship* is known to be the only social decision scheme satisfying strategyproofness and *ex post* efficiency. When also allowing indifferences, *random serial dictatorship (RSD)* is a well-known generalization of random dictatorship that retains both properties. *RSD* has been particularly successful in the special domain of random assignment where indifferences are unavoidable. While *executing RSD* is obviously feasible, we show that *computing* the resulting probabilities is #P-complete and thus intractable, both in the context of voting and assignment.

Keywords: Social choice theory, random serial dictatorship, random priority, computational complexity, assignment problem.

JEL: C63, C70, C71, and C78

1. Introduction

Social choice theory studies how a group of agents can make collective decisions based on the—possibly conflicting—preferences of its members. In the most general setting, there is a set of abstract *alternatives* over which each agent entertains *preferences*. A *social decision scheme* aggregates these preferences into a probability distribution (or *lottery*) over the alternatives.

Perhaps the most well-known social decision scheme is *random dictatorship*, in which one of the agents is uniformly chosen at random and then picks his most preferred alternative. Gibbard [11] has shown that random dictatorship is the only social decision scheme that is strategyproof and *ex post* efficient, i.e., it never puts positive probability on Pareto dominated alternatives. Note that random dictatorship is only well-defined when there are no ties in the agents' preferences. However, ties are unavoidable in many important domains of social

*Corresponding author

Email addresses: haris.aziz@nicta.com.au (Haris Aziz), brandtf@in.tum.de (Felix Brandt), brill@in.tum.de (Markus Brill)

choice such as assignment, matching, and coalition formation since agents are assumed to be indifferent among all outcomes in which their assignment, match, or coalition is the same [e.g., 14].

In the presence of ties, random dictatorship is typically extended to *random serial dictatorship (RSD)*, where dictators are invoked sequentially and ties between most-preferred alternatives are broken by subsequent dictators.¹ *RSD* retains the important properties of *ex post* efficiency and strategyproofness and is well-established in the context of random assignment [see e.g., 15, 1, 5, 10].

In this paper, we focus on two important domains of social choice: (1) the *voting setting*, where alternatives are candidates and agents' preferences are given by rankings over candidates, and (2) the aforementioned *assignment setting*, where each alternative corresponds to an assignment of houses to agents and agents' preferences are given by rankings over houses. Whereas agents' preferences over alternatives are listed explicitly in the voting setting, this is not the case in the assignment setting. However, preferences over houses can be easily extended to preferences over assignments by assuming that each agent only cares about the house assigned to himself and is indifferent between all assignments in which he is assigned the same house. As a consequence, the assignment setting is a special case of the voting setting. However, due to the different representations, *computational* statements do not carry over from one setting to the other.

In various settings, the probability that a social decision scheme assigns to an alternative is interpreted as the fraction of time or another resource that is allocated to the alternative [see, e.g., 13, 6]. Similarly, in the assignment setting, the probability with which an agent is allocated a certain object is often viewed as the fraction of the object that this agent receives or the fraction of time that the agent is allowed to use the shared object. As a consequence, computing *RSD* lotteries is of great importance and has applications in domains such as scheduling of resources [see, e.g., 10].

We examine the computational complexity of *RSD* and show that computing the *RSD* lottery is $\#P$ -complete *both* in the voting setting and in the assignment setting. Loosely speaking, $\#P$ is the counting equivalent of NP —the class of decision problems whose solutions can be verified in polynomial time. $\#P$ -completeness is commonly seen as strong evidence that a problem cannot be solved in polynomial time.

As mentioned above, neither of the two results implies the other. We furthermore present a polynomial-time algorithm to compute the *support* of the *RSD* lottery in the voting setting. This is not possible in the assignment setting, because the support of the *RSD* lottery might be of exponential size. However, we can decide in polynomial time whether a given alternative (i.e., an assignment) is contained in the support or not.

¹*RSD* is referred to as *random priority* by Bogomolnaia and Moulin [5].

2. Preliminaries

In the general social choice setting, there is a set $N = \{1, \dots, n\}$ of *agents*, who have preferences over a finite set A of *alternatives*. The preferences of agent $i \in N$ are represented by a complete and transitive *preference relation* $R_i \subseteq A \times A$. The interpretation of $(a, b) \in R_i$, usually denoted by $a R_i b$, is that agent i values alternative a at least as much as alternative b . In accordance with conventional notation, we write P_i for the strict part of R_i , i.e., $a P_i b$ if $a R_i b$ but not $b R_i a$, and I_i for the symmetric part of R_i , i.e., $a I_i b$ if $a R_i b$ and $b R_i a$. A *preference profile* $R = (R_1, \dots, R_n)$ is an n -tuple containing a preference relation R_i for every agent $i \in N$.

A preference relation R_i is *linear* if $a P_i b$ or $b P_i a$ for all distinct alternatives $a, b \in A$. A preference relation R_i is *dichotomous* if $a R_i b R_i c$ implies $a I_i b$ or $b I_i c$.

We let Π^N denote the set of all permutations of N and write a permutation $\pi \in \Pi^N$ as $\pi = \pi(1) \dots \pi(n)$. For $k \leq n$, we furthermore let $\pi|_k$ denote the prefix of π of length k , i.e., $\pi|_k = \pi(1) \dots \pi(k)$.

If R_i is a preference relation and $B \subseteq A$ a subset of alternatives, then $\max_{R_i}(B) = \{a \in B : a R_i b \text{ for all } b \in B\}$ is the set of most preferred alternatives from B according to R_i . Hence, $a I_i b$ for all $a, b \in \max_{R_i}(B)$ and $a P_i b$ for all $a \in \max_{R_i}(B), b \in B \setminus \max_{R_i}(B)$.

In order to define the social decision scheme known as *random serial dictatorship* (*RSD*), let us first define its deterministic variant *serial dictatorship* (*SD*). For a given preference profile R and a permutation $\pi \in \Pi^N$, $SD(R, \pi)$ is defined via the following procedure. Agent $\pi(1)$ chooses the set of most preferred alternatives from A , $\pi(2)$ chooses his most preferred alternatives from the refined set and so on until all agents have been considered. The resulting set of alternatives is returned. Formally, $SD(R, \pi)$ is defined inductively via $SD(R, \pi|_0) = A$ and $SD(R, \pi|_i) = \max_{R_{\pi(i)}}(SD(R, \pi|_{i-1}))$.

Throughout this paper, we assume that the preferences of the agents are such that there is no pair $a, b \in A$ with $a \neq b$ and $a I_i b$ for all $i \in N$.² This assumption ensures that the set $SD(R, \pi)$ is always a singleton. We will usually write $SD(R, \pi) = a$ instead of $SD(R, \pi) = \{a\}$.

We are now ready to define *RSD*. For a given preference profile R , *RSD* returns $SD(R, \pi)$, where π is chosen uniformly at random from Π^N . The probability $RSD(R)(a)$ of alternative $a \in A$ is thus proportional to the number of permutations π for which $SD(R, \pi) = a$:

$$RSD(R)(a) = \frac{1}{n!} |\{\pi \in \Pi^N : SD(R, \pi) = a\}|.$$

We refer to the probability $RSD(R)(a)$ as the *RSD probability* of alternative a and to the probability distribution $RSD(R)$ as the *RSD lottery*.

²In the assignment setting, this assumption always holds if the agents have linear preferences over houses. *SD* (and *RSD*) can be defined without this assumption [see, e.g., 3].

Our proofs leverage the fact that a certain matrix related to the Pascal triangle has a non-zero determinant.

Lemma 1 (Bacher [4]). *The $n \times n$ matrix $M = (m_{ij})_{i,j}$ given by $m_{ij} = (i+j-2)!$ has a non-zero determinant. That is,*

$$\det \begin{pmatrix} 0! & 1! & \cdots & (n-1)! \\ 1! & 2! & \cdots & n! \\ \vdots & \vdots & \ddots & \vdots \\ (n-1)! & n! & \cdots & (2n-2)! \end{pmatrix} \neq 0.$$

3. Voting Setting

A *voting problem* is given by a triple (N, A, R) , where $N = \{1, \dots, n\}$ is a set of agents, A is a set of alternatives, and $R = (R_1, \dots, R_n)$ is a preference profile that contains, for each agent i , a preference relation on the set of alternatives. The goal is to choose an alternative that is socially acceptable according to the preferences of the agents.

If each agent has a unique most preferred alternative, the *RSD* lottery can be computed in linear time. Therefore, computational aspects of *RSD* only become interesting when at least some of the agents express indifferences among their most preferred alternatives. The straightforward approach to compute the *RSD* lottery involves the enumeration of permutations. This approach obviously takes exponential time. At first sight, it seems that even finding the support of the *RSD* lottery requires the enumeration of all permutations. However, we outline a surprisingly simple algorithm that checks in polynomial time whether a given alternative a is contained in the support (Algorithm 1).

Algorithm 1 Is the RSD probability of alternative a positive?

```

1  $N' \leftarrow N$ 
2  $A' \leftarrow A$ 
3  $\pi^* \leftarrow$  empty list
4 while  $N' \neq \emptyset$  do
5   if  $a \notin \max_{R_i}(A')$  for all  $i \in N'$  then
6     return “no”
7   else
8     Take the smallest  $i \in N'$  such that  $a \in \max_{R_i}(A')$ 
9      $N' \leftarrow N' \setminus \{i\}$ 
10     $A' \leftarrow \max_{R_i}(A')$ 
11    Append  $i$  to  $\pi^*$ 
12  end if
13 end while
14 return “yes”

```

The algorithm is based on a greedy approach and maintains a working set of alternatives A' and a working set of agents N' , which are initialized as A

and N , respectively. If no agent in N' has a as a most preferred alternative in A' , then the algorithm returns “no.” Otherwise let $i \in N'$ be the smallest index such that agent i has a as a most preferred alternative in A' .³ The set A' is refined by deleting all alternatives that are not among the most preferred alternatives in A' according to agent i . The process is then repeated until all agents have been considered.

The following lemma will be essential for showing the correctness of Algorithm 1.

Lemma 2. *Consider a permutation $\pi \in \Pi^N$ and let $1 \leq k < j \leq n$ be such that $a \in \max_{R_{\pi(j)}} SD(R, \pi|_k)$. Define another permutation $\pi' \in \Pi^N$ by moving $\pi(j)$ to position $k+1$, i.e., $\pi' = \pi(1) \dots \pi(k)\pi(j)\pi(k+1) \dots \pi(j-1)\pi(j+1) \dots \pi(n)$. If $SD(R, \pi) = a$, then $SD(R, \pi') = a$.*

Proof. Assume for contradiction that $SD(R, \pi') \neq a$. Then, there exists $i \leq n$ such that $a \in SD(R, \pi'|_{i-1})$ and $a \notin SD(R, \pi'|_i)$. That is, $a \notin \max_{R_{\pi'(i)}} (SD(R, \pi'|_{i-1}))$. Since π and π' agree on the first k positions and $SD(R, \pi) = a$, it follows that $i > k$.

Now consider the set that agent $\pi'(i)$ faces when it is his turn in the original permutation π . This set, call it B , is identical to the set $SD(R, \pi'|_{i-1})$ the agent faces in permutation π' , except that agent $\pi(j)$ might not have refined the set yet. Thus $SD(R, \pi'|_{i-1}) \subseteq B$. Since a was not among the most preferred alternatives in $SD(R, \pi'|_{i-1})$ (according to the preferences of agent $\pi'(i)$), it follows that a is not among the most preferred alternatives in B . The latter statement however contradicts the assumption that $SD(R, \pi) = a$. \square

Theorem 1. *In the voting setting, the support of the RSD lottery can be computed in polynomial time.*

Proof. It is obvious that Algorithm 1 runs in polynomial time. Since the number of alternatives is linear, we can run Algorithm 1 for each alternative and return the set of alternatives for which the algorithm returns “yes.” We now show that Algorithm 1 returns “yes” if and only if a is in the support of the RSD lottery.

If Algorithm 1 returns “yes,” it also constructs a corresponding permutation π^* such that for each $j \in N$, a is one of the most preferred alternatives of $\pi^*(j)$ in the working set of alternatives. Hence, $a \in SD(R, \pi^*)$.

For the other direction, assume that a is in the support of the RSD lottery. Then, there exists a permutation π^a such that $SD(R, \pi^a) = a$. This permutation can be transformed into the permutation π^* constructed by Algorithm 1 by only using steps that are covered by Lemma 2. In particular, we start with permutation π^a and move agent $\pi^a(1)$ to position 1. Then, agent $\pi^a(2)$ is moved to position 2, and so on until all agents are in position.

Repeated application of Lemma 2 yields that $SD(R, \pi^*) = a$. Therefore, Algorithm 1 returns “yes.” \square

³One may choose any $i \in N'$ for which a is a most preferred alternative in A' and the algorithm still works. Our choice of i with the smallest index simplifies the proof.

We now show that the problem of computing the actual RSD lottery is intractable, even when preferences are severely restricted.

Theorem 2. *In the voting setting, computing the RSD probability of an alternative is $\#P$ -complete, even for dichotomous preferences.*

Proof. We show that computing the probability $RSD(R)(a)$ of RSD choosing alternative $a \in A$ is $\#P$ -complete. Membership of this problem in the complexity class $\#P$ is straightforward. For hardness, we present a polynomial-time Turing reduction from the $\#P$ -complete problem $\#SETCOVERS-k$:

Given a set U and a collection $S = \{S_1, \dots, S_n\}$ of subsets of U , count the number of set covers of U of size k . Here, $S' \subseteq S$ is a *set cover* if $\bigcup_{S_i \in S'} S_i = U$ and the *size* of a set cover S' is $|S'|$.

$\#P$ -completeness of $\#SETCOVERS-k$ follows from the fact that counting set covers of arbitrary size is $\#P$ -complete which in turn follows from the result of Vadhan [16] that even the problem of counting *vertex covers* (a special case of set covers) is $\#P$ -complete.

We are now in a position to outline our Turing reduction.⁴ Given an instance (U, S) of $\#SETCOVERS-k$, we construct a preference profile R^k for each $k \in \{1, \dots, n\}$. The set of alternatives in R^k is $A = U \cup \{a\}$ where $a \notin U$ and the set of agents is $N^k = \{1, \dots, n+k\}$. The dichotomous preferences of the agents are specified such that

$$\begin{aligned} \text{for all } i \leq n, \quad \max_{R_i}(A) &= (U \setminus S_i) \cup \{a\}, \quad \text{and} \\ \text{for all } i > n, \quad \max_{R_i}(A) &= U. \end{aligned}$$

We now construct a system of equations in order to show that efficient computability of RSD probabilities implies efficient computability of the number of set covers of a given size. For each $k \in \{1, \dots, n\}$, let p_k denote the number of permutations $\pi \in \Pi^{N^k}$ for which $SD(R^k, \pi) = a$. Furthermore, let x_j denote the number of set covers of U of size j . We show that for all $k \in \{1, \dots, n\}$,

$$p_k = RSD(R^k)(a) \times (n+k)! = \sum_{j=1}^n j! \times k \times (n+k-j)! \times x_j. \quad (1)$$

The first equality follows directly from the definition of RSD . In order to verify that the third term in equation (1) equals p_k , envision $j+1$ as the earliest position in the permutation at which an agent $\ell \in \{n+1, \dots, n+k\}$ is present. If alternative a is chosen by SD for a given permutation, then it must be the case that the j agents preceding ℓ in the permutation must have already filtered out all the alternatives in U which means that their corresponding sets must cover U .

⁴Similar reduction techniques are used by Valiant [17].

(If the first j agents in a permutation do not filter out all the elements in U , then the agent in position $j + 1$ ensures that a is *not* chosen since all the alternatives in U are strictly more preferred to a by this agent.) In this case, neither the ordering of the first j agents nor the identity of $\ell \in \{n + 1, \dots, n + k\}$ matters. Hence, the number of permutations in which the earliest $\ell \in \{n + 1, \dots, n + k\}$ is at position $j + 1$ and a gets selected is $x_j \times j! \times k \times (n + k - j - 1)!$. Therefore, the total number of permutations of N^k in which a gets selected is $\sum_{j=1}^n j! \times k \times (n + k - j - 1)! \times x_j$.

It follows from (1) that for all $k \in \{1, \dots, n\}$,

$$\sum_{j=1}^n j! \times (n - j + k - 1)! \times x_j = RSD(R^k)(a) \times \frac{(n + k)!}{k}.$$

We get the following system of equations.

$$\begin{pmatrix} (n)!(0)! & \cdots & 1!(n - 1)! \\ \vdots & \ddots & \vdots \\ (n)!(n - 1)! & \cdots & 1!(n - 1 + n - 1)! \end{pmatrix} \begin{pmatrix} x_n \\ \vdots \\ x_1 \end{pmatrix} = \begin{pmatrix} RSD(R^1)(a) \times (n + 1)!/1 \\ \vdots \\ RSD(R^n)(a) \times (2n)!/n \end{pmatrix}$$

If the RSD probabilities can be computed in polynomial time, so can the $n \times 1$ matrix on the right hand side of the equation. Furthermore, Lemma 1 implies that the $n \times n$ matrix on the left hand side of the equation is non-singular. Hence, the values x_1, \dots, x_n can be computed in polynomial time via Gaussian elimination. The largest constants in the system of equations are of order $(2n)!$ and can be represented by $O(n \log n)$ bits. Since the variables x_1, \dots, x_n can be computed in polynomial time, it follows that $\#\text{SETCOVERS-}k$ can be solved in polynomial time. \square

Corollary 1. *The following problem is NP-hard: Given an alternative $a \in A$ and $q \in (0, 1)$, is the RSD probability of a greater than or equal to q ?*

Proof. Let $p = RSD(R)$ and note that $p(a) = c/n!$ for some integer $c \in [0, n!]$. Therefore $p(a)$ can take $n! + 1$ different values. If there was a polynomial-time algorithm to check whether $p(a) \geq q$, then we can actually compute $p(a)$ via binary search in $\log(n! + 1) < \log n^n = n \log n$ queries. \square

4. Assignment Setting

We now turn to the setting in which random serial dictatorship is most commonly studied. An *assignment problem* is given by a triple (N, H, R) , where $N = \{1, \dots, n\}$ is a set of agents, H is a set of houses with $|H| = n$, and $R = (R_1, \dots, R_n)$ is a preference profile that contains, for each agent i , a *linear* preference relation on the set of houses. The goal is to find an assignment of agents to houses.

A *deterministic* assignment is a one-to-one mapping $\sigma : A \rightarrow H$. Apart from deterministic assignments, assignment mechanisms are usually allowed to

return randomized assignments, i.e., lotteries over the set of deterministic assignments. Every randomized assignment yields a *fractional assignment* that specifies, for every agent i and every house h , the probability p_{ih} with which house h is assigned to agent i . The fractional assignment can be seen as a compact representation of the randomized assignment.

Observe that the assignment setting corresponds to a special case of the general social choice problem where the set A of alternatives is given by the set of deterministic assignments and agents' preferences over A are obtained by extending their preferences over H in such a way that each agent is indifferent between all deterministic assignments in which he is assigned the same house.⁵

Executing *RSD* in the assignment setting has a particularly natural interpretation: *choose an ordering of the agents uniformly at random and let every agent select his most preferred among all remaining houses.*

In this section, we examine the computational complexity of *RSD* in the assignment setting. Listing the probability of each deterministic assignment explicitly is prohibitive because the number of deterministic assignments is exponential in the size of the problem instance. However, it can be checked in polynomial time whether a distinguished deterministic assignment is realized with a positive probability.

Theorem 3. *It can be checked in polynomial time whether some deterministic assignment is in the support of the RSD lottery.*

We omit the straightforward proof, which invokes a simplified version of Algorithm 1.

We now show that computing the fractional assignment is intractable.⁶ We first present a useful lemma. Fix a house $h \in H$ and let s_j denote the number of sets $N' \subseteq N$ with $|N'| = j$ such that there exists a deterministic assignment in which each agent in N' gets a house that he prefers to h .

Lemma 3. *Computing s_j is #P-complete.*

Proof. Colbourn et al. [9] proved that the following problem is #P-complete:

Given an undirected and unweighted bipartite graph $G = (S \cup T, E)$ with $E \subseteq S \times T$, compute x , the number of subsets $B \subseteq S$ such that $(B \cup T, E)$ contains a perfect matching.

We propose a polynomial-time Turing reduction from computing x to computing s_j . Consider an assignment problem (N, H, R) in which $N = S$, $H = T \cup \{h\}$, and for any $i \in S$ and $h' \in T$, $h' P_i h$ if and only if $\{i, h'\} \in E$. Then, $x = \sum_{j=1}^{|S|} s_j$. Therefore, if there exists a polynomial-time algorithm to compute s_j , then there exists a polynomial-time algorithm to compute x . \square

⁵While agents' preferences over H are linear, their extended preferences over A are not.

⁶It was brought to our attention that, in recent unpublished work, Daniela Saban and Jay Sethuraman proved the statement in Theorem 4 using a different proof. They also showed that checking whether an agent gets a house with non-zero probability is NP-complete.

Theorem 4. *Computing the probability with which an agent gets a certain house under RSD is #P-complete.*

Proof. For an instance $G = (N, H, R)$ of the assignment problem, we denote by $(p_{ih}(G))_{i \in N, h \in H}$ the fractional assignment resulting from RSD.

Consider an instance $G_0 = (N, H, R)$ of the assignment problem and fix an agent $i \in N$. Let h be the most preferred house of agent i . We show that computing the probability p_{ih} is #P-complete.

By G_k we will denote a modification of G_0 in which k additional agents N_k and houses have been added. The k new agents have the same preferences over H as agent i . The k additional houses are each less preferred than houses in H by all agents in $N \cup N_k$ in any arbitrary order.

We now present a polynomial-time reduction from computing s_j to computing the probability p_{ih} in G_0 . In problem G_k , for i to get house h , it is clear that i should be earlier in the permutation than all agents in N_k . Furthermore, agents in $N \setminus \{i\}$ which are earlier than i , should all be able to get a house that they prefer to house h . Based on this insight, we obtain the following equations.

$$\begin{aligned} p_{ih}(G_0) &= \frac{1}{n!} \sum_{j=0}^{n-1} s_j \times j! \times (n-j-1)! \\ p_{ih}(G_1) &= \frac{1}{(n+1)!} \sum_{j=0}^{n-1} s_j \times j! \times (n-j-1+1)! \\ &\vdots \\ p_{ih}(G_{n-1}) &= \frac{1}{(2n-1)!} \sum_{j=0}^{n-1} s_j \times j! \times (n-j-1+n-1)! \end{aligned}$$

The system of equations can be written as

$$M \begin{pmatrix} s_{n-1} \\ s_{n-2} \\ \vdots \\ s_0 \end{pmatrix} = \begin{pmatrix} p_{ih}(G_0) \times n! \\ p_{ih}(G_1) \times (n+1)! \\ \vdots \\ p_{ih}(G_{n-1}) \times (2n-1)! \end{pmatrix},$$

where

$$M = \begin{pmatrix} (n-1)!0! & (n-2)!1! & \cdots & 0!(n-1)! \\ (n-1)!1! & (n-2)!2! & \cdots & 0!n! \\ \vdots & \vdots & \ddots & \vdots \\ (n-1)!(n-1)! & (n-2)!n! & \cdots & 0!(2n-2)! \end{pmatrix}.$$

Lemma 1 yields that the $n \times n$ matrix M is non-singular. The remaining arguments are similar to those in the proof of Theorem 2. \square

As in the voting setting, we get an NP-hardness result as an immediate corollary. The proof mirrors that of Corollary 1.

Corollary 2. *The following problem is NP-hard: Given an agent $i \in N$, a house $h \in H$, and $q \in (0, 1)$, is the probability p_{ih} greater or equal to q ?*

Theorem 4 also implies that computing the fractional assignment achieved by the *employment by lotto* method [2, 12] is #P-complete. This is due to the fact that employment by lotto is an extension of *RSD* to two-sided matching in which the set of feasible matchings is a subset of stable matchings. Employment by lotto reduces to *RSD* in the assignment model if one side is indifferent among all matchings. Another corollary is that computing the fractional assignment achieved by the *draft* mechanism [see e.g., 7, 8] is #P-complete since the draft mechanism is equivalent to *RSD* for the case in which each agent can get a maximum of one house.

Although dichotomous preferences are a very restricted class of preferences, they are not a subset of strict preferences. The proof of Theorem 4 can be adapted to prove that computing the fractional assignment achieved by *RSD* is #P-complete even if agents have dichotomous preferences over houses.

Acknowledgments

This material is based upon work supported by the Australian Government's Department of Broadband, Communications and the Digital Economy, the Australian Research Council, the Asian Office of Aerospace Research and Development through grant AOARD-124056, and the Deutsche Forschungsgemeinschaft under grants BR 2312/7-1 and BR 2312/10-1.

References

- [1] Abdulkadiroğlu, A., Sönmez, T., 1998. Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica* 66 (3), 689–702.
- [2] Aldershof, B., Carducci, O. M., Lorenc, D. C., 1999. Refined inequalities for stable marriage. *Constraints* 4, 281–292.
- [3] Aziz, H., Brandt, F., Brill, M., 2013. On the tradeoff between economic efficiency and strategyproofness in randomized social choice. In: *Proceedings of the 12th International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*. IFAAMAS, pp. 455–462.
- [4] Bacher, R., 2002. Determinants of matrices related to the Pascal triangle. *Journal de théorie des nombres de Bordeaux* 14, 19–41.
- [5] Bogomolnaia, A., Moulin, H., 2001. A new solution to the random assignment problem. *Journal of Economic Theory* 100 (2), 295–328.
- [6] Bogomolnaia, A., Moulin, H., Stong, R., 2005. Collective choice under dichotomous preferences. *Journal of Economic Theory* 122 (2), 165–184.

- [7] Budish, E., 2012. Matching “versus” mechanism design. *SIGecom Exchanges* 11 (2), 4–15.
- [8] Budish, E., Cantillion, E., 2012. The multi-unit assignment problem: Theory and evidence from course allocation at Harvard. *American Economic Review* 102 (5), 2237–2271.
- [9] Colbourn, C. J., Provan, J. S., Vertigan, D., 1995. The complexity of computing the Tutte polynomial on transversal matroids. *Combinatorica* 15 (1), 1–10.
- [10] Crès, H., Moulin, H., 2001. Scheduling with opting out: Improving upon random priority. *Operations Research* 49 (4), 565–577.
- [11] Gibbard, A., 1977. Manipulation of schemes that mix voting with chance. *Econometrica* 45 (3), 665–681.
- [12] Klaus, B., Klijn, F., 2006. Procedurally fair and stable matching. *Economic Theory* 27, 431–447.
- [13] Moulin, H., 2003. *Fair Division and Collective Welfare*. The MIT Press.
- [14] Sönmez, T., Ünver, M. U., 2011. Matching, allocation, and exchange of discrete resources. In: Benhabib, J., Jackson, M. O., Bisin, A. (Eds.), *Handbook of Social Economics*. Vol. 1. Elsevier, Ch. 17, pp. 781–852.
- [15] Svensson, L.-G., 1994. Queue allocation of indivisible goods. *Social Choice and Welfare* 11, 323–330.
- [16] Vadhan, S., 1997. The complexity of counting in sparse, regular, and planar graphs. *SIAM Journal on Computing* 31, 398–427.
- [17] Valiant, L. G., 1979. The complexity of enumeration and reliability problems. *SIAM Journal on Computing* 8 (3), 410–421.