

Pareto Optimal Allocation under Compact Uncertain Preferences

Haris Aziz

UNSW Sydney and Data61
Sydney, Australia
haris.aziz@data61.csiro.au

Peter Biro

Hungarian Academy of Sciences
Budapest, Hungary
birop@econ.core.hu

Ronald de Haan

ILLIC, University of Amsterdam
Amsterdam, the Netherlands
R.deHaan@uva.nl

Baharak Rastegari

University of Bristol, UK
baharak.rastegari@bristol.ac.uk

Abstract

The assignment problem is one of the most well-studied settings in multi-agent resource allocation. Aziz, de Haan, and Rastegari (2017) considered this problem with the additional feature that agents preferences involve uncertainty. In particular, they considered two uncertainty models neither of which is necessarily compact. In this paper, we focus on three uncertain preferences models whose size is polynomial in the number of agents and items. We consider several interesting computational questions with regard to Pareto optimal assignments. We also present some general characterization and algorithmic results that apply to large classes of uncertainty models.

1 Introduction

Multi-agent resource allocation is a widely applicable topic within multi-agent systems. One of the most fundamental settings in this domain is the so called assignment problem: n agents express preferences over n items and these items are then allocated among the agents in a suitable manner (Abdulkadiroğlu and Sönmez 1999; Aziz *et al.* 2016b; Bogomolnaia and Moulin 2001; Cechlárová *et al.* 2015; Gärdenfors 1973; Svensson 1994; 1999).

When making such an allocation, a social designer may consider several goals to identify a desirable allocation. One widely-used goal, that is considered by economists as the most fundamental requirement for any reasonable solution, is that the allocation should be Pareto optimal, i.e., there should not be another allocation in which each agent is weakly better off and at least one agent is strictly better off.

Pareto optimality is easy to achieve in classical allocation settings via sequential allocation or other matching algorithms. However, the problem may become more challenging and interesting when there is some uncertainty in the agents' preferences. Aziz, de Haan, and Rastegari (Aziz *et al.* 2017b) initiated a computational study of Pareto optimality under uncertain preferences. They discussed several reasons to consider uncertain preferences including lack of information or communication; cost of eliciting further information; a record of changing or mixed choices in the past, or the fact that the agents are in fact virtual or 'bidding' agents who represent the composition of preferences of the real agents they represent. Another possible

reason for uncertain preferences could be that agents' preferences over items may be based on several criteria, and that given each individual criterion they may have certain preferences over items, but they may not be perfectly clear on how much weight to assign to each criterion (Miyazaki and Okamoto 2017). Uncertainty in preferences has already been studied in voting (Dopazo and Martínez-Céspedes 2017; Hazon *et al.* 2012) and stable matching (Aziz *et al.* 2016a; 2017a; Miyazaki and Okamoto 2017; Chen *et al.* 2018). Similarly, in auction theory, it is standard to examine Bayesian settings in which there is probability distribution over the types of the agents.

In the context of Pareto optimal allocation, Aziz *et al.* (2017b) considered two uncertain preferences models that are not necessarily compact. In this paper, we significantly extend their study by considering further compact and natural uncertain preferences models and examining their computational aspects. If an assignment is PO (Pareto optimal) with probability one, we will call it *certainly PO*. If it is PO with non-zero probability, we call it *possibly PO*. We consider the following compact uncertainty models:

- **Compact Indifference Model:** Each agent reports a single weak preference list that allows for ties. Each complete linear order extension of this weak order is assumed to be equally likely.*
- **Pairwise Model:** Each agent reports independent pairwise probabilities over pairs of items. If i prefers item o over item o' with probability p , then she prefers o' over o with probability $1 - p$.
- **Ranking Model:** Each agent reports probabilities for each item being in each rank.[†] The input for each agent can be viewed as a bistochastic matrix, i.e, the sum of probabilities of items being in a given rank is 1, and the sum of probabilities of an item being in one of the n ranks is 1 as well.

*This can be viewed as a House Allocation problem with Ties (HRT), where any preference list obtained by breaking ties arbitrarily is possible, and all possible preferences have the same likelihood of being realized.

[†]Under a linear preference ordering, the most preferred item is ranked 1st, the second most preferred item is ranked 2nd, and so on, with the least preferred item being ranked n th.

A preference profile specifies preferences of each agent over items. Previous work (Aziz *et al.* 2017b) on Pareto optimal allocation had considered the following models: **Lottery Model** (for each agent, we are given a probability distribution over linear preferences) and **Joint Probability Model** (a probability distribution over linear preference profiles is specified).

We consider six problems proposed by Aziz *et al.* (2017b). The two most natural computational problems are:

- **PO-PROBABILITY**: what is the probability that a given assignment is PO?
- **ASSIGNMENTWITHHIGHESTPO-PROBABILITY**: compute an assignment with the highest probability of being PO.

We consider simpler problems than **PO-PROBABILITY**:

- **ISPO-PROBABILITYNON-ZERO**: for a given assignment, is the probability of being PO non-zero?
- **ISPO-PROBABILITYONE**: for a given assignment, is the probability of being PO one?

We also consider two problems connected to **ASSIGNMENTWITHHIGHESTPO-PROBABILITY**:

- **EXISTS CERTAINLY PO-ASSIGNMENT** asks whether there exists an assignment that has probability one for being PO
- **EXISTS POSSIBLY PO-ASSIGNMENT** asks whether there exists an assignment that is PO with non-zero probability.

Note that the latter problem is trivial for all uncertainty models in which the induced ‘certainly preferred’ relation is acyclic (Aziz *et al.* 2017b). An agent certainly prefers an item o to o' if she prefers o to o' with probability 1.

We say that a given uncertainty model is *independent* if any uncertain preference profile L under the model can be written as a product of uncertain preferences L_a for all agents a , where all L_a ’s are independent. All the five uncertainty models except the joint probability model are independent.

Contributions In this paper, we significantly extend the work initiated on Pareto optimal allocation under uncertain preferences in two ways: (i) we present several general insights and results that apply to large classes of uncertainty models, and (ii) we present a comprehensive set of results for three compact uncertainty models that have not been considered in the context of Pareto optimal allocation. Our technical results are summarized in Table 1. Missing proofs and the complete version of the shortened proofs can be found in the supplementary material. Note that our most interesting technical results are computational hardness results which therefore carry over to any settings in which an agent may find certain items unacceptable and/or an agent may be genuinely indifferent between two or more items.

Our algorithmic results for the compact indifference and pairwise models also extend to the setting in which agents are allowed to declare some items unacceptable, as well as the setting in which there are unequal number of items and agents. To see this, note that the proof in Aziz *et al.* 2017b (Theorem 1) and our proofs for Theorems 2 and 3 go through without any change for these two generalizations. Extending the ranking model to capture either of these two generalizations is, on the

other hand, not straightforward. This is because we heavily use the assumption that the uncertainty matrix is bistochastic when defining various concepts for the ranking model.

Discussion of the Models The compact indifference model explains scenarios where an agent may express indifference between some items because she does not have sufficient knowledge about their differences. The model is neutral to the relative ordering of these items and assumes that all linear orders consistent with the weak preferences are equiprobable. The model was considered by Aziz *et al.* (2016a) in the context of the two-sided marriage problem and with stability as the main concern. Although the model is quite restrictive, our hardness results for the model underline the fact that even uncertainty in restrictive models can lead to intractability.

The pairwise uncertainty model is well-studied in social sciences, in particular psychology, where people are asked to make repeated pairwise comparisons between different items or experiences. The model was formally studied in a related but different setting of the two-sided marriage problem (Aziz *et al.* 2017a), where the focus was on stability rather than Pareto optimality. The model is applicable to scenarios where the system has a record of similar pairwise choices and uses that record to find an outcome that has a high probability of being Pareto optimal.

The ranking model appeals to the idea that agents often ascribe ranks to items but they may not always be sure about the exact rank of each item (see e.g., Dopazo and Martínez-Céspedes 2017; Mazurek 2017). This notable approach to fuzzy rankings is applicable to scenarios where the system has a record of past ranking information and uses that fuzzy or aggregate record to find an outcome that has a high probability of being Pareto optimal.

Any compact indifference preference profile can be represented by a lottery preference profile albeit with possibly an exponential blowup. In Section 6 we will discuss how the ranking model has connections with the lottery model.

In both the compact indifference model and the ranking model, we assume that the underlying preferences of agents are linear orders. Hence the certainly preferred part of the relation is acyclic under these two models. In the pairwise probability model, however, we do not assume that comparisons are transitive and even allow agents to have cycles in their certainly preferred relations. Cyclic and intransitive preferences may happen in practice, e.g., when an agent is a virtual agent representing the preferences of a group of agents or a committee, or when agents’ valuation functions over alternatives are complex and multidimensional.

2 Preliminaries

An instance of the (*deterministic*) *assignment problem* is a triple $(N, O, >)$ where N is the set of n agents $\{1, \dots, n\}$, $O = \{o_1, \dots, o_n\}$ is the set of items, and the preference profile $> = (>_1, \dots, >_n)$ specifies complete and asymmetric preferences $>_i$ of each agent i over O . In the classical assignment problem, agents’ preferences are also assumed to be transitive, hence resulting in linearly ordered preferences. Let $>_S$ denote the preference profile of agents in the set $S \subset N$.

Problems	Compact indifference	Pairwise Probability	Ranking
PO-PROBABILITY	#P-complete	#P-complete	#P-complete
IsPO-PROBABILITYNON-ZERO IsPO-PROBABILITYONE	in P in P	in P in P	in P in P
EXISTSPOSSIBLYPO-ASSIGNMENT	in P (always exists)	NP-complete	in P (always exists)
EXISTS CERTAINLY PO-ASSIGNMENT	NP-complete	NP-complete	NP-complete
ASSIGNMENT WITH HIGHEST PO-PROB	NP-hard	NP-hard	NP-hard

Table 1: Summary of results.

An *assignment* is an allocation of items to agents such that each agent is allocated a unique item. We will represent assignments by a permutation over O so that an item in the i -th position in the permutation is given to agent i . For example, We refer by abc the assignment in which agent 1 gets a , agent 2 gets b and agent 3 gets c . For a given assignment ω , let $\omega(i)$ denote the item allocated to agent i .

An assignment ω is *PO (Pareto optimal)* if there exists no other assignment μ such that $\mu(j) \succ_j \omega(j)$ for some agent j and $\mu(i) = \omega(i)$ or $\mu(i) \succ_i \omega(i)$ for all agents i . If such an assignment μ exists, then we say that μ Pareto dominates ω .

In this work, we allow agents to express uncertainty in their preferences and consider three uncertainty models. For each agent i we define the *certainly preferred* relation \succ_i^{cert} . We write $b \succ_i^{cert} c$ if and only if agent i prefers b over c with probability 1. Checking $b \succ_i^{cert} c$ is straightforward in the pairwise and compact indifference model. We will show in Section 6 that for the ranking model it can be checked in polynomial time whether $b \succ_i^{cert} c$.

An agent i *possibly prefers* an item b to item c if and only if i prefers b over c with nonzero probability. Given an agent i , a subset of items $O' \subseteq O$, and an item $a \in O'$, we say that a is a *possibly most preferred* item for i among the items in O' if there is a deterministic preference compatible with the uncertain preferences in which a is certainly preferred over other items in O' . Checking whether a is a possibly most preferred item in O' is straightforward in the pairwise and compact indifference model. We will show in Section 6 that this problem can be solved in polynomial time for the ranking model as well.

3 General Results

We first note a well-known characterizations of PO assignments in deterministic settings. For any assignment ω , the *corresponding trading graph* is a graph on vertex set $N \cup O$, in which each item points to its owner and each agent points to items more preferred than her allocated item. The assignment ω admits a (*trading*) *cycle* $\langle o_0, i_0, o_1, i_1, \dots, o_{k-1}, i_{k-1}, o_0 \rangle$ if for all $j \in \{0, \dots, k-1\}$ we have that $o_j = \omega(i_j)$ and $o_{j+1 \bmod k} \succ_{i_j} o_j$. The following fact is well-known: an assignment is PO if and only if its corresponding trading graph does not admit a cycle (Abraham *et al.* 2004).

The trading cycle approach described above can be extended in two different ways when considering uncertain preferences: (1) each agent can point to items that are *certainly more preferred* to her current item, or (2) each agent can point to items that are *possibly more preferred* to her current item. Note that computing the certainly preferred relation directly gives us the possibly preferred relation. Aziz *et al.* (2017b) used approach (2) and characterized certainly PO assignments as those whose trading graph corresponding to ‘possibly preferred’ part of the relation does not admit a cycle. They then used the characterization to prove the following theorem.

Theorem 1 (Aziz *et al.* (2017b)) *For any independent uncertainty model in which the certainly preferred relation can be derived in polynomial time, IsPO-PROBABILITYONE can be solved in polynomial time.*

All the three compact models we consider are independent and, as remarked in Section 2, the certainly preferred relation can be derived in polynomial time in all three of them. Therefore, as a corollary, we get the following statement.

Corollary 1 *For the compact indifference, pairwise, and the ranking uncertainty models, IsPO-PROBABILITYONE can be solved in polynomial time.*

Analogous to the characterization of certainly PO assignments (Aziz *et al.* 2017b), one may presume that an assignment is possibly PO if and only if the trading graph corresponding to ‘certainly preferred’ part of the relation does not admit a cycle. However, it can be shown that the latter condition is necessary but not sufficient for possible Pareto optimality.

Next, we present a characterization of possibly PO assignments. We do this by first presenting a classic characterization of PO assignments for deterministic and linear preferences (Abdulkadiroğlu and Sönmez 1998) that is defined with respect to the outcomes of serial dictatorship. Serial Dictatorship (SD) is an assignment mechanism that is specified with respect to a permutation π over N : each agent is considered in the order stipulated by the permutation, and is allocated the item she most prefers among those that have not been

allocated yet. We will denote by $SD(N, O, \succ, \pi)$ the outcome of applying serial dictatorship with respect to permutation π over assignment problem (N, O, \succ) .

Fact 1 (Abdulkadiroğlu and Sönmez (1998)) *An assignment is PO if and only if it is an outcome of serial dictatorship.*

In the literature, SD is defined for linearly ordered preferences. We noted earlier that under pairwise preferences, we allow agents to have cycles in their preferences; e.g., an agent may certainly prefer b to c , c to d , and d to b . There is nothing to prevent one from running SD on deterministic pairwise preferences. However, when an agent's turn arrives she may not find an item she most prefers, due to the existence of a cycle, in which case SD must abort without any solution. It can thus be shown that Fact 1 applies to pairwise preferences as well. We present the following general theorem that applies to any uncertainty model and is a general characterization of possibly PO assignments. **This theorem is an immediate consequence of Fact 1 if all realizable deterministic preferences are linear. We are providing a detailed proof to cover the case in which cycles may be present in deterministic preferences.**

Theorem 2 *For any uncertainty model, an assignment ω is possibly PO if and only if there is some permutation π and some preference profile that has non-zero probability under which when serial dictatorship is applied with respect to π , each agent in her turn gets item $\omega(i)$.*

Proof: (\Leftarrow) Follows from Fact 1. If some deterministic preference profile \succ' has non-zero probability such that ω is the outcome of $SD(N, O, \succ', \pi)$, then ω is possibly PO.

(\Rightarrow) Since ω is possibly PO, it is PO with respect to some realizable deterministic profile \succ' . We show that $\omega = SD(N, O, \succ', \pi)$ for some permutation of agents π . Take any partial allocation $\omega' \subset \omega$ that allocates to a subset of agents $S \subset N$; that is, $\omega'(i) = \omega(i)$ for all $i \in S$. Denote the items that are allocated in ω' by O' . (Note that ω' and hence S and O' can be empty sets.) We claim that there exists some agent $i \in N \setminus S$ such that, with respect to \succ' , $\omega(i)$ is the most preferred item of agent i among the set of items $O \setminus O'$ (note that, by the definition of ω' and O' , $\omega(i)$ is in $O \setminus O'$). Suppose for a contradiction that there exists no such an agent. Therefore, every agent $j \in N \setminus S$ is interested in, and in the corresponding trading graph points to, an item (or more) that is held (according to ω) by another agent in $N \setminus S$. This implies the existence of a trading cycle where some agents in $N \setminus S$ can exchange items among themselves to get a more preferred item than in ω , implying that ω is not PO with respect to the deterministic profile \succ' , a contradiction. Therefore, we have established that starting with $\omega' = \emptyset$, we can obtain ω by iteratively finding an agent who has not been allocated yet and, according to \succ' , has $\omega(i)$ as her most preferred item among the unallocated ones. Such an agent always exists as proved above. Let the order in which the agents are chosen be π . It is easy to verify that $\omega = SD(N, O, \succ', \pi)$. \square

We call an uncertainty model *reasonable* if, for any subset of items $O' \subseteq O$ and any agent $i \in N$, it can be checked in

polynomial time whether $o \in O'$ is a possibly most preferred item for i among items in O' . It is easy to verify that the lottery, compact indifference, and the pairwise models are all reasonable. Although it is not obvious, we will show in Section 6 that the ranking model is reasonable as well. Next, we present another general result that applies to a large class of uncertainty models.

Theorem 3 *For any independent and reasonable uncertainty model, ISPO-PROBABILITYNON-ZERO can be solved in polynomial time.*

Proof: Consider an assignment ω that we want check whether it is possibly PO. We use the following algorithm that builds a permutation of agents π such that serial dictatorship produces ω given π , if and only if ω is possibly PO. To start with, we initialize the set of remaining items to O , the remaining agents to N , and the permutation of the agents π to an empty list. We then repeat the following procedure until no more items are left, or the procedure returns no. Check if there exists some agent i such that $\omega(i)$ is a possibly most preferred item for i amongst the available items. If no such agent exists, return no. Otherwise, if such an i exists, append i to the permutation π , remove i from the set of remaining agents, and remove $\omega(i)$ from the set of available items. Let \succ'_i denote a preference of agent i that has $\omega(i)$ as the most preferred remaining item.

It is easy to verify that if the algorithm returns π , then $SD(N, O, \succ', \pi) = \omega$. It can be shown that if the algorithm returns no, then ω is PO with zero probability. Consider the first point in the algorithm where for no agent i we have $\omega(i)$ as a possibly most preferred available item for i . This means that no remaining agent gets her most preferred item (for any possible deterministic preferences) among the available items. Therefore, for each realization of the preference profiles, each of the remaining agents is interested in, and points to, another item held by another agent among the remaining agents. This implies the existence of a trading cycle for each realization of the preference profiles, where some remaining agents can exchange items among themselves to get a more preferred item than in ω . Thus ω is PO with probability zero. \square

We note that the theorem above generalizes Theorem 5 from (Aziz *et al.* 2017b) that only applies to the lottery model. It also gives us the following corollary.

Corollary 2 *For the compact indifference, pairwise, and the ranking uncertainty models, ISPO-PROBABILITYNON-ZERO can be solved in polynomial time.*

4 Compact Indifference Model

Any compact indifference preference profile admits a possibly PO assignment that can be computed as follows: break the ties arbitrarily and then run SD. What about the existence of certainly PO assignments? Although the compact indifference model is one of the most restricted and structured uncertainty models we consider, EXISTCERTAINLYPO-ASSIGNMENT is NP-complete for this model. To prove this, and several other claims in the paper, we reduce from the NP-complete problem

SERIALDICTATORSHIPFEASIBILITY (Saban and Sethuraman 2015) that is defined as follows: given an instance of the assignment problem (N, O, \succ') , an agent i and an item $o \in O$, does there exist a permutation of agents for which SD awards o to i ?

Theorem 4 *For the compact indifference uncertainty model, EXISTCERTAINLYPO-ASSIGNMENT is NP-complete.*

Proof: EXISTCERTAINLYPO-ASSIGNMENT is in NP because it can be checked in polynomial time whether a given assignment is certainly PO or not (Theorem 1).

To prove NP-hardness, we reduce from SERIALDICTATORSHIPFEASIBILITY. Let (N, O, \succ', i, o) be an instance of this problem with n agents and n items. W.l.o.g., suppose that $i = 1$. We construct an instance (N', O', \succ) of the compact indifference uncertainty model as follows. Let o', o'' be fresh items not appearing in O , let $O' = O \cup \{o', o''\}$, and let $N' = N \cup \{1', n + 1\}$.

Let $O_{1, \succ' o}$ denote the set of items that agent 1 prefers over item o with respect to \succ'_1 —i.e., those items o_ℓ such that $o_\ell \succ'_1 o$. Moreover, let $O_{1, \prec' o}$ denote the set of items that agent 1 likes less than item o —i.e., $O_{1, \prec' o} = O \setminus (\{o\} \cup O_{1, \succ' o})$. Also, for each $1 < j \leq n$, let j_1, \dots, j_n be the indices of items such that $o_{j_1} \succ'_j \dots \succ'_j o_{j_n}$.

We define the preferences of the agents in N' over the items in O' as follows, where items in round brackets are tied.

- Agent 1: $(O_{1, \succ' o} \cup \{o'\}) \succ_1 o \succ_1 (O_{1, \prec' o} \cup \{o''\})$.
- Agent 1': $(O_{1, \succ' o} \cup \{o'\}) \succ_{1'} o \succ_{1'} (O_{1, \prec' o} \cup \{o''\})$.
- Agents $1 < j \leq n$: $o_{j_1} \succ_j \dots \succ_j o_{j_n} \succ_j (o', o'')$
- Agent $n + 1$: (O') .

We claim that there is a permutation π of the agents in N' under which agent $i = 1$ gets item o when serial dictatorship is run on (N', O', \succ) if and only if there is an assignment that is certainly PO for (N, O, \succ') . \square

Corollary 3 *For the compact indifference uncertainty model, ASSIGNMENTWITHHIGHESTPO-PROB is NP-hard.*

Theorem 5 *For the compact indifference uncertainty model, PO-PROBABILITY is #P-complete.*

Proof: It is straightforward to show that PO-PROBABILITY is in #P. We show #P-hardness by reducing from the #P-complete problem Monotone-#2SAT that is defined as follows: count the number of satisfying assignments for a 2CNF formula that contains no negation (Valiant 1979).

Let φ be a monotone 2CNF formula with clauses c_1, \dots, c_m and variables x_1, \dots, x_n . We construct an instance of PO-PROBABILITY as follows. Consider agents $1, \dots, 2n$, and items o_1, \dots, o_{2n} and take the assignment ω where $\omega(i) = o_i$ for all $1 \leq i \leq 2n$. Agents' preferences are constructed as follows. Take an arbitrary $1 \leq i \leq n$. Consider the set $\{j_1, \dots, j_u\}$ of indices j such that the clause $(x_i \vee x_j)$ occurs in φ . Suppose that $j_1 < j_2 < \dots < j_u$, in order to fix an (arbitrary) order over these indices. Then agent i has the following preference: $(o_{n+i}, o_i) \succ_i \dots$, where the remaining items appear in arbitrary order. Moreover, agent $n+i$ has the following preference:

$o_{j_1} \succ_{n+i} \dots \succ_{n+i} o_{j_u} \succ_{n+i} o_{n+i} \succ_{n+i} \dots$, where the remaining items appear in arbitrary order. Note that agents $n + 1, \dots, 2n$ have certain preferences.

The instantiations of linear preferences to uncertain agents $1, \dots, n$ correspond one-to-one to truth assignments to variables x_1, \dots, x_n : for each $1 \leq i \leq n$, setting the preference of agent i to $o_i \succ_i o_{n+i} \succ_i \dots$ corresponds to setting x_i to true, and setting agent i 's preference to $o_{n+i} \succ_i o_i \succ_i \dots$ corresponds to setting x_i to false.

We show that a truth assignment α to the variables x_1, \dots, x_n satisfies φ if and only if the corresponding profile of linear preferences for the agents $1, \dots, 2n$ leads the assignment ω to be PO.

(\Rightarrow) Suppose that α satisfies φ . We show that the assignment ω admits no trading cycle under the profile corresponding to α . Suppose, to derive a contradiction, that ω admits a trading cycle. By construction of the preferences of the agents, any trading cycle must involve the sequence $\langle o_i, i, o_{n+i}, n+i, o_j, j, o_{n+j}, n+j \rangle$. That is, agent i must prefer o_{n+i} over o_i , agent j must prefer o_{n+j} over o_j , and agent $n+i$ must prefer o_j over o_{n+i} . The former two statements are the case if and only if α sets both x_i and x_j to false. The latter statement is the case if and only if $(x_i \vee x_j)$ is a clause of φ . This is a contradiction to our assumption that α satisfies φ . Thus, we can conclude that ω is PO under the profile corresponding to α .

(\Leftarrow) Conversely, suppose that α does not satisfy φ . That is, there is a clause $(x_i \vee x_j)$ of φ such that α sets both x_i and x_j to false. We show that the assignment ω admits a trading cycle under the profile corresponding to α . Since α sets x_i and x_j to false, we know that $o_{n+i} \succ_i o_i$ and $o_{n+j} \succ_j o_j$. Moreover, since $(x_i \vee x_j)$ is a clause of φ , we know that $o_j \succ_{n+i} o_{n+i}$ and $o_i \succ_{n+j} o_{n+j}$. Thus, $\langle o_i, i, o_{n+i}, n+i, o_j, j, o_{n+j}, n+j \rangle$ is a trading cycle under this preference profile, and hence ω is not PO.

Each possible preference profile occurs with probability 2^{-n} . The number of satisfying truth assignments φ is then exactly equal to 2^n times the probability that assignment ω is PO. Thus, PO-PROBABILITY is #P-hard. \square

5 Pairwise Model

A possibly PO assignment exists for the pairwise model, if the 'certainly preferred' relation of each agent is acyclic: one can derive a possible linear order consistent with the certainly preferred relation and then run serial dictatorship. In the presence of cycles, however, there is no guarantee that an instance admits a possibly PO assignment.

Theorem 6 *Under the pairwise uncertainty model, if the certainly preferred relation is cyclic, then there may not exist a possibly PO assignment.*

Proof: We prove this by showing that the following example with three agents and three items does not admit a possibly PO assignment. The pairwise preferences for the agents 1, 2 and 3 over the items a, b and c are as depicted. All the three

agents have the same certainly preferred relations.

- 1 : $a \succ_1^{cert} c, c \succ_1^{cert} b, b \succ_1^{cert} a$
- 2 : $a \succ_2^{cert} c, c \succ_2^{cert} b, b \succ_2^{cert} a$
- 3 : $a \succ_3^{cert} c, c \succ_3^{cert} b, b \succ_3^{cert} a$

Take assignment abc . It is not PO because it is Pareto dominated by bca . By symmetry no other assignment is PO as well. \square

We next show, with an argument similar to that of the proof of Theorem 3 in (Aziz *et al.* 2017b), that the problem of checking whether a PO assignment exists is NP-complete under pairwise preferences even in the absence of uncertainty.

Theorem 7 *Given an instance of the assignment problem with deterministic pairwise preferences, the problem of deciding if a PO assignment exists is NP-complete.*

Proof: Under deterministic preferences, possibly PO and certainly PO coincide. The problem is in NP because it can be checked in polynomial time whether a given assignment is certainly PO (Corollary 1).

To prove NP-hardness we reduce from SERIALDICTATORSHIPFEASIBILITY. Let (N, O, \succ, i, o) be an instance of this problem with n agents and n items. We construct an instance (N', O', \succ) of deterministic pairwise preferences as follows. Let o', o'' be fresh items not appearing in O , let $O' = O \cup \{o', o''\}$, and let $N' = N \cup \{n+1, n+2\}$.

Intuitively, the (linear) preference \succ_i^{cert} is obtained from \succ'_i by replacing o with $o \succ o' \succ o''$. The preference \succ_j^{cert} for each $j \in N \setminus \{i\}$ is obtained from \succ'_j by replacing o with the cycle $o \succ o' \succ o'' \succ o$. Finally, the preferences \succ_{n+1}^{cert} and \succ_{n+2}^{cert} are constructed by having the cycle $o \succ o' \succ o'' \succ o$ on the top, followed by the items in $O \setminus \{o\}$ (in arbitrary order). Formally:

- For agent i :
 - For each pair of items o_1, o_2 such that $o \notin \{o_1, o_2\}$, let $o_1 \succ_i^{cert} o_2$ if and only if $o_1 \succ'_i o_2$.
 - Let $o \succ_i^{cert} o', o \succ_i^{cert} o''$ and $o' \succ_i^{cert} o''$.
 - For each item $o_1 \in O \setminus \{o\}$, let $o_1 \succ_i^{cert} o, o_1 \succ_i^{cert} o'$, and $o_1 \succ_i^{cert} o''$ if and only if $o_1 \succ'_i o$.
 - For each item $o_1 \in O \setminus \{o\}$, let $o \succ_i^{cert} o_1, o' \succ_i^{cert} o_1$, and $o'' \succ_i^{cert} o_1$ if and only if $o \succ'_i o_1$.
- For each agent $j \in N \setminus \{i\}$:
 - For each pair of items o_1, o_2 such that $o \notin \{o_1, o_2\}$, let $o_1 \succ_j^{cert} o_2$ if and only if $o_1 \succ'_j o_2$.
 - Let $o \succ_j^{cert} o', o' \succ_j^{cert} o''$ and $o'' \succ_j^{cert} o$.
 - For each item $o_1 \in O \setminus \{o\}$, let $o_1 \succ_j^{cert} o, o_1 \succ_j^{cert} o'$, and $o_1 \succ_j^{cert} o''$ if and only if $o_1 \succ'_j o$.
 - For each item $o_1 \in O \setminus \{o\}$, let $o \succ_j^{cert} o_1, o' \succ_j^{cert} o_1$, and $o'' \succ_j^{cert} o_1$ if and only if $o \succ'_j o_1$.
- For both agents $j \in \{n+1, n+2\}$:
 - For each pair of items $o_\ell, o_{\ell'} \in O \setminus \{o\}$, let $o_\ell \succ_j^{cert} o_{\ell'}$ if and only if $\ell > \ell'$.

- For each item $o_\ell \in O \setminus \{o\}$, let $o \succ_j^{cert} o_\ell, o' \succ_j^{cert} o_\ell$, and $o'' \succ_j^{cert} o_\ell$.
- Let $o \succ_j^{cert} o', o' \succ_j^{cert} o''$ and $o'' \succ_j^{cert} o$.

We claim that there is a permutation π under which agent i gets item o when serial dictatorship is run on (N, O, \succ) if and only if there is an assignment that is PO for (N', O', \succ) . (\Rightarrow) Suppose that there is a permutation π under which agent i gets item o when serial dictatorship is run on (N, O, \succ) (resulting in assignment ω). Construct the permutation π' for agents N' obtained from π by placing the agents $n+1, n+2$ directly after agent i . Then $\text{Gen-SD}(N', O', \succ^{cert}, \pi')$ returns the assignment $p' = p \cup \{n+1 \mapsto o', n+2 \mapsto o''\}$. This assignment p' is PO.

(\Leftarrow) Suppose that there is an assignment p' for (N', O', \succ^{cert}) that is PO. By Theorem 2, we know that there exists a permutation π of the agents in N' such that executing Gen-SD with this permutation on (N', O', \succ^{cert}) yields p' . Let $\pi = (\pi_0, i, \pi_1)$, i.e., π_0 is the sequence of agents appearing in π before agent i , and π_1 is the sequence of agents appearing in π after agent i . By construction of \succ^{cert} , we know that o is picked before o' and o'' , because o' and o'' are dominated in every preference order whenever o has not yet been picked. Moreover, we know that agent i must pick item o , because whenever o, o', o'' have not yet been picked, item o is dominated in every preference order except agent i 's. Thus, after the agents in π_0 have picked the unique undominated remaining item in their preference order, all items that agent i prefers to item o (in \succ'_i) have been picked. Note that agents $n+1$ and $n+2$ are not in π_0 since, as long as item o is not picked, both of these agents are confused. Thus, (π_0, i) can be extended (by appending the rest of agents in N in arbitrary order) to a permutation π' under which agent i gets item o when serial dictatorship is run on (N, O, \succ) with π' . \square

Since possibly PO and certainly PO assignments coincide with PO assignments for deterministic pairwise preferences, we get the following corollaries.

Corollary 4 *For the pairwise uncertainty model, EXISTS-POSSIBLYPO-ASSIGNMENT and EXISTS-CERTAINLYPO-ASSIGNMENT are NP-complete even if pairwise preferences are all deterministic.*

The proof of Theorem 7 exploits the fact that certainly preferred relations can be cyclic under the pairwise uncertainty model. What if we disallow cycles in certainly preferred relations? Our next result shows that checking whether a certainly PO assignment exists is NP-complete for instances of pairwise uncertainty model, even if the certainly preferred part of the preferences is acyclic.

Theorem 8 *For the pairwise uncertainty model, EXISTS-CERTAINLYPO-ASSIGNMENT is NP-complete even if the certainly preferred part of the preferences is acyclic.*

As a corollary, for the pairwise model, ASSIGNMENT-WITH-HIGHEST-PO-PROB is NP-hard even if the certainly preferred part of the preferences is acyclic.

Theorem 9 For the pairwise uncertainty model, PO-PROBABILITY is #P-complete.

Proof: It is straightforward to show that PO-PROBABILITY is in #P. Hardness for #P can be shown by following the reduction used in the proof of Theorem 5. In this reduction, each agent either has a certain linear preference, or a preference where they are indifferent between their two most preferred items (and have a certain linear preference over all other items). These preferences can straightforwardly be expressed in the pairwise model, hence PO-PROBABILITY is #P-hard under this model as well. \square

6 Ranking Model

In the ranking model, each agent i reports a bistochastic matrix M_i of size $n \times n$. The rows of the matrix correspond to indices of items and the columns to the rank of the items. Note that if the matrix is a permutation matrix (with only 0-1 entries), then the ranking model degenerates to a linear preference. We say that a permutation matrix P is *consistent* with a bistochastic matrix Q if for each $P_{ij} = 1$, we have that $Q_{ij} > 0$. We say that a linear preference $>_i$ is consistent with a ranking preference if the permutation matrix representing $>_i$ is consistent with the ranking preference matrix. For any ranking preference matrix M_i for agent i , we will also consider a corresponding ‘consistency graph’ which is a bipartite graph $(O \cup R, E)$ where $R = \{1, \dots, n\}$ is the set of possible ranks and $(o, r) \in E$ if i expresses non-zero probability for o to be in rank r .

By Birkhoff’s theorem (Lovász and Plummer 2009), any bistochastic matrix can be represented as a convex combination of at most quadratic number of permutation matrices. Hence, a ranking preference profile can always be represented as a lottery preference profile. This representation is not necessarily unique; that is, a ranking preference profile may have several or possibly even exponentially many lottery preference profile representation. Due to this reason, it is not clear whether the PO probability of a matching is well-defined under the ranking model. In view of this we clarify the definitions of possibly PO and certainly PO that we are using for the ranking model. We say that an assignment is certainly PO under the ranking model if it is certainly PO for all lottery preference profiles that represent the ranking preference profile. Likewise, we say that an assignment is possibly PO under the ranking model if it is possibly PO for some lottery preference profile that represents the ranking preference profile. Next, we argue that Corollaries 1 and 2 regarding ISPO-PROBABILITYONE and ISPO-PROBABILITYNONZERO hold under the ranking model by showing that “possibly preferred” (and hence “certainly preferred”) and “possibly most preferred item in a set” can be checked in polynomial time.

Given a ranking preference profile $>$, we say that $b >_i^{\text{cert}} c$ if $b >'_i c$ in each linear preference profile $>'$ that is in the support of some lottery preference profile that represents $>$. Note that a linear preference is in the support of some lottery preference that represents a ranking preference if and only if the linear preference is consistent with the ranking preference. The negation of $b >_i^{\text{cert}} c$ holds if there exist $j, k \in [n]$ such that $j < k$ and there is a linear preference consistent with the

ranking preferences in which b gets rank k and c gets rank j . The relation $b \not>_i^{\text{cert}} c$ can be tested by checking whether there exist $j, k \in [n]$ such that $j < k$ and the consistency graph corresponding to the ranking preferences admits a perfect matching in which (b, k) and (c, j) are part of the matching.

Given a ranking preference profile $>$, we say that a is a possibly most preferred item of agent i among a set of items O' , if there is a linear preference consistent with $>$ in which a is the most preferred item in set O' . The latter can be checked in polynomial time as follows. For each possible rank j for item a , we construct a corresponding graph $G_j = (O \cup R, E)$ which is derived from the consistency graph by removing all edges involving a except (a, j) , as well as removing all edges (o', k) where $o' \in O' \setminus \{a\}$ and $k \leq j$. Then a is a possibly most preferred item of i in O' if and only if there exists a perfect matching for some G_j .

Any ranking preference profile admits a possibly PO assignment that can be computed by finding a linear preference profile consistent with the ranking preference profile and then running serial dictatorship.

Theorem 10 For the ranking uncertainty model, a possibly PO assignment always exists and can be computed in polynomial time.

On the other hand, just like the pairwise and the compact indifference models, EXISTS CERTAINLY PO-ASSIGNMENT is NP-complete under the ranking model. The proof is an adaptation of a result by Aziz *et al.* (2017b, Theorem 6).

Theorem 11 For the ranking uncertainty model, EXISTS CERTAINLY PO-ASSIGNMENT is NP-complete.

Next we show that PO-PROBABILITY is #P-complete for instances for which PO probability is well defined due to the existence of a unique lottery model representation.

Theorem 12 For the ranking uncertainty model, PO-PROBABILITY is #P-complete.

Proof: It is straightforward to show that PO-PROBABILITY is in #P. Hardness for #P can be shown by using the reduction used in the proof of Theorem 5. In this reduction, each agents either has a certain linear preference, or a preference where they are indifferent between their two most preferred items (and have a certain linear preference over all other items). These preferences can straightforwardly be expressed in the ranking model. Thus, for the ranking model, PO-PROBABILITY is #P-hard as well. \square

7 Conclusions

We presented some general characterization and algorithmic results that apply to large classes of uncertainty models. We then extended the study to three uncertainty models that are especially compact (at most polynomial in the number of agents and items). For the pairwise model, it will be interesting to see if hardness results still hold if we impose stronger types of stochastic transitivity properties (Fishburn 1973). Another interesting direction is considering ex ante Pareto

optimality with respect to plausible utility functions (Aziz *et al.* 2015). We observe that most of our computational results are similar for the three compact models we consider. One interesting contrast is that EXISTSPOSSIBLYPO-ASSIGNMENT is polynomial-time solvable for the Compact Indifference model and the Ranking model whereas it is NP-complete for the Pairwise model.

References

- A. Abdulkadiroğlu and T. Sönmez. Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica*, 66(3):689–701, 1998.
- A. Abdulkadiroğlu and T. Sönmez. House allocation with existing tenants. *Journal of Economic Theory*, 88:233–260, 1999.
- D.J. Abraham, K. Cechlárová, D.F. Manlove, and K. Mehlhorn. Pareto optimality in house allocation problems. In *Proceedings of ISAAC '04: the 15th Annual International Symposium on Algorithms and Computation*, volume 3341 of *Lecture Notes in Computer Science*, pages 3–15. Springer, 2004.
- H. Aziz, F. Brandl, and F. Brandt. Universal Pareto dominance and welfare for plausible utility functions. *Journal of Mathematical Economics*, 60:123–133, 2015.
- H. Aziz, P. Biró, S. Gaspers, R. de Haan, N. Mattei, and B. Rastegari. Stable matching with uncertain linear preferences. In *Proceedings of the 9th International Symposium on Algorithmic Game Theory (SAGT)*, pages 195–206, 2016.
- H. Aziz, P. Biró, J. Lang, J. Lesca, and J. Monnot. Optimal reallocation under additive and ordinal preferences. In *Proceedings of the 15th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 402–410, 2016.
- H. Aziz, P. Biró, T. Fleiner, S. Gaspers, R. de Haan, N. Mattei, and B. Rastegari. Stable matching with uncertain pairwise preferences. In *Proceedings of the 16th Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 344–352, 2017.
- H. Aziz, R. de Haan, and B. Rastegari. Pareto optimal allocation under uncertain preferences. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1472–1474, 2017.
- A. Bogomolnaia and H. Moulin. A new solution to the random assignment problem. *Journal of Economic Theory*, 100(2):295 – 328, 2001.
- K. Cechlárová, P. Eirínakis, T. Fleiner, D. Magos, D. F. Manlove, I. Mourtos, E. Oceláková, and B. Rastegari. Pareto optimal matchings in many-to-many markets with ties. In *Proceedings of the 8th International Symposium on Algorithmic Game Theory (SAGT)*, Lecture Notes in Computer Science (LNCS), pages 27–39. Springer-Verlag, 2015.
- J. Chen, R. Niedermeier, and P. Skowron. Stable marriage with multi-modal preferences. In *Proceedings of the 19th Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 269–286, 2018.
- E. Dopazo and M. L. Martínez-Céspedes. Rank aggregation methods dealing with ordinal uncertain preferences. *Expert Systems with Applications*, 78(C):103–109, July 2017.
- P. C. Fishburn. Binary choice probabilities: on the varieties of stochastic transitivity. *Journal of Mathematical Psychology*, 10:327–352, 1973.
- P. Gärdenfors. Assignment problem based on ordinal preferences. *Management Science*, 20:331–340, 1973.
- N. Hazon, Y. Aumann, S. Kraus, and M. Wooldridge. On the evaluation of election outcomes under uncertainty. *Artificial Intelligence*, 189:1–18, 2012.
- L. Lovász and M. D. Plummer. *Matching Theory*. AMS Chelsea Publishing, 2009.
- J. Mazurek. Fuzzy rankings: Properties and applications. *CoRR*, abs/1703.05201, 2017.
- S. Miyazaki and K. Okamoto. Jointly stable matchings. In Yoshio Okamoto and Takeshi Tokuyama, editors, *28th International Symposium on Algorithms and Computation (ISAAC 2017)*, volume 92 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 1–56:12. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2017.
- D. Saban and J. Sethuraman. The complexity of computing the random priority allocation matrix. *Mathematics of Operations Research*, 40(4):1005–1014, 2015.
- L.-G. Svensson. Queue allocation of indivisible goods. *Social Choice and Welfare*, 11:323–330, 1994.
- L.-G. Svensson. Strategy-proof allocation of indivisible goods. *Social Choice and Welfare*, 16:557–567, 1999.
- L. G. Valiant. The complexity of enumeration and reliability problems. *SIAM Journal on Computing*, 8(3):410–421, 1979.