

Pareto Optimal Allocation under Uncertain Preferences

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

Ronald de Haan
 University of Amsterdam
 Amsterdam, the Netherlands
 R.deHaan@uva.nl

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

Abstract

The assignment problem is one of the most well-studied settings in social choice, matching, and discrete allocation. We consider this problem with the additional feature that agents' preferences involve uncertainty. The setting with uncertainty leads to a number of interesting questions including the following ones. How to compute an assignment with the highest probability of being Pareto optimal? What is the complexity of computing the probability that a given assignment is Pareto optimal? Does there exist an assignment that is Pareto optimal with probability one? We consider these problems under two natural uncertainty models: (1) the lottery model in which each agent has an independent probability distribution over linear orders and (2) the joint probability model that involves a joint probability distribution over preference profiles. For both of these models, we present a number of algorithmic and complexity results highlighting the difference and similarities in the complexity of the two models.

1 Introduction

Multi-agent resource allocation and dealing with uncertainty are two big topics in AI. In this paper we examine optimal allocation of resources under uncertain preferences.

When preferences of agents are aggregated to identify a desirable social outcome, *Pareto optimality* is a minimal requirement. Pareto optimality stipulates that there should not be another outcome that is at least as good for all agents and better for at least one agent. We take Pareto optimality as a central concern and consider a richer version of the classical *assignment problem* where the twist is that agents may express uncertainty in their preferences. The assignment problem is a fundamental setting in which n agents express preferences over n items and each agent is to be allocated one item. The setting is a classical one in discrete allocation. Its axiomatic and computational aspects have been well-studied [Abdulkadiroğlu and Sönmez, 1999; Abraham *et al.*, 2005; Aziz *et al.*, 2015; 2016b; Bogomolnaia and Moulin, 2001; Gärdenfors, 1973; Svensson, 1994;

1999]. Our motivation for studying assignment with uncertain preferences is that agents' preferences may not be completely known because of lack of information or communication. In some settings, eliciting preferences from agents may be costly, so a central planner may want to only obtain, and provide a recommendation based on, a subset of the complete orders [Rastegari *et al.*, 2013; 2014; Drummond and Boutilier, 2014]. Another possible motivation is that agents are in fact virtual or 'bidding' agents who are each representing a group of real agents and the virtual agent's probabilistic preferences simply represent the composition of preferences of the real agents it represents.

Our work is inspired by the recent work of Aziz *et al.* [2016a] who examined the stable marriage problem under uncertain preferences. Uncertainty in preferences has already been studied in voting [Hazon *et al.*, 2012]. Similarly, in auction theory, it is standard to examine Bayesian settings in which there is probability distribution over the types of the agents. Although computational aspects of Pareto optimal outcomes have been intensely studied in various settings such as assignment, matching, housing markets, and committee voting [Abraham *et al.*, 2005; Aziz and de Keijzer, 2012; Aziz *et al.*, 2015; 2016c; Erdil and Ergin, 2015; Krysta *et al.*, 2014; Manlove, 2013; Saban and Sethuraman, 2013], there has not been much work on Pareto optimality under uncertain preferences. In the presence of uncertainty, one can relax the goal of computing a Pareto optimal outcome and focus on computing outcomes that have the *highest probability* of being Pareto optimal. We will abbreviate Pareto optimal as PO. If an assignment is Pareto optimal with probability one, we will call it certainly PO.

We consider the following uncertainty models:

- **Lottery Model:** For each agent, we are given a probability distribution over linear preferences.
- **Joint Probability Model:** A probability distribution over linear preference profiles* is specified.

The most natural computational problems that we will consider are as follows.

- **PO-PROBABILITY:** what is the probability that a given assignment is PO?

*A preference profile specifies (deterministic) preferences of each agent over items.

- **ASSIGNMENTWITHHIGHESTPO-PROBABILITY**: compute an assignment with the highest probability of being PO.

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

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

We also consider a problem connected to **ASSIGNMENTWITHHIGHESTPO-PROBABILITY**: **EXISTS CERTAINLY PO-ASSIGNMENT** asks whether there exists an assignment that is PO with probability one. Note that **EXISTS POSSIBLY PO-ASSIGNMENT**—the problem of checking whether there exists some PO assignment with non-zero probability—is trivial for all uncertainty models in which the induced *certainly preferred* relation is acyclic. An agent certainly prefers an item to another if the preference is with probability 1. The reason for the triviality is that the certainly preferred relation can be completed in a way so that it is transitive, and then for the completed deterministic preferences there exists at least one PO assignment.

Note that the product of the independent uncertain preferences in the lottery model results in a probability distribution over preference profiles and hence can be represented in the joint probability model. However, the change in representation can result in a blowup. Thus whereas the joint probability model is more general than the lottery model, it is not as compact. In view of this, complexity results for one model do not directly carry over to results for the other model.

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 [Aziz *et al.*, 2016a]. Note that the lottery model is independent but the joint probability model is not.

Contributions We initiate the first study of computational aspects of Pareto optimal allocation under uncertain preferences. We show that for both the lottery model and the joint probability model, **EXISTS CERTAINLY PO-ASSIGNMENT** is NP-complete. We also prove that **ASSIGNMENTWITHHIGHESTPO-PROBABILITY** is NP-hard for both models. In view of these results, we see that as we move from deterministic preferences to uncertain preferences, the complexity of computing Pareto optimal assignments jumps significantly. On the other hand, we show that for a general class of independent uncertainty models, both problems **ISPO-PROBABILITYNON-ZERO** and **ISPO-PROBABILITYONE** can be solved in linear time. Whereas **PO-PROBABILITY** is polynomial-time solvable for the joint probability model, we prove that the problem is #P-complete for the lottery model. This problem becomes polynomial-time solvable for the lottery model if there is a constant number of uncertain agents. Moreover, we show that the problem **PO-PROBABILITY** for the lottery model can be solved in fixed-parameter tractable time when parameterized by the number of uncertain agents. Our results are summarized in Table 1.

2 Preliminaries

The setting we consider is the *assignment problem* which is a triple (N, O, \succ) where N is the set of n agents $\{1, \dots, n\}$, $O = \{o_1, \dots, o_n\}$ is the set of n items, and $\succ = (\succ_1, \dots, \succ_n)$ is a preference profile that specifies complete, asymmetric, and transitive preferences \succ_i of each agent i over O . We write $o \succ_i o'$ if $o \succ_i o'$ or $o' = o$. We will denote by $\mathcal{R}(O)$ the set of all complete and transitive relations over the set of items O . We will denote by \succ_S the preference profile of agents in the set $S \subset N$.

An *assignment* is an allocation of items to agents, represented as an $n \times n$ matrix $[p(i)(o_j)]_{1 \leq i \leq n, 1 \leq j \leq n}$ such that for all $i \in N$, and $o_j \in O$, $p(i)(o_j) \in \{0, 1\}$; for all $i \in \{1, \dots, n\}$, $\sum_{j \in N} p(i)(o_j) = 1$; and for all $j \in \{1, \dots, n\}$, $\sum_{i \in N} p(i)(o_j) = 1$. The aforementioned constraints ensure that each item is allocated wholly (items are indivisible), each agent is allocated exactly one item, and each item is allocated to exactly one agent. An agent i gets item o_j if and only if $p(i)(o_j) = 1$. Each row $p(i) = (p(i)(o_1), \dots, p(i)(o_n))$ represents the *allocation* of agent i . With an abuse of notation we write $p(i) = o_j$ if $p(i)(o_j) = 1$ for some item o_j , and $p(i) = \emptyset$ otherwise. We always assume that every agent prefers any item to \emptyset , i.e. to being allocated no item.

An assignment p is *Pareto optimal* if there does not exist another assignment q such that $q(i) \succ_i p(i)$ for all $i \in N$ and $q(i) \succ_i p(i)$ for some $i \in N$. If such an assignment q exists, then we say that q Pareto dominates p .

We first note a couple of well-known characterisations of Pareto optimal assignments. The characterisations give rise to linear-time algorithms for computing and verifying Pareto optimal assignments in settings with deterministic preferences.

An assignment p 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 $p(i_j)(o_j) = 1$ and $o_{j+1 \bmod k} \succ_{i_j} o_j$. Intuitively, each agent i_j in the trading cycle prefers the next item o_{j+1} (o_0 if $j = k-1$) to his assigned item o_j . Thus, the existence of a trading cycle implies that there is a set of agents who all benefit from trading their assigned items among themselves.

Fact 1 (Folklore) *An assignment is Pareto optimal if and only if it does not admit a trading cycle.*

We will also use the following characterization of Pareto optimal discrete assignments [Abdulkadiroğlu and Sönmez, 1998] that is defined with respect to outcomes of serial dictatorship. Serial dictatorship is a straightforward greedy assignment mechanism that is specified with respect to a permutation π over N . The mechanism takes each agent in turn, according to the permutation π , and allocates him the most preferred item on his preference list that has 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 2 (Abdulkadiroğlu and Sönmez [1998]) *An assignment is Pareto optimal if and only if it is an outcome of serial dictatorship.*

Problems	Lottery Model	Joint Probability Model
PO-PROBABILITY	#P-complete (Thm 8) but in FPT w.r.t. parameter: # uncertain agents (Thm 9)	in P (Thm 2)
IsPO-PROBABILITYNON-ZERO IsPO-PROBABILITYONE	in P (Thm 5) in P (Thm 4)	in P (Thm 2) in P (Thm 2)
EXISTSPOSSIBLYPO-ASSIGNMENT	in P (trivially)	in P (trivially)
EXISTS CERTAINLY PO-ASSIGNMENT	NP-complete (Thm 6)	NP-complete (Thm 3)
ASSIGNMENT WITH HIGHEST PO-PROB	NP-hard (Cor 3)	NP-hard (Cor 2)

Table 1: Summary of results.

Fact 2 also follows from Proposition 1 of Brams and King [2005]. The facts above show that when preferences are deterministic, a Pareto optimal assignment can be computed and verified easily [Abraham *et al.*, 2005]. In this paper we focus on finding and verifying Pareto optimal assignments in settings where agents have uncertain preferences.

Example 1 Consider the following assignment problem in which agent 1 has uncertain preferences in the lottery model.

$$\begin{array}{ll}
 1 : & a, b, c \text{ (0.6)} & 2 : & b, a, c \\
 & b, a, c \text{ (0.4)} & 3 : & c, b, a
 \end{array}$$

The same uncertain preferences can be also represented in the joint probability model with a probability distribution over two preference profiles.

Consider the assignment q in which 1 gets a , 2 gets b , and 3 gets c . The probability of this assignment being PO is 1. This can be verified by considering each of the two possible preference profiles, and testing that no other assignment Pareto dominates q under either of them. On the other hand, the assignment p in which 1 gets b , 2 gets a , and 3 gets c has 0.4 probability of being PO. This is because p is not PO if the first possible preference list of agent 1 is realized, i.e., if $a \succ_1 b \succ_1 c$. To see this, notice that p admits a trading cycle $\langle b, 1, a, 2, b \rangle$, implying that 1 and 2 prefer to trade their items and be both better off. If 1 and 2 trade their assigned items, we get assignment q . Assignment q Pareto dominates assignment p when the first possible preference list of agent 1 is realized.

Assignment p is PO if the second possible preference list of agent 2 is realized, i.e., if $b \succ_2 a \succ_2 c$. Note that p is the outcome of applying serial dictatorship with respect to permutation $\pi = 1, 2, 3$: (1) the first agent in the permutation, agent 1, is allocated his most preferred item, item b , (2) agent 2 is allocated his most preferred available item, that is a , as b has been already allocated to 1, (3) agent 3 is allocated his most preferred available item, that is c .

3 Joint Probability Model

We first observe that PO-PROBABILITY can be solved easily for the joint probability model.

Theorem 2 For the joint probability model, PO-PROBABILITY can be solved in polynomial time.

Proof: The probability that a given assignment p is PO is equivalent to the total probability weight of the preference profiles for which p is PO. This can be calculated as follows. For each preference profile, we test (in polynomial time) whether p is PO. We then add the probabilities of those profiles for which p is PO. The sum of the probabilities is the probability that the assignment p is PO. \square

Corollary 1 For the joint probability model, IsPO-PROBABILITYNON-ZERO and IsPO-PROBABILITYONE can be solved in polynomial time.

What about EXISTS CERTAINLY PO-ASSIGNMENT? This problem is equivalent to checking whether the sets of PO assignments associated with each possible preference profile have a non-empty intersection. We show that this problem is NP-complete even when the probability distribution is over two linear preference profiles.

We reduce from the NP-complete problem SERIALDICTATORSHIPFEASIBILITY that is defined as follows: check whether there exists a permutation of agents for which serial dictatorship gives a particular item o to a particular agent i [Saban and Sethuraman, 2015].

For linear preference profiles, the set of Pareto optimal allocations are characterized by those that can be achieved via some serial dictatorship. Thus it follows that the following problem is also NP-complete: check whether there exists a Pareto optimal allocation in which a specified agent i gets a specified item o .

Theorem 3 For the joint probability model, EXISTS CERTAINLY PO-ASSIGNMENT is NP-complete even when the probability distribution is over two linear preference profiles.

Proof: EXISTS CERTAINLY PO-ASSIGNMENT is in NP because it can be checked in polynomial time whether a given assignment is certainly PO or not (Theorem 2).

To prove NP-hardness, we reduce from SERIALDICTATORSHIPFEASIBILITY problem. Let (N, O, \succ, i, o) be an instance of

this problem. We construct a joint probability over two preference profiles. One of the profiles is the same as \succ . In the other preference profile \succ' , agent i has o as the most preferred item and has the same order of preference over all other items as in \succ_i . Each agent $j \in N \setminus \{i\}$, on the other hand, has o as the least preferred item and has the same order of preference over all other items as in \succ_j .

Our first observation is that an assignment is PO under profile \succ' only if i receives o in it (**Observation 1**). We now prove that we have a yes instance of SERIALDICTATORSHIPFEASIBILITY if and only if our constructed joint probability instance admits a certainly PO assignment. Since SERIALDICTATORSHIPFEASIBILITY is NP-complete, this will entail that EXISTSCERTAINLYPO-ASSIGNMENT is NP-complete.

Assume that there exists a certainly PO assignment p . Then, p must be PO under \succ' implying that, by Observation 1, i receives o in p . Assignment p must also be PO under profile \succ which implies that there exists an assignment (i.e. p) that is PO under profile \succ and in which i receives o . In light of Fact 2, this implies that there exists a serial dictatorship the outcome of which under profile \succ is p . Hence, we have a yes instance of SERIALDICTATORSHIPFEASIBILITY.

Now consider the case when we have a yes instance of SERIALDICTATORSHIPFEASIBILITY. This means that there is a permutation π under which i gets o when serial dictatorship is run. Let us call this assignment by p . Due to Fact 2, p is PO under preference profile \succ . So it only remains to show that p is PO under \succ' . Due to Fact 2, it is sufficient to prove that for profile \succ' , there exists a corresponding permutation of agents under which the outcome of serial dictatorship is p .

In fact, we show that $SD(N, O, \succ', \pi) = p$. That is, the outcome of applying serial dictatorship with permutation π is p even if the preference profile is \succ' instead of \succ . This proof is by induction on the rounds of serial dictatorship and by showing that at the end of any given round, all agents whose turn has already come up are assigned the same items as in p . Detailed proof is removed due to shortage of space. Thus p is PO under both possibly realizable preference profiles. \square

Corollary 2 *For the joint probability model, ASSIGNMENTWITHHIGHESTPO-PROB is NP-hard.*

Proof: We show that an algorithm to solve ASSIGNMENTWITHHIGHESTPO-PROB can be used to solve the NP-complete problem EXISTSCERTAINLYPO-ASSIGNMENT with only a polynomial increase in the running time. An algorithm for ASSIGNMENTWITHHIGHESTPO-PROB can compute an assignment p with the highest probability of being PO. By Corollary 1, it can be checked in polynomial time whether p is PO with probability one or not. If p is PO with probability one, then we know that we have a yes instance of EXISTSCERTAINLYPO-ASSIGNMENT. Otherwise, we have a no instance of EXISTSCERTAINLYPO-ASSIGNMENT. \square

4 Lottery Model

Recall that a given uncertainty model is *independent* if any uncertain preference profile L under the model can be writ-

ten as a product of uncertain preferences L_a for all agents a , where all L_a 's are independent.

We first define the *certainly preferred* relation \succ_i^{certain} for agent i . We write $b \succ_i^{\text{certain}} c$ if and only if agent i prefers b over c with probability 1.

Theorem 4 *For any independent uncertainty model, IsPO-PROBABILITYONE can be solved in polynomial time.*

Proof sketch: Given an assignment p , we create a trading cycle graph G with agents and items as vertices. Each item points to its owner agent and each agent i points to any item o such that $p(i) \not\succeq_i^{\text{certain}} o$. We show that p is PO with probability one if and only if G does not contain a cycle. \square

Since the lottery model is an independent uncertainty model, Theorem 4 applies to it. Next we prove that for the lottery uncertainty model, IsPO-PROBABILITYNON-ZERO can be solved in polynomial time.

Theorem 5 *For the lottery model, IsPO-PROBABILITYNON-ZERO can be solved in polynomial time.*

Proof: Consider an assignment p for which we want to check whether it is PO with non-zero probability. We use the following algorithm that builds a permutation of agents π such that serial dictatorship produces p given π , if and only if p is PO with non-zero probability. 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 $p(i)$ is the most preferred available item of i in at least one of his preference lists.
- If no such agent exists, return no. Otherwise, if such an i exists, give agent i item $p(i)$, append i to the permutation π , remove i from the set of remaining agents, and remove $p(i)$ from the set of available items. Let \succ_i denote the preference of agent i that has $p(i)$ as the most preferred remaining item (if more than one of such preferences exists, select one arbitrarily).

It is easy to verify that if the algorithm returns π , then $SD(N, O, \succ, \pi) = p$. It remains to show that if the algorithm returns no, then p is PO with zero probability. Consider the first point in the algorithm where for no agent i we have $p(i)$ as the most preferred available item of i in at least one of his preference lists. This means that no remaining agent gets his most preferred item (for any preference list) 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 p . Thus p is PO with probability zero. \square

In contrast to the positive results above, the problem of checking whether there exists an assignment that is PO with probability one is NP-complete. The proof is very similar to the proof of Theorem 3 and is hence removed. Note that Theorem 3 does not directly imply the following claim and an independent argument, albeit similar one, is required.

Theorem 6 *For the lottery model, EXISTS CERTAINLY PO-ASSIGNMENT is NP-complete.*

Corollary 3 *For the lottery model, ASSIGNMENT WITH HIGHEST PO-PROB is NP-hard.*

We now turn to the problem of computing the probability that a given assignment is PO. We first present a polynomial-time solution for a restricted setting, and then show that PO-PROBABILITY is #P-complete for the lottery model in general.

Theorem 7 *For the lottery model, if the number of uncertain agents is constant, then PO-PROBABILITY is polynomial-time solvable.*

Proof: Let p be a given assignment. Let constant k denote the number of uncertain agents, and let the maximum number of preferences for any uncertain agent be ℓ . Therefore, the maximum number of preference profiles that are realizable is ℓ^k which is still polynomial in the input since $k = O(1)$. For each possible preference profile \succ , it is easy to compute the probability that \succ occurs by simply computing the product of the probabilities of the preferences chosen for the uncertain agents. In this way, we can reduce PO-PROBABILITY for the lottery model to PO-PROBABILITY for the joint probability model which can be solved in polynomial time (Theorem 2). \square

Theorem 8 *For the lottery model, PO-PROBABILITY is #P-complete, even when restricted to the case where each agent has at most two possible preferences.*

Proof sketch: 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 $1, \dots, n$ and items o_1, \dots, o_n , and take the assignment p where $p(i) = o_i$.

We construct the preferences of the agents as follows. Take an arbitrary agent i . Consider the set $\{j_1, \dots, j_u\}$ of indices j such that the clause $(x_i \vee x_j)$ occurs in φ . (W.l.o.g. this set $\{j_1, \dots, j_u\}$ is non-empty.) Suppose that $j_1 < j_2 < \dots < j_u$, in order to fix an (arbitrary) order over these indices. With probability $\frac{1}{2}$, agent i has o_i at the top of his preference list, followed by the rest of the items in arbitrary order. With probability $\frac{1}{2}$, agent i has the following preference: $o_{j_1} \succ_i \dots \succ_i o_{j_u} \succ_i o_i \succ_i \dots$, where the remaining items appear in arbitrary order after o_i .

This way, the possible preference profiles correspond one-to-one to the possible truth assignments over x_1, \dots, x_n .

Namely, taking the preference $o_i \succ_i \dots$ for agent i corresponds to setting x_i to 1, and taking the other preference for agent i corresponds to setting x_i to 0. Moreover, each possible preference profile occurs with probability $1/2^n$.

We show that the number of satisfying assignments for φ is equal to the number of preference profiles under which p is Pareto optimal. In particular, we can show that p is PO under a preference profile if and only if the corresponding truth assignment T satisfies φ .

The number of satisfying truth assignments of φ is then exactly equal to 2^n times the probability that assignment p is Pareto optimal. Thus, PO-PROBABILITY is #P-hard, even when restricted to the case where each agent has at most two possible preferences. \square

We showed that when there are only a constant number of uncertain agents, we can compute the PO probability in polynomial time for the lottery model (Theorem 7). However, the order of the polynomial that upper bounds the running time of our proposed algorithm grows with the number of uncertain agents. In particular, when k is the number of uncertain agents, and ℓ is the maximum number of possible preference lists for these uncertain agents, the running time of the algorithm outlined in the proof of Theorem 7 is $\Omega(\ell^k)$. We improve on this result by showing that there exists a fixed-parameter tractable algorithm that computes the PO probability for the lottery model. That is, we provide an algorithm that runs in time $f(k)n^c$ for some computable function f and some fixed constant c independent of k , where n denotes the input size. In other words, we show that the parameterized problem k -PO-PROBABILITY, where the parameter is the number of uncertain agents, is fixed-parameter tractable for the lottery model.

Theorem 9 *For the lottery model, k -PO-PROBABILITY can be solved in fixed-parameter tractable time.*

Proof: Take an arbitrary instance of the problem k -PO-PROBABILITY, consisting of agents $1, \dots, n$, items o_1, \dots, o_n , and an assignment σ . W.l.o.g. assume that the assignment gives each agent i item o_i , and that the uncertain agents are agents $1, \dots, k$. For each uncertain agent i , let $\succ_{i,1}, \dots, \succ_{i,u_i}$ denote the different possible preferences for agent i . Additionally, assume w.l.o.g. that for each of the uncertain agents $1, \dots, k$, each of the possible preferences for these agents occurs with probability ℓ/d where the numerator ℓ can vary between different agents and different possible preferences, but where the denominator d is common among all agents and all possible preferences. In other words, all probabilities mentioned in the instance are rational numbers that share a common denominator d . If this were not the case, we could straightforwardly transform the instance in polynomial time to an equivalent instance that does satisfy this property. Also, assume w.l.o.g. that σ admits no trading cycle that involves only certain agents. If this were the case, then σ is PO with probability zero, and we can filter out such trivial instances using the result of Theorem 5.

We now show how to compute the probability that σ is PO in fixed-parameter tractable time. Our computation will proceed in three stages:

- (1) We construct a directed graph G with $O(ku2^{k^2})$ vertices, where the edges are weighted. Here u denotes the maximum number of possible preferences for any uncertain agent.
- (2) We count the number of homomorphisms f of a directed path P_{2k+2} of length $2k + 2$ to this graph G , where each homomorphism is counted multiple times according to (the product of) the weights on the edges in $f(P_{2k+2})$. This can be done in polynomial time using an extension of a known algorithm [Flum and Grohe, 2004; 2006].
- (3) We divide the weighted total number of homomorphisms of P_{2k+2} to G by the number d^k to obtain the probability that the given assignment σ is PO.

We begin with stage (1), and we construct the weighted, directed graph G . Let $\Pi = \{o_1, \dots, o_k\}^2$ be the set of all possible pairs (o_i, o_j) of items among o_1, \dots, o_k . We define the set V of vertices of G as follows. First, we define an auxiliary set V' : $V' = \{1, \dots, k + 1\} \cup \{(i, >_{i,j}) \mid i \in [k], j \in [u_i]\}$. Then, we define the set V of vertices as follows: $V = \{s, t\} \cup \{(v', \Pi') \mid v' \in V', \Pi' \subseteq \Pi\}$. That is, the graph G has vertices s and t , and 2^{k^2} copies of each element in V' (one for each $\Pi' \subseteq \Pi$). Intuitively, the vertices s and t will act as source and target for each homomorphism of P_{2k+2} to G .

The sets $\Pi' \subseteq \Pi$ will be intuitively used to memorize the ‘trading paths’ (i.e., paths in the trading cycle graph) that result from particular choices of the preference lists $>_{i,j}$ chosen for the agents $1, \dots, k$. That is, each $(o_i, o_j) \in \Pi'$ corresponds to a path from o_i to o_j in the directed graph with vertices o_1, \dots, o_n where there is an edge from $o_{i'}$ to $o_{j'}$ if and only if agent i' prefers item $o_{j'}$ to item $o_{i'}$.

We construct the set E of (weighted and directed) edges as follows.

- We add an edge with weight 1 from s to $(1, \emptyset)$.
- For each $i \in [k]$, $j \in [u_i]$, and $\Pi' \subseteq \Pi$, we add an edge from (i, Π') to $(i, >_{i,j}, \Pi')$. This edge has weight ℓ , where the possible preference list $>_{i,j}$ for agent i occurs with probability ℓ/d .
- For each $i \in [k]$, $j \in [u_i]$, and $\Pi' \subseteq \Pi$, we add an edge with weight 1 from $(i, >_{i,j}, \Pi')$ to the vertex $(i + 1, \Pi'')$, where the set Π'' (with $\Pi' \subseteq \Pi'' \subseteq \Pi$) consists of all pairs $(o_{i'}, o_{j'})$ for which there is a path from $o_{i'}$ to $o_{j'}$ in the ‘trading cycle graph,’ constructed using $>_{i,j}$ and Π' . To construct Π'' , consider the following graph $G_{\Pi', >_{i,j}}$. The vertices of this graph are o_1, \dots, o_n . For each pair $(o_{i'}, o_{j'})$ of vertices among o_{k+1}, \dots, o_n , there is an edge from $o_{i'}$ to $o_{j'}$ if and only if agent j prefers item $o_{j'}$ to item $o_{i'}$. Moreover, for each $(o_{i'}, o_{j'}) \in \Pi'$, we add an edge from $o_{i'}$ to $o_{j'}$. Finally, for each agent $o_{i'}$ among o_{k+1}, \dots, o_n , we add an edge from o_i to $o_{i'}$ if and only if $o_{i'} >_{i,j} o_i$. We then let $\Pi'' \subseteq \Pi$ be the set of all pairs $(o_{i'}, o_{j'})$ such that there is a path from $o_{i'}$ to $o_{j'}$ in $G_{\Pi', >_{i,j}}$. Clearly, $\Pi' \subseteq \Pi'' \subseteq \Pi$.
- For each $\Pi' \subseteq \Pi$ such that $(o_i, o_i) \notin \Pi'$ for all $i \in [k]$, we add an edge with weight 1 from $(k + 1, \Pi')$ to t .

Using this construction, we can show that the choices $>_{1,j_1}, \dots, >_{k,j_k}$ of preferences for the agents $1, \dots, k$ that make the

assignment Pareto optimal are in one-to-one correspondence with the homomorphisms f from P_{2k+2} to G . (Due to space reasons, we omit a formal proof of this claim.)

We count each such homomorphism f in a weighted fashion as follows—this is phase (2). Take a homomorphism f from P_{2k+2} to G . Its weight in the grand total is the product of the weights for each edge in $f(P_{2k+2})$. The only edges in $f(P_{2k+2})$ that have weight > 1 are edges from (i, Π') to $(i, >_{i,j}, \Pi')$. Such an edge has weight ℓ , where the probability that $>_{i,j}$ occurs is ℓ/d . Then, the total weighted sum of all homomorphisms is equal to $p \cdot d^k$, where p is the probability that the given assignment is Pareto optimal. Therefore, in order to compute p , we only need to take the weighted sum of all homomorphisms, and divide it by d^k —this is phase (3).

All that remains is to argue how we can compute the weighted sum of all homomorphisms f from P_{2k+2} to G in polynomial time. We can do this by extending a known polynomial-time algorithm to count the number of homomorphisms of a graph whose treewidth is bounded by a fixed constant into another graph [Flum and Grohe, 2006, Theorem 14.7]. Since paths have treewidth 1, counting the number of homomorphisms from a path to another graph can be done in polynomial time using this algorithm. This algorithm uses a dynamic programming approach to count the number of homomorphisms. This dynamic programming technique can straightforwardly be extended to take into account the weights of the homomorphisms. (We omit a detailed description of the extended algorithm.)

This concludes our description of the fixed-parameter tractable algorithm that solves k -PO-PROBABILITY for the lottery model. \square

An additional interesting problem that one might consider is to decide, given two assignments, which one has a higher probability of being PO. We conjecture that this problem is PP-complete for the lottery model.

5 Conclusions

Computing Pareto optimal outcomes is an active line of research in economics and computer science. We see that as we move from deterministic preferences to uncertain preferences, the complexity of computing Pareto optimal outcomes jumps significantly even though the input for problems we study may not be compact. The computational hardness results carry over to more complex models in which there may be more items than agents, agents may have capacities, and items may have copies. For future work, we have started looking into other uncertainty models of compact indifference and pairwise [Aziz *et al.*, 2016a]. An orthogonal but equally interesting direction will be to consider other fairness, stability, or efficiency desiderata.

Acknowledgments

We are grateful to the anonymous IJCAI reviewers for their helpful comments. Haris Aziz is supported by a Julius Career Award. Ronald de Haan is supported by the Austrian Science Fund (FWF), projects P26200 and J4047.

References

- [Abdulkadiroğlu and Sönmez, 1998] Atila Abdulkadiroğlu and Tayfun Sönmez. Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica*, 66(3):689–701, 1998.
- [Abdulkadiroğlu and Sönmez, 1999] Atila Abdulkadiroğlu and Tayfun Sönmez. House allocation with existing tenants. *Journal of Economic Theory*, 88(2):233–260, 1999.
- [Abraham *et al.*, 2005] David J. Abraham, Katarína Cechlárová, David F. Manlove, and Kurt Mehlhorn. Pareto optimality in house allocation problems. In *Proceedings of the 16th International Symposium on Algorithms and Computation (ISAAC)*, volume 3341 of *Lecture Notes in Computer Science (LNCS)*, pages 1163–1175, 2005.
- [Aziz and de Keijzer, 2012] Haris Aziz and Bart de Keijzer. Housing markets with indifference: a tale of two mechanisms. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI)*, pages 1249–1255, 2012.
- [Aziz *et al.*, 2015] Haris Aziz, Simon Mackenzie, Lirong Xia, and Chun Ye. Ex post efficiency of random assignments. In *Proceedings of the 14th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 1639–1640, 2015.
- [Aziz *et al.*, 2016a] Haris Aziz, Péter Biró, Serge Gaspers, Ronald de Haan, Nicholas Mattei, and Baharak Rastegari. Stable matching with uncertain linear preferences. In *Proceedings of the 9th International Symposium on Algorithmic Game Theory (SAGT)*, pages 195–206, 2016.
- [Aziz *et al.*, 2016b] Haris Aziz, Péter Biró, Jérôme Lang, Julien Lesca, and Jérôme 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.
- [Aziz *et al.*, 2016c] Haris Aziz, Jérôme Lang, and Jérôme Monnot. Computing Pareto Optimal Committees. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 60–66, 2016.
- [Bogomolnaia and Moulin, 2001] Anna Bogomolnaia and Hervé Moulin. A new solution to the random assignment problem. *Journal of Economic Theory*, 100(2):295–328, 2001.
- [Brams and King, 2005] Steven J. Brams and Daniel L. King. Efficient fair division: Help the worst off or avoid envy? *Rationality and Society*, 17(4):387–421, 2005.
- [Drummond and Boutilier, 2014] Joanna Drummond and Craig Boutilier. Preference elicitation and interview minimization in stable matchings. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI)*, pages 645–653. AAAI Press, 2014.
- [Erdil and Ergin, 2015] Aytekin Erdil and Haluk Ergin. Two-sided matching with indifference. June 2015.
- [Flum and Grohe, 2004] Jörg Flum and Martin Grohe. The parameterized complexity of counting problems. *SIAM Journal on Computing*, 33(4):892–922, 2004.
- [Flum and Grohe, 2006] Jörg Flum and Martin Grohe. *Parameterized Complexity Theory*. Texts in Theoretical Computer Science. Springer-Verlag, 2006.
- [Gärdenfors, 1973] Peter Gärdenfors. Assignment problem based on ordinal preferences. *Management Science*, 20:331–340, 1973.
- [Hazon *et al.*, 2012] Noam Hazon, Yonatan Aumann, Sarit Kraus, and Michael Wooldridge. On the evaluation of election outcomes under uncertainty. *Artificial Intelligence*, 189:1–18, 2012.
- [Krysta *et al.*, 2014] Piotr Krysta, David Manlove, Baharak Rastegari, and Jinshan Zhang. Size versus truthfulness in the house allocation problem. In *Proceedings of the 15th ACM Conference on Economics and Computation (ACM-EC)*, pages 453–470. ACM Press, 2014.
- [Manlove, 2013] David Manlove. *Algorithmics of Matching Under Preferences*. World Scientific Publishing Company, 2013.
- [Rastegari *et al.*, 2013] Baharak Rastegari, Anne Condon, Nicole Immorlica, and Kevin Leyton-Brown. Two-sided matching with partial information. In *Proceedings of the 14th ACM Conference on Electronic Commerce (ACM-EC)*, pages 733–750. ACM Press, 2013.
- [Rastegari *et al.*, 2014] Baharak Rastegari, Anne Condon, Nicole Immorlica, Robert Irving, and Kevin Leyton-Brown. Reasoning about optimal stable matchings under partial information. In *Proceedings of the 15th ACM Conference on Economics and Computation (ACM-EC)*, pages 431–448. ACM Press, 2014.
- [Saban and Sethuraman, 2013] Daniela Saban and Jay Sethuraman. House allocation with indifference: a generalization and a unified view. In *Proceedings of the 14th ACM Conference on Electronic Commerce (ACM-EC)*, pages 803–820. ACM Press, 2013.
- [Saban and Sethuraman, 2015] Daniela Saban and Jay Sethuraman. The complexity of computing the random priority allocation matrix. *Mathematics of Operations Research*, 40(4):1005–1014, 2015.
- [Svensson, 1994] Lars-Gunnar Svensson. Queue allocation of indivisible goods. *Social Choice and Welfare*, 11:323–330, 1994.
- [Svensson, 1999] Lars-Gunnar Svensson. Strategy-proof allocation of indivisible goods. *Social Choice and Welfare*, 16(4):557–567, 1999.
- [Valiant, 1979] Leslie G. Valiant. The complexity of enumeration and reliability problems. *SIAM Journal on Computing*, 8(3):410–421, 1979.