Popular and Dominant Matchings with Uncertain and Multimodal Preferences

Gergely Csáji[1,2]
[1] HUN-REN KRTK KTI
[2] ELTE Eötvös Loránd University
csaji.gergely@krtk.hun-ren.hu

Abstract

We study the Popular Matching (PM) problem in multiple models, where the preferences of the agents in the instance may change or may be unknown or uncertain. In particular, we study an Uncertainty model, where each agent has a possible set of preference lists, a Multilayer model, where there are layers of preference profiles, and a Robust popularity model, where any agent may move some other agents up or down some places in his preference list.

We study both one-sided (only one class of the agents have preferences) and two-sided bipartite markets. In the one-sided model, we show that all our problems can be solved in polynomial time by utilizing the structure of popular matchings. We also obtain nice structural results. With two-sided preferences, we show that all three above models lead to NP-hard questions for popular matchings. By using the connection between dominant matchings and stable matchings, we show that in the robust and uncertainty models, a certainly dominant matching in all possible preference profiles can be found in polynomial time, whereas in the multilayer model, the problem remains NP-hard for dominant matchings too.

We also answer an open question about d-robust stable matchings.

1 Introduction

Matching markets form a key area in computer science, economics and mathematics literature because of their demonstrated usefulness in many practical and theoretical applications, as well as their rich and beautiful structure. The intensive study of the area started after the seminal paper of Gale and Shapley [1962]. They introduced the famous stable marriage problem and provided a simple algorithm that always finds a stable matching. Intuitively, stability means that no two agents mutually prefer each other to their assigned partners. Stability turned out to be an essential property of matchings in many real-world scenarios such as resident allocation to hospitals (HR), student allocation to schools/universities and many more. For a nice overview of the field of preference-based matching markets, see for example [Manlove, 2013].

A problem with stable matchings is that their size may not be as large as some applications require. For example, in school choice it is important that all students are assigned to some school. In fact, a stable matching may only be half the size of a maximum matching.To overcome this issue, a relaxation of stability, called popularity was introduced by Gardenfors [1975]. Roughly saying, a matching is popular if it does not lose a head-to-head comparison to any other matching. He showed that stable matchings are popular, implying that such matchings also always exist. Furthermore, their size may even be two times the size of stable matchings. While with stable matchings we wish to avoid local blocking structures, with popular matching the goal is only a global form of stability (or fairness).

In recent matching market literature, popular matchings became a central area, partly because of their beautiful connection to stable matchings (see e.g. [Kavitha, 2014]). For a survey about popular matchings, see [Cseh *et al.*, 2017].

In some cases, the participating agents in a matching market may have unknown or partially known preferences to the central authority, or even the agents themselves might not be sure about their own preferences, i.e. they are uncertain. Alternatively, agents may rank their possible partners differently according to different criteria, which gives rise to multilayer preference profiles, one for each such criteria. These kinds of problem have only recently come into attention in the related literature. In such problems, we want to compute stable or popular matchings, that are certainly stable or popular based on the information that we know about the preferences of the agents. Such problems have been studied for stable matchings ([Aziz *et al.*, 2016; Chen *et al.*, 2018]) and Pareto-optimal matchings ([Aziz *et al.*, 2019]). However, so far there has been no work on similar questions for popular matchings. In our paper, we consider the corresponding questions for popularity that have been studied for stability and Pareto-optimality in the literature and provide a comprehensive overview of their complexity.

1.1 Related Work

Popularity was introduced in [Gärdenfors, 1975] for two-sided matching markets. He showed that every stable matching is popular.

Abraham et al. [2007] introduced popularity for the one-sided preference model, also referred to as the House Allocation (HA) model. They provided an algorithm for finding a maximum size popular matching in the model.

Since popular matchings are easy to find, interest shifted to finding maximum size popular matchings in two-sided markets too. Huang and Kavitha [2013] proved that even maximum size popular matchings can be found in polynomial time. Furthermore, in [Cseh and Kavitha, 2018] it was shown that a nonempty subset of maximum size popular matchings, called dominant matchings of an instance I are in correspondence with the stable matchings of another instance I'.

Models where the preferences of agents can vary or may be uncertain have been studied in many papers in the recent literature.

Uncertain preferences have been studied by many papers (e.g. [Aziz *et al.*, 2016; Aziz *et al.*, 2019]). In [Aziz *et al.*, 2016] they studied both a Lottery model, where there was a probability distribution for each different agent over some possible preferences, and a Joint Probability model, where the probability distribution was only over certain preference profiles. In both frameworks, they studied when can we find a matching that is stable with probability one (that is, stable in all possible preferences). In [Aziz *et al.*, 2019], they studied similar questions for one-sided preferences and Pareto optimality as a condition instead of stability. Our work is very similar to these, as we consider many of the same problems for one-sided and two-sided popularity as they considered for Pareto-optimality and stability.

Miyazaki and Okamoto [2017] studied the jointly stable matching problem, where the goal is to find stable matchings across multiple instances with the same agents and contracts, but different preferences, which is equivalent to finding a matching that is stable with probability 1, if each possible preference profile has some positive probability. They strengthened the result of [Aziz *et al.*, 2016] and showed that it is NP-hard even for two preference profiles.

Multilayer preferences in bipartite matching markets have been studied in [Chen *et al.*, 2018]. They defined α-global and α-pair stability. The former means that we aim to find a matching that is stable in at least α layers, and the latter means that each edge can block in at most α layers. The bipartite graph of the instances in each layer is the same; hence, if α is the number of layers, then α-global stability becomes equivalent to deciding whether there is a stable matching across multiple preference profiles. They proved that deciding whether an α-global or α-pair stable matching exists is NP-hard even with strict preferences. Bentert et al. [2023] studied some similar questions for multilayer preferences, where in each layer the preferences of the agents are approval based, i.e. their preference list is a single tie over some acceptable partners.

Chen et al. [2021] studied a notion of d-globally-robust stability. This means that the agents may change the position of an other agent in their preference list, by moving them up or down so that the number of such swaps in the preference profile are at most d. They have shown that finding a d-globally-robust stable matching can be done in polynomial time. They posed the open question of whether finding a d-robust stable matching, where each agent is allowed to make d swaps in their preference lists independently, is solvable in polynomial time. We answer this question positively.

Steindl and Zehavi ([2021; 2022]) studied multilayer preferences in the one-sided model. They also considered similar problems, e.g. problems where we aim to find matchings that satisfy many properties like Pareto-optimality in multiple layers.

1.2 Our Results

We consider an Uncertainty, a Multilayer and a Robust-popularity model for popular and dominant matchings.

In the Uncertainty model, each agent has a set of possible preference lists, and the goal is to find a matching that is popular for any possible choices of the preferences, which we call *certainly popular*. We show that deciding if there is such a matching is NP-hard with two-sided preferences, even if the preference lists on one side of the instance are fixed. If the preferences are one-sided, we show that the problem becomes solvable, and we give an algorithm that finds a certainly popular matching as well as a nice characterization of all certainly popular matchings.

In the Multilayer model, there are a set of layers of preference profiles and we aim to find matchings that are popular in all the layers, which again we call certainly popular matchings. We show that in the two-sided preference model, this is also NP-hard. On the positive side, we give an algorithm that finds a certainly popular matching in one-sided instances.

In the Robust popularity model, we know a preference list for each agent, but these may differ from their original preferences by at most k swaps, where a swap is just the switching of the order of two adjacent agents in a preference list. We show that finding a k-robust popular matching in such a model is NP-hard with two-sided preferences, but is solvable with one-sided ones.

Regarding dominant matchings, we show that in the Uncertainty model and the Robust popularity model, we can find a certainly (resp. k-robust) dominant matching in polynomial time, by utilizing the connection between stable and dominant matchings. For the Multilayer model, we show that the problem becomes NP-hard, even in restricted settings.

The results for the two-sided market are discussed in Section 3, while the one-sided market is considered in Section 4. Our results are summarized in Table 1.

2 Preliminaries

In this section, we describe the basics of two-sided and one-sided matching markets and our studied problems.

2.1 Two-sided Markets

In *two-sided markets*, we are given a bipartite graph $G = (A, B; E)$, where the vertices represent the agents and the edges represent the possible contracts or acceptability relations. Furthermore, for each agent $u \in A \cup B$, we are given

	Uncertainty	Multilayer	Robust
One-sided popular	P ([Thm 10])	P ([Thm 10])	P ([Thm **??**], Appendix)
Two-sided popular	NP-c ($\forall k \geq 2$, [Thm 5])	NP-c ($\forall k \geq 2$, [Thm 5])	NP-c ($\forall k \geq 1$, [Thm 6])
Dominant	P ([Thm 4])	NP-c ($\forall k \geq 2$, [Thm 2])	P ([Thm 4])

Table 1: A summary of our results. P means the problem is solvable in polynomial time, while NP-c denotes NP-complete.

a strict linear order (called a *preference list)* $\succ_u$ over $V(u)$, where $V(u)$ denotes the neighbors of u in G, i.e. the acceptable partners of u. We call the set of these preferences the *preference profile* of the instance and denote it by L, so $L = \{\succ_u \mid u \in A \cup B\}$.

We say that $M \subseteq E$ is a *matching*, if $|M(u)| \leq 1$ for all $u \in A \cup B$, where $M(u) \subseteq V(u)$ denotes the partner(s) of u in M. If u does not have any partners in M, then $M(u) = \emptyset$, which is considered strictly worse than any acceptable partner $v \in V(u)$. Depending on the context, we may also refer by $M(u)$ to the edge incident to u in M.

Stability. In two-sided matching markets, the most well-known optimality concept is stability. Let M be a matching. We say that an edge (a, b) *blocks* M, if $b \succ_a M(a)$ and $a \succ_b M(b)$ both hold. We say that M is *stable*, if no edge blocks M. A well-known result is that a stable matching always exists and can be found in linear time ([Gale and Shapley, 1962]).

Popularity. A relaxation of stability is popularity. To define the popularity of a matching, take two arbitrary matchings M and N. Then, we let each agent $u \in A \cup B$ cast a vote $\mathit{vote}^L_u(M, N)$ as follows. If $M(u) \succ_u N(u)$, then $\mathit{vote}^L_u(M, N) = +1$, which means that u prefers M to N, if $M(u) = N(u)$, then $\mathit{vote}^L_u(M, N) = 0$ and if $N(u) \succ_u M(u)$, then $\mathit{vote}^L_u(M, N) = -1$, meaning that u prefers N to M. We sum up these votes to obtain $\Delta^L(M, N) = \sum_{u \in A \cup B} \mathit{vote}^L_u(M, N)$.

We say that a matching M is *popular* (with respect to the preference profile L), if for any other matching N we have $\Delta^L(M, N) \geq 0$. In other words, a matching M is popular if it does not lose a head-to-head comparison with any other matching. Otherwise, if $\Delta^L(M, N) < 0$, we say that N *dominates* M.

Dominant matchings. We say that a matching M is *dominant*, if (i) M is popular and (ii) for any other matching N with $|N| > |M|$ it holds that $\Delta^L(M, N) > 0$, or equivalently $\Delta^L(N, M) < 0$. Dominant matchings form a nonempty subset of maximum size popular matchings, and can be found in polynomial time ([Cseh and Kavitha, 2018]). The motivation behind this definition is the following. A matching, where the same number of people are getting better and worse, but more people get matched may still be a more desirable solution and hence it would "dominate" the original matching in a weaker but still practical sense.

We will denote the problem of deciding whether an instance I of a two-sided matching market admits a popular matching by PM.

We study three different models. We start with describing the Multilayer model, where we assume that we are given multiple instances of PM, $I_1, I_2, \ldots, I_k$, with the same underlying bipartite graph, but with different preference profiles $L_1, L_2, \ldots, L_k$. Our goal is to decide whether there is a matching that is popular in all preference profiles (so $\Delta^{L_i}(M, N) \geq 0 \; \forall i \in [k]$ for all matchings N). We call such a matching a *certainly popular matching*. This model corresponds to the Multilayer preference model in [Chen *et al.*, 2018], defined for stability instead of popularity with layers $L_1, L_2, \ldots, L_k$, hence we call it POPULAR-MULTILAYER. It also corresponds to finding a certainly stable matching in the joint probability model in [Aziz *et al.*, 2016].

POPULAR-MULTILAYER

In: A number k, a bipartite graph $G = (A, B; E)$, and preference profiles $L_i = \{\succ^i_u \mid u \in A \cup B\}$ for $i \in [k]$.

Q: Is there a matching M that is certainly popular?

The problem DOMINANT-MULTILAYER is defined analogously, but instead of a certainly popular matching, we are looking for a certainly dominant matching, that is, a matching that is dominant in all of the preference profiles.

We proceed by defining the Uncertaintly model. Instead of one bipartite graph with multiple preference profiles, now we assume that for each agent u, there is a set of possible preference lists $P_u = \{\succ^1_u, \ldots, \succ^{k_u}_u\}$. This can be interpreted as each agent having preferences uncertain/unknown to us, such that we know some information about them, but not everything. In this model, our objective is to find a matching that is popular for any preference profile $L \subseteq \times_{u \in A \cup B} P_u$, that is, for any choices $\succ_u \in P_u$ for the individual preferences of each agent u. We call such a matching a *certainly popular matching*. This corresponds to the lottery model in [Aziz *et al.*, 2016] for stability.

POPULAR-UNCERTAIN

In: A number k and a bipartite graph $G = (A, B; E)$, with a set of preference lists P_u for each $u \in A \cup B$ with $|P_u| \leq k$.

Q: Is there a matching M that is certainly popular?

The problem DOMINANT-UNCERTAIN is defined analogously.

The next model is the Robust popularity model. Let $k \geq 1$. Here, we have a fixed $\succ_u$ preference list for each $u \in A \cup B$ and we know that agent u's real preference list differs from $\succ_u$ by at most k-swaps, where a swap is the switching if the order of two adjacent entries in the preference lists. For example, $a_2 \succ a_3 \succ a_1$ can be obtained from $a_1 \succ a_2 \succ a_3$, with two consecutive swaps (a_1, a_2) and (a_1, a_3). We say that agent v is *k-better (k-worse)* than v' for u, if the number of swaps required to make v worse (better) than v' is

k, respectively. We say that v is *at least k-better (worse)* than v', if the number of swaps required to make v worse (better) than v' is at least k. For two preference lists $\succ_u^1, \succ_u^2$, we let $d^s(\succ_u^1, \succ_u^2)$ be the minimum number of swaps required to reach $\succ_u^2$ from $\succ_u^1$. We say that a matching M is *k-robust popular*, if M is popular in any preference profile $L = \{\succ_u^L | u \in A \cup B\}$, where $d^s(\succ_u, \succ_u^L) \leq k$ for all $u \in A \cup B$.

POPULAR-ROBUST

In: A number k and a bipartite graph $G = (A, B; E)$ with preference lists $\succ_u$, $u \in A \cup B$.

Q: Is there a matching M that is k-robust popular?

We obtain the problem DOMINANT-ROBUST similarly, where we are looking for a k-robust dominant matching instead, that is, a matching which is dominant even if each agent may do at most k swaps in his preference list.

2.2 One-sided Markets

One-sided markets, or (capacitated) *House Allocation (*HA*) markets* as often called in the literature, are similar to two-sided markets. Here, also, we are given a bipartite graph $G = (A, B; E)$. However, now only the agents $a \in A$ have preference lists $\succ_a$ over their acceptable houses $V(a) \subseteq B$, so $L = \{\succ_a | a \in A\}$. We also allow the agents in B -who we call houses from now on- to have a positive integer capacity $\mathbf{q}[b]$.

In this model, we say that $M \subseteq E$ is a *matching*, if it holds that $|M(a)| \leq 1$ for any agents $a \in A$ and $|M(b)| \leq \mathbf{q}[b]$ for any house $b \in B$, where $M(a)$ and $M(b)$ denote the house(s) allocated to a and the agent(s) allocated to b in M, respectively. If $|M(b)| = \mathbf{q}[b]$, then we say that b is *filled (or saturated)*.

Popularity for the HA model is defined analogously, with the only difference being that only the agents in A participate in the voting procedure, so $\Delta^L(M, N)$ is defined as $\sum_{a \in A} vote_a^L(M, N)$. We say that a matching M is *popular*, if $\Delta^L(M, N) \geq 0$ for any matching N.

In the HA model, popular matchings may fail to exists, but it is possible to find a maximum size popular matching in polynomial time, if one exists [Abraham *et al.*, 2007].

The corresponding variants of POPULAR-MULTILAYER, POPULAR-UNCERTAIN and POPULAR-ROBUST for the one-sided case are called HA-POPULAR-MULTILAYER, HA-POPULAR-UNCERTAIN and HA-POPULAR-ROBUST, respectively. Their definition is analogous to the two-sided case; hence the precise definitions of these problems are deferred to the Appendix.

3 Two-sided Markets

In this section, we investigate the questions discussed for two-sided markets.

3.1 Verification

First of all, we show that deciding if a matching M is certainly popular or k-robust popular can be done in polynomial time, hence the problems are in NP.

Theorem 1. *Verifying if a matching M is certainly popular (resp. dominant) or k-robust-popular (resp. dominant) can be done in polynomial time in all of* POPULAR-MULTILAYER, DOMINANT-MULTILAYER, POPULAR-UNCERTAIN, DOMINANT-UNCERTAIN, POPULAR-ROBUST *and* DOMINANT-ROBUST.

Proof. The case of POPULAR-MULTILAYER and DOMINANT-MULTILAYER is trivial, because we only need to verify that M is popular (resp. dominant) in polynomially many preference profiles in the input size, and verifying whether a matching is popular or dominant can be done in polynomial time ([Huang and Kavitha, 2013; Cseh and Kavitha, 2018]).

We may assume without loss of generality that $|A| = |B|$. Otherwise, we can add isolated agents to the instance, whose vote will always be 0 when comparing any two matchings and the other agents' votes are unaffected.

Let I be an instance of POPULAR-UNCERTAIN or POPULAR-ROBUST with $G = (A, B, E)$ and let $\mathcal{L}$ denote the set of possible preference profiles (in the case of POPULAR-ROBUST, these are the profiles, where each preference list differs from the original by at most k swaps). Let M be a matching. First, we extend G to be a complete bipartite graph by adding all missing edges. Call this new graph H. Then, we create a weight function $w : E(H) \to \mathbb{Z}$ as follows. For each $e = (a, b) \in E(G)$, we let $w(e) = vote_a^*(M, e) + vote_b^*(M, e)$. Furthermore, $vote_a^*(M, e)$ is defined to be -1, if there is a possible preference list, where a prefers his partner in e, 0 if $M(a) = e$ and $+1$, if a always prefers his partner in M (so in the k-robust case, this corresponds to that partner being at least $(k+1)$-better). For an edge $e = (a, b) \notin E(G)$, we let $w(e) = vote_a^*(M, \emptyset) + vote_b^*(M, \emptyset)$ (here, $vote_a^*(M, \emptyset) = +1$, except the case when $M(a) = \emptyset$, when it is 0).

We claim that M is popular in all possible preference profiles if and only the minimum weight perfect matching in H has weight at least 0. Indeed, for any perfect matching N' of H, if we let $N = N' \cap E(G)$, then $w(N')$ is exactly $\min_{L \in \mathcal{L}} \Delta^L(M, N)$ by the definition of the weights and the fact that since N' is perfect, we count the smallest possible vote (i.e. it is -1, whenever it can be -1 in a possible preference list) for every agent in $A \cup B$ exactly once. Also, for any maximal matching $N \subseteq E(G)$, we can extend it to a perfect matching N' greedily (H is complete and $|A| = |B|$), using edges from $E(H) \setminus E(G)$ such that $w(N')$ is $\min_{L \in \mathcal{L}} \Delta^L(M, N)$. Hence, M is popular if and only if there is no perfect matching $N' \subseteq E(H)$ such that $w(N') < 0$. This can be verified in polynomial time with various methods, e.g. Egerváry's algorithm or Linear Programming.

For the case of DOMINANT-UNCERTAIN and DOMINANT-ROBUST, we can decide similarly if M is certainly popular or k-robust popular and if not, we can conclude that it is not certainly dominant or k-robust dominant, respectively. Therefore, we may assume that M is popular in any possible preference profile. To decide whether there is a larger matching N, such that $\Delta^L(M, N) = 0$ in a possible preference profile L, we subtract $\varepsilon = \frac{1}{|M|+1}$ from each original edge's

($e \in E(G)$) weight and check whether there is a perfect matching with weight at most -1. If there is a perfect matching N' of weight at most -1, then for $N = N' \cap E(G)$, we have that $|N| \geq |M|+1$, as M is popular. Also, $|N| \leq 2|M|$, because any popular matching is maximal. Therefore, the weight of N' cannot be at least 1 with respect to the original weights (so without subtracting the ε values), so it is exactly 0. Therefore, N is larger than M, but $\Delta^L(M,N) = 0$ in a possible preference profile L, as $w(N') = 0$, so M is not certainly dominant or k-robust dominant, respectively. In the other direction, if M is not dominant in all possible preference profiles, then there is a matching N with $|N| \geq |M|+1$ and $\Delta^L(M,N) = 0$ in a possible preference profile L, so N can be extended to a perfect matching N' with weight at most -1. □

3.2 Dominant Matchings

We proceed to our results about dominant matchings in two-sided markets. Our first result is the hardness of DOMINANT-MULTILAYER. The proof is deferred to the Appendix.

Theorem 2. DOMINANT-MULTILAYER *is NP-complete, even for a fixed constant $k = 2$.*

Before the other results, we describe a useful reduction that connects dominant and stable matchings. Let $I = (G, L)$, $G = (A, B, E)$ be an instance of PM. We create an instance I', by adding two parallel edges $x(e), y(e)$ for each edge $e \in E$. Then, we create the preference profile in I' (over the edges now, as we have parallel edges in the new instance) from L, by the following rule. For each agent $a \in A$, we first rank all the edges of type x incident to it, in the same order as a's preference over his neighbors, then we rank all incident edges of type y, again in the order of a's preference list. So, if a had preference list $b_1 \succ b_2$, then it becomes $x((a,b_1)) \succ x((a,b_2)) \succ y((a,b_1)) \succ y((a,b_2))$. For the agents $b \in B$, we do it similarly, by interchanging the order of the x and y types. The following theorem connects the dominant matchings in I with the stable matchings in I'.

Theorem 3 ([Kavitha, 2014], [Kamiyama, 2020])**.** *Every stable matching M' in I' gives a dominant matching $M = f(M')$ in I, where $f(M') = \{e \mid |M' \cap \{x(e), y(e)\}| \geq 1\}$. Furthermore, for every dominant matching M in I, there is a matching M' in I' such that $M = f(M')$.*

The only problem with Theorem 3 is that for a dominant matching M, the stable matching M' with $f(M') = M$ also depends on the preference profile in I. Hence, finding a certainly dominant matching does not reduce trivially to finding a certainly stable matching, as for a certainly dominant matching, its preimage in I' may be different for each preference profile, so there may not be a certainly stable matching in I' that gives M.

We use this connection to the stable matching problem to solve DOMINANT-UNCERTAIN and DOMINANT-ROBUST. First, we show that DOMINANT-ROBUST reduces in polynomial time to DOMINANT-UNCERTAIN.

Let I be an instance of DOMINANT-ROBUST. We create an instance I' of DOMINANT-UNCERTAIN with the same underlying bipartite graph. Then, we create the set of preference lists for each agent. Let $P_u = \{\succ_u^v \mid v \in V(u)\}$, where $\succ_u^v$ is the preference list obtained from $\succ_u$ by swapping v down k times in the preference list (a swap down of the worst entry in the preference list is just considered as doing nothing). For example, if $\succ_u = v_1 \succ v_2 \succ v_3 \succ v_4$ and $k = 2$, then $\succ_u^{v_1}$ is $v_2 \succ v_3 \succ v_1 \succ v_4$.

Clearly, for a fixed matching M, it is k-robust dominant in DOMINANT-ROBUST if and only if it is dominant even if every agent swaps down his partner in M (if there is any) by k places. Also, for each agent u and each possible partner v of u, we have created a preference list $\succ_u^v$, where any agent that can be better than v with at most k swaps is better than v at the same time. Hence, $\min_{\succ_u' \mid d(\succ_u, \succ_u') \leq k}\{vote_{\succ_u'}(M,N)\} = \min_{v \in V(u)}\{vote_{\succ_u^v}(M,N)\}$ holds for each $u \in A \cup B$ and matchings M, N, so a matching M is certainly dominant in I if and only if it is certainly dominant in I'.

Remark 1. *We note that with the same idea, we get that we can reduce the problem of finding a k-robust stable matching, that is, a matching which remains stable even if each agent can apply k swaps in his preference list, to finding a certainly stable matching in polynomial time. As the latter problem can be solved in polynomial time (see [Aziz* et al.*, 2016]), we get that this problem also admits a polynomial algorithm, answering an open question of [Chen* et al.*, 2021].*

We are ready for our next result.

Theorem 4. DOMINANT-UNCERTAIN *and* DOMINANT-ROBUST *can be solved in polynomial time.*

Proof. We have seen that it is enough to give an algorithm for DOMINANT-UNCERTAIN.

Let I be an instance of DOMINANT-UNCERTAIN. We create an instance I' of STABLE-UNCERTAIN, where for each edge e, we create two parallel edges $x(e)$ and $y(e)$ as in Theorem 3, and for each agent u, his possible preference lists are created from his original preference lists $\succ_u^i$, such that if $u \in A$, then we first rank the x type edges according to $\succ_u^i$ and then the y type edges in the same order, and if $u \in B$, then we rank the y types first and the x types last.

By Theorem 3, we see that if there is a certainly stable matching in I', then that gives us a certainly dominant matching M in I. Therefore, our goal is to show that if there is a certainly dominant matching M in I, then there is a certainly stable matching M' in I'. For this, let M be certainly dominant. We first create a strict preference list $\succ_u'$ for each $u \in A \cup B$ in I, such that for any $v \in V(u)$, if $v \succ_u^i M(u)$ for a possible preference list $\succ_u^i$ of u, then $v \succ_u' M(u)$, and if $M(u) \succ_u^i v$ for all possible preference lists, then $M(u) \succ_u' v$. Otherwise, the preferences between agents that are both better or both worse than $M(u)$ are created arbitrarily. As M is certainly dominant, it is also dominant with the preferences $\succ_u'$, $u \in A \cup B$, because $\min_{i \in k}\{vote_{\succ_u^i}(M,N)\} = vote_{\succ_u'}(M,N)$ for any matching N.

By Theorem 3, this matching M gives a stable matching M' in I' with respect to the extensions of the $\succ_u'$ preferences, which we also denote $\succ_u'$ by abuse of notation. We claim that it is certainly stable. Suppose that an edge $x(a,b)$ blocks M' with the extension of some preference lists of a and b. Then, if a is not matched by an x type edge, then a prefers this edge in the extension of $\succ'_a$ too. Also, b must be unmatched or matched with an x-type edge ($x(a,b)$ cannot dominate y type edges for b) and in the latter case, there must be a preference list $\succ_b^i$, such that $a \succ_b^i M(b)$. However, then we have that $x(a,b) \succ'_b M'(b)$, which contradicts that M' is stable with these preferences. If a is matched with an x-type edge, then the fact that $x(a,b)$ blocks implies that $b \succ_a^i M(a)$ for some i and so $x(a,b) \succ'_a M'(a)$. Similarly as before, we get that $x(a,b) \succ'_b M'(b)$ also holds and therefore M' is not stable, a contradiction. The case for edges of type y is analogous.

Hence, we conclude that M' is certainly stable, and therefore the theorem follows from the fact that a certainly stable matching can be found in polynomial time ([Aziz *et al.*, 2016]). □

3.3 Popular Matchings

We conclude the section with our results for popular matchings. Due to space constraints, their proofs are deferred to the Appendix.

Theorem 5. POPULAR-MULTILAYER *and* POPULAR-UNCERTAIN *are NP-complete, even if uncertainty occurs only on one side and $k = 2$ is a fixed constant.*

Theorem 6. POPULAR-ROBUST *is NP-complete, even for a fixed constant $k = 1$.*

4 One-sided Markets

In this section, we discuss our results for the HA model.

First, we describe a characterization of popular matchings in HA instances from [Manlove and Sng, 2006]. Let $(G, \mathbf{q}, L)$ be an instance with preference profile L. We can add a "last-resort" house $\hat{b}$ to B with capacity $|A|$ that is appended to the end of each applicant's preference list. Then, the set of popular matchings in this new instance are in one-to-one correspondence with the original ones: we just say that each unmatched agent a is now at an imaginary worst house, the "last-resort". Hence, from now on we may assume that each instance of HA-POPULAR-MULTILAYER and HA-POPULAR-UNCERTAIN contains such a house. Very importantly though, the new inputs are always such that this house must always be last for every agent in all possible preference lists or preference profiles in the input. This is without loss of generality, since this house did not exist originally.

Next, for each applicant $a \in A$, define the *first house* $f_L(a)$ to be the house that is the most preferred by a. For a house b, we call an agent a an *admirer of* b in L, if $b = f_L(a)$. Denote the set of admirers of b by $AD_L(b)$. Next, we define the *pseudo-second house* $s_L(a)$ to be either $f_L(a)$, whenever $b = f_L(a)$ has at most as many admirers as its capacity $\mathbf{q}[b]$, or otherwise as the house most preferred by a among the houses that have strictly less admirers than their capacity. Note that because of the last resort house, $s_L(a)$ exists.

Make a graph $G_L = (A, B, E_L)$ by adding the edges $(a, f_L(a))$ and $(a, s_L(a))$ for each applicant $a \in A$. In the case where $f_L(a) = s_L(a)$ we only add one of the parallel edges.

Theorem 7 ([Manlove and Sng, 2006])**.** *A matching M is popular if and only if the following conditions hold:*

- *Every edge of M is from E_L,*
- *M is A-perfect, so it matches every $a \in A$ and*
- *every house $b \in B$ that can be filled with admirers of b only is filled with admirers of b only.*

4.1 Multilayer and Uncertainty Models

Let us consider an instance I of HA-POPULAR-MULTILAYER or HA-POPULAR-UNCERTAIN.

For a house b, let $CA(b)$ denote the set of *certain admirers* of b, that is, the agents $a \in A$ that consider b first in any possible preference profile/preference list. Let $PA(b)$ denote the set of *possible admirers* of b, that is, the agents $a \in A$ that consider b best in some possible preference profiles/preference lists. Finally, let $H_a(b)$ be the set of houses that are better than b in some possible preference lists of a. All these sets can be easily computed in polynomial time in both the HA-POPULAR-MULTILAYER and the HA-POPULAR-UNCERTAIN problems.

Using Theorem 7 we get the following result.

Theorem 8. *Let $G = (A, B; E)$ be a bipartite graph, with $\boldsymbol{q}[b]$ capacities for each $b \in B$. Let $\mathcal{L} = \{L_1, L_2, \dots, L_k\}$ be a set of preference profiles for A. Then, a matching M is popular in any $L \in \mathcal{L}$ if and only if*

- *M is an A-perfect matching in the graph $\hat{G} = (A, B, \hat{E})$, where $\hat{E} = \cap_{i=1}^n E_{L_i}$ and*
- *for any $L \in \mathcal{L}$, any house b for which $|AD_L(b)| \geq \boldsymbol{q}[b]$ is filled by agents from $AD_L(b)$.*

For the uncertainty model, we prove a different characterization in the appendix that we are going to utilize.

Theorem 9. *Let $G = (A, B; E)$ be a bipartite graph with $\boldsymbol{q}[b]$ capacities for each $b \in B$. Let $P_{a_1}, \dots, P_{a_n}$ be the sets of preference lists for $a_i \in A$. Then, a matching M is popular in any $L \in \times_{i=1}^n P_{a_i}$ if and only if*

- *M is an A-perfect matching in the graph $\hat{G} = (A, B, \hat{E})$, where $\hat{E} = \cap_{i=1}^n E_{L_i}$,*
- *any house b that has at least $\boldsymbol{q}[b]$ possible admirers is filled with possible admirers,*
- *any house b that has at least $\boldsymbol{q}[b]$ certain admirers is filled with certain admirers,*
- *for each agent a and each possible pseudo-second house b of a, if $(a, b) \in M$, then each $b' \in H_a(b)$ is filled with certain admirers.*

Using Theorems 8 and 9, we show that HA-POPULAR-MULTILAYER and HA-POPULAR-UNCERTAIN are polynomial-time solvable. It also follows that the verification problems are solvable too in both cases.

Theorem 10. HA-POPULAR-MULTILAYER *and* HA-POPULAR-UNCERTAIN *are solvable in polynomial time.*

Proof. Let I be an instance of either HA-POPULAR-MULTILAYER or HA-POPULAR-UNCERTAIN. Let $\mathcal{L}$ denote the set of possible preference profiles, which in the case of HA-POPULAR-UNCERTAIN is $\mathcal{L} = \times_{a \in A} P_a$, which can have exponentially many profiles in the input size, whereas in the case of HA-POPULAR-MULTILAYER only polynomially many (since each of them is part of the input).

Let us start with the case of HA-POPULAR-MULTILAYER. By Theorem 8, it is enough to decide whether $\hat{G} = (A, B, \hat{E})$, $\hat{E} = \cap_{L \in \mathcal{L}} E_L$ admits an A-perfect matching or not, satisfying that for any $L \in \mathcal{L}$, any house b for which $|AD_L(b)| \geq \mathbf{q}[b]$ is filled by agents from $AD_L(b)$. We can compute $AD_L(b)$ for each $(b, L) \in B \times \mathcal{L}$ in polynomial time, and we can also compute the graph $\hat{G}$ in polynomial time. Now, for each edge $(a, b) \in \hat{E}$, we check whether there is a preference profile L, such that $|AD_L(b)| \geq \mathbf{q}[b]$, but $a \notin AD_L(b)$, and if so, then we delete the edge (a, b). If there remains an A-perfect matching M in the graph, then it clearly satisfies the first condition and also the second, as otherwise if $(a, b) \in M$, but a is not an admirer of b in L and $|AD_L(b)| \geq \mathbf{q}[b]$, then the edge (a, b) has been deleted, contradiction. For the other direction, if there is no A-perfect matching remaining, then there can be no matching that is popular in every $L \in \mathcal{L}$, because such a matching would be A-perfect and none of its edges would have been deleted.

Next, consider the case of HA-POPULAR-UNCERTAIN. We start by computing the sets $CA(b), PA(b)$ for each house $b \in B$ and the sets $H_a(b)$ for each edge $(a, b) \in E$.

First, we show that we can decide whether $e \in \hat{E}$ or not, so we can construct $\hat{G}$ in polynomial time.

By Theorem 7, an edge $e = (a, b)$ is included in $\hat{G}$ if and only if $b \in \{s_L(a), f_L(a)\}$ for each possible preference profile $L \in \mathcal{L}$. If b is first in all of a's possible preference lists, then $(a, b) \in \hat{E}$ clearly. If not, then $(a, b) \in \hat{E}$, if and only if for each preference list $\succ_a^i$, where b is not first, each house $b' \succ_a^i b$ for a must have at least $\mathbf{q}[b']$ many admirers in any possible preference profile containing $\succ_a^i$. This means that $|CA(b')| \geq \mathbf{q}[b']$. Using the sets $CA(b)$, we can check for a given $a \in A$ and $b \in B$, whether the conditions for $(a, b) \in \hat{E}$ hold for each $\succ_a^i \in P_a$ preference list of a and, therefore, we can construct $\hat{G}$ in polynomial time.

Then, by point two of Theorem 9, we know that the edges (a, b) such that $|PA(b)| \geq \mathbf{q}[b]$, but $a \notin PA(b)$ cannot be included in any matching that is certainly popular, so we can safely delete them from $\hat{G}$. Similarly, we can delete the edges (a, b), such that $|CA(b)| \geq \mathbf{q}[b]$, but $a \notin CA(b)$ by point three of Theorem 9.

Finally, for each agent $a \in A$, and each remaining (a, b) edge, if b is not a first house of a in some preference list, then we check whether for each $b' \in H_a(b)$, $|CA(b')| \geq \mathbf{q}[b']$, which can be done in polynomial time. If not, then b' cannot be filled with certain admirers and hence a cannot be at b in a certainly popular matching by Theorem 9 point four, so we can delete (a, b) from the graph. Otherwise, for each such b', we already deleted each edge (a', b') when $a' \notin CA(b')$, so whenever a matching in the remaining graph fills b', it must fill it with certain admirers.

Therefore, we conclude that there is a certainly popular matching if and only if the remaining graph has a matching that is (i) A-perfect and (ii) fills every $b \in B$ such that $|PA(b)| \geq \mathbf{q}[b]$ by Theorem 9. Such a matching, if exists, can be found efficiently with, for example, the Hungarian-method. □

Our results regarding the Robust popularity model are deferred to the appendix.

5 Conclusions and Open Questions

In this paper, we considered models where agents' preferences may vary or may be uncertain for popular matching problems both in one-sided and two-sided preferences. We gave a comprehensive overview of the complexity of our studied problems.

For future work, it may be interesting to consider similar questions for other well-known notions. We also left open the complexity of DOMINANT-MULTILAYER and POPULAR-ROBUST if only one side has uncertain preferences.

Acknowledgements

Gergely Csáji acknowledges the financial support by the Hungarian Academy of Sciences, Momentum Grant No. LP2021-1/2021, by the Hungarian Scientific Research Fund, OTKA, Grant No. K143858 and by the Ministry of Culture and Innovation of Hungary from the National Research, Development and Innovation fund, financed under the KDP-2023 funding scheme (grant number C2258525).

References

[Abraham *et al.*, 2007] David J Abraham, Robert W Irving, Telikepalli Kavitha, and Kurt Mehlhorn. Popular matchings. *SIAM Journal on Computing*, 37(4):1030–1045, 2007.

[Aziz *et al.*, 2016] Haris Aziz, Péter Biró, Serge Gaspers, Ronald De Haan, Nicholas Mattei, and Baharak Rastegari. Stable matching with uncertain linear preferences. In *Algorithmic Game Theory: 9th International Symposium, SAGT 2016, Liverpool, UK, September 19–21, 2016, Proceedings 9*, pages 195–206. Springer, 2016.

[Aziz *et al.*, 2019] Haris Aziz, Péter Biró, Ronald de Haan, and Baharak Rastegari. Pareto optimal allocation under uncertain preferences: uncertainty models, algorithms, and complexity. *Artificial Intelligence*, 276:57–78, 2019.

[Bentert *et al.*, 2023] Matthias Bentert, Niclas Boehmer, Klaus Heeger, and Tomohiro Koana. Stable matching with multilayer approval preferences: approvals can be harder than strict preferences. *Games and Economic Behavior*, 2023.

[Chen *et al.*, 2018] Jiehua Chen, Rolf Niedermeier, and Piotr Skowron. Stable marriage with multi-modal preferences. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, pages 269–286, 2018.

[Chen *et al.*, 2021] Jiehua Chen, Piotr Skowron, and Manuel Sorge. Matchings under preferences: Strength of stability and tradeoffs. *ACM Transactions on Economics and Computation*, 9(4):1–55, 2021.

[Cseh and Kavitha, 2018] Ágnes Cseh and Telikepalli Kavitha. Popular edges and dominant matchings. *Mathematical Programming*, 172(1-2):209–229, 2018.

[Cseh *et al.*, 2017] Ágnes Cseh, Chien-Chung Huang, and Telikepalli Kavitha. Popular matchings with two-sided preferences and one-sided ties. *SIAM Journal on Discrete Mathematics*, 31(4):2348–2377, 2017.

[Gale and Shapley, 1962] David Gale and Lloyd S Shapley. College admissions and the stability of marriage. *American Mathematical Monthly*, 69(1):9–15, 1962.

[Gärdenfors, 1975] Peter Gärdenfors. Match making: assignments based on bilateral preferences. *Behavioral Science*, 20(3):166–173, 1975.

[Huang and Kavitha, 2013] Chien-Chung Huang and Telikepalli Kavitha. Popular matchings in the stable marriage problem. *Information and Computation*, 222:180–194, 2013.

[Kamiyama, 2020] Naoyuki Kamiyama. Popular matchings with two-sided preference lists and matroid constraints. *Theoretical Computer Science*, 809:265–276, 2020.

[Kavitha, 2014] Telikepalli Kavitha. A size-popularity tradeoff in the stable marriage problem. *SIAM Journal on Computing*, 43(1):52–71, 2014.

[Manlove and Sng, 2006] David F Manlove and Colin TS Sng. Popular matchings in the capacitated house allocation problem. In *European Symposium on Algorithms*, pages 492–503. Springer, 2006.

[Manlove, 2013] David Manlove. *Algorithmics of matching under preferences*, volume 2. World Scientific, 2013.

[Miyazaki and Okamoto, 2017] Shuichi Miyazaki and Kazuya Okamoto. Jointly stable matchings. *Leibniz International Proceedings in Informatics, LIPIcs*, 92:56–1, 2017.

[Steindl and Zehavi, 2021] Barak Steindl and Meirav Zehavi. Parameterized analysis of assignment under multiple preferences. In *Multi-Agent Systems: 18th European Conference, EUMAS 2021, Virtual Event, June 28–29, 2021, Revised Selected Papers 18*, pages 160–177. Springer, 2021.

[Steindl and Zehavi, 2022] Barak Steindl and Meirav Zehavi. Verification of multi-layered assignment problems. *Autonomous Agents and Multi-Agent Systems*, 36(1):15, 2022.