



Two-Stage Stochastic Stable Matching

Yuri Faenza^(✉), Ayoub Foussoul^(✉), and Chengyue He^(✉)

IEOR, Columbia University, New York, USA
 {yf2414, af3209, ch3480}@columbia.edu

Abstract. We introduce and study a two-stage stochastic stable matching problem between students and schools. A decision maker chooses a stable matching in a marriage instance; then, after some agents enter or leave the market following a probability distribution $\mathcal{D}$, chooses a stable matching in the new instance. The goal is, roughly speaking, to maximize the expected quality of the matchings across the two stages and minimize the expected students' discontent for being downgraded to a less preferred school in the second-stage. We consider both the case when $\mathcal{D}$ is given explicitly and when it is accessed via a sampling oracle. In the former case, we give a polynomial time algorithm. In the latter case, we show that, unless $P = NP$, no algorithm can find the optimal value or the optimal solution of the problem in polynomial-time. On the positive side, we give a pseudopolynomial algorithm that computes a solution of arbitrarily small additive error. Our techniques include the use of a newly defined poset of stable pairs, which may be of independent interest.

Keywords: Stable Matching · Two-stage stochastic programming · Poset

1 Introduction

Stability is a fundamental concept in matching markets problems when we do not only wish to optimize a global objective function, but we also care that the output solution is fair at the level of individual agents. Since its introduction in the seminal work by Gale and Shapley [7], stability has been employed in many real-world applications, including matching medical residents to hospitals, assigning students to schools, and matching organ donors to recipients. An instance $(A, B, \succ)$ of the marriage model (see, e.g., [8, 12]) consists of a two-sided market of students A and schools B , and for every $a \in A$ (resp., $b \in B$), a strict order $\succ_a$ (resp., $\succ_b$) over $B^+ = B \cup \{\emptyset\}$ (resp., $A^+ = A \cup \{\emptyset\}$), where $\emptyset$ denotes the outside option. An element $v \in A \cup B$ is called an *agent*. In the stable matching problem, the goal is to find a matching¹ M that is *stable*: a matching with no *blocking pairs* or *blocking agents*. More precisely, let $M(v) \in A \cup B \cup \{\emptyset\}$

¹ Throughout the paper, a matching in a marriage instance $(A, B, \succ)$ refers to a subset of pairs of $A^+ \times B^+$ where every agent $v \in A \cup B$ is paired with either a unique agent from the other side or with the outside option $\emptyset$.

$a^1: \boxed{b^1} \succ_{a^1} b^2 \succ_{a^1} b^3 \succ_{a^1} \underline{b^4} \succ_{a^1} b^5 \succ_{a^1} \emptyset$	$b^1: a^4 \succ_{b^1} a^5 \succ_{b^1} a^3 \succ_{b^1} a^2 \succ_{b^1} a^1 \succ_{b^1} \emptyset$
$a^2: \boxed{b^2} \succ_{a^2} b^1 \succ_{a^2} b^4 \succ_{a^2} \underline{b^3} \succ_{a^2} b^5 \succ_{a^2} \emptyset$	$b^2: a^3 \succ_{b^2} a^5 \succ_{b^2} a^4 \succ_{b^2} a^1 \succ_{b^2} a^2 \succ_{b^2} \emptyset$
$a^3: \boxed{b^3} \succ_{a^3} b^4 \succ_{a^3} b^1 \succ_{a^3} \underline{b^2} \succ_{a^3} b^5 \succ_{a^3} \emptyset$	$b^3: a^2 \succ_{b^3} a^5 \succ_{b^3} a^1 \succ_{b^3} a^4 \succ_{b^3} a^3 \succ_{b^3} \emptyset$
$a^4: \boxed{b^4} \succ_{a^4} b^3 \succ_{a^4} b^2 \succ_{a^4} \underline{b^1} \succ_{a^4} b^5 \succ_{a^4} \emptyset$	$b^4: a^1 \succ_{b^4} a^5 \succ_{b^4} a^2 \succ_{b^4} a^3 \succ_{b^4} a^4 \succ_{b^4} \emptyset$
$a^5: \boxed{b^5} \succ_{a^5} b^3 \succ_{a^5} b^2 \succ_{a^5} b^1 \succ_{a^5} b^4 \succ_{a^5} \emptyset$	$b^5: a^5 \succ_{b^5} a^1 \succ_{b^5} a^2 \succ_{b^5} a^3 \succ_{b^5} a^4 \succ_{b^5} \emptyset$

Fig. 1. Example adapted from [6, Example 8.1]. Consider the first-stage instance I_1 given above, where $a^1, \dots, a^5$ are the students. Then $M_0 = \{a^1b^1, a^2b^2, a^3b^3, a^4b^4, a^5b^5\}$ and $\underline{M} = \{a^1b^4, a^2b^3, a^3b^2, a^4b^1, a^5b^5\}$ are two stable matchings of I_1 , with M_0 being student-optimal (the partner of each student a in M_0 and $\underline{M}$ is, respectively, boxed and underlined in a 's preference list). If b^5 leaves the market in the second stage, the only stable matching is $M' = \underline{M} \setminus \{a^5b^5\}$, hence, $\underline{M}$ minimizes the number of students downgraded in the second stage.

denote the partner of an agent $v \in A \cup B$ in M , then a pair $ab \in A \times B$ is called a blocking pair if both a and b prefer each other to their partner in M , that is, $b \succ_a M(a)$ and $a \succ_b M(b)$, and an agent $v \in A \cup B$ is called a blocking agent if v prefers the outside option to its partner in M , that is, $\emptyset \succ_v M(v)$.

While the classical stable matching problem assumes static and fully known input, in many applications the input changes over time as new agents enter or leave the market. For instance, each year, and after an assignment of the students to San Francisco public middle schools has been decided, around 20% of the students who were allotted a seat choose not to use it, mostly to join private schools instead [3]. Also, new students move to the city and need to be allotted a seat. On the other side of the market, the schools may have unforeseen budget cuts or expansions, leading to a change in the number of available seats.

In all of the above scenarios, a second-stage reallocation of at least part of the seats is required. For instance, while a subset of students might commit to their assigned seats in the first round of matching (and hence leave the market along with their allotted seats), the original assignment needs to be adjusted for the rest of the students after new agents arrive (or old agents depart) in order to maintain stability. It is therefore important that the initial stable matching, besides being of good quality, is adaptable to a changing environment with small adjustments so that it leads to a good quality second round matching and to small dissatisfaction of the students present in both rounds from being downgraded to a less preferred school. It is easy to find examples (see Fig. 1) where, in order to achieve such goal, one might want to go beyond the student optimal stable matching², which is the obvious solution that is almost always employed (see, e.g., [1, 11]).

Motivated by these considerations, in this paper we introduce and study a two-stage stochastic stable matching problem, where agents enter and leave the market between the two stages. The goal of the decision maker is, roughly

² The student optimal matching is a stable matching that all students prefer to every other stable matching.

speaking, to maximize the expected quality of the matchings across the two stages and minimize the expected students' discontent from being downgraded to a less preferred school when going from the first to the second stage. Below, we make these ideas formal.

1.1 The Model

An instance of the two-stage stochastic stable matching problem is as follows.

First-Stage Instance: a marriage instance $I_1 = (A, B, \succ)$ of agents present in the first stage.

Second-Stage Instance: a distribution $\mathcal{D}$ over subsets of $A \cup B$, from which the agents present in the second stage are sampled³.

Cost: cost functions $c_1, c_2 : A^+ \times B^+ \rightarrow \mathbb{Q}$ and a penalty coefficient $\lambda \in \mathbb{Q}_+$. Let S denote the set of agents sampled from $\mathcal{D}$ in the second-stage, and $I_2 = (A \cap S, B \cap S, \succ)$ denote the corresponding second-stage instance⁴. Given a first- (resp., second-) stage matching M_1 (resp., M_2), we measure the quality of M_1 (resp., M_2) by $c(M_1) = \sum_{ab \in M_1} c_1(ab)$ (resp., $c(M_2) = \sum_{ab \in M_2} c_2(ab)$), and measure the dissatisfaction of students for being moved to a less preferred school between M_1 and M_2 by

$$d_S(M_1, M_2) = \lambda \sum_{a \in A \cap S} [R_a(M_2(a)) - R_a(M_1(a))]^+. \quad (1)$$

Here, $R_a : B^+ \rightarrow \mathbb{N}$ is the *rank function* of a , that is, $R_a(b) = i$ iff b is the i -th most preferred choice of a (among the schools B and the outside position) and $[x]^+ = \max\{0, x\}$. The coefficient λ is the per unit of rank change dissatisfaction of a student from switching to a school of higher rank between M_1 and M_2 . This is a natural measure of dissatisfaction where students are unhappy to be downgraded to a less preferred school in the second-stage, such that the worse the school is, the more unsatisfied they are.

Objective Function: The goal is to solve the two-stage stochastic problem

$$\min_{M_1 \in \mathcal{M}_{I_1}} c_1(M_1) + \mathbb{E}_{S \sim \mathcal{D}} \left[\min_{M_2 \in \mathcal{M}_{I_2}} c_2(M_2) + d_S(M_1, M_2) \right], \quad (2\text{STO})$$

where $I_1 = (A, B, \succ)$ is the first-stage marriage instance, S is the set of agents of the second-stage sampled from $\mathcal{D}$, $I_2 = (A \cap S, B \cap S, \succ)$ is the second-stage instance, and $\mathcal{M}_I$ denotes the set of stable matchings of a marriage instance I .

In (2STO), we therefore wish to select a first-stage matching M_1 such that the cost of M_1 plus the expected cost that we have to pay in the second-stage

³ This setting subsumes the more general setting where agents enter and leave the market (see Remark 1).

⁴ Given subsets $A' \subset A$ and $B' \subset B$, we use $(A', B', \succ)$ to denote the marriage instance $(A', B', \succ')$ where $\succ'$ is the restriction of the collection of orders $\succ$ to agents $A' \cup B'$.

is minimized. The second-stage cost is given by the cost of the second-stage matching plus the total dissatisfaction of the students for being downgraded to a less preferred school between the first- and second-stage. We note that once a first-stage matching M_1 is fixed and for every fixed second-stage scenario S , the second-stage problem

$$\min_{M_2 \in \mathcal{M}_{I_2}} c_2(M_2) + d_S(M_1, M_2)$$

is a minimum weight stable matching problem and hence can be solved efficiently (see, e.g., [8]).

Remark 1 (Generalizations). We remark that our model captures the more general setting where agents can also enter the market in the second stage. In particular, for every entering student a in an instance $\mathcal{I}$, one can consider an instance $\mathcal{I}'$ with an extra dummy school b' such that a and b' are the most preferred partners of each other. Hence, a, b' are matched to each other in each stable matching of the first stage. Then a entering the market in the second stage is equivalent to b' leaving the market. It is easy to see that c_2 can be adjusted so as to discount the dissatisfaction of a in the new model and obtain the original objective function. An entering school is modelled analogously.

Also, our algorithm and analysis extend seamlessly to capture the following generalization of the dissatisfaction function (1): given $\lambda \in \mathbb{Q}_+$, and a set of non-negative scores for the pairs $w : A^+ \times B^+ \rightarrow \mathbb{Q}_+$, with the property that $w(ab) \geq w(ab')$ if $b \succ_a b'$, define

$$d_S(M_1, M_2) = \lambda \sum_{a \in A \cap S} [w(aM_2(a)) - w(aM_1(a))]^+.$$

This more general version of the dissatisfaction function captures, for example, the setting when the per unit of rank change dissatisfaction from changing school is student-dependent (in which case $w(ab) = \alpha_a \cdot R_a(b)$ for some set of non-negative weights $\{\alpha_a\}_{a \in A}$) or where a student does not care if the change happens between, say, their top 5 schools (in which case $w(ab_1) = w(ab_2) = \dots = w(ab_5)$, where $b_1, \dots, b_5$ are the top 5 schools in a 's list) but cares if it happens between one of their top 5 schools and the rest. For simplicity of exposition, we restrict the presentation to the unweighted case, and briefly discuss the changes needed for the general case when we present our algorithm (Sect. 2.2).

1.2 Our Results and Techniques

Explicit Second-Stage Distribution. Assume first that the second-stage distribution is given explicitly by the list of all possible second-stage scenarios Θ and their respective probabilities of occurrence $\{p_S\}_{S \in \Theta}$. Let $\mathbf{N}^{\text{exp}}(\mathcal{I})$ denote the input size⁵ of instance $\mathcal{I}$ of (2STO) in this setting. Our first main result is a polynomial time algorithm for (2STO) in this input model.

⁵ For $x \in \mathbb{Q}$, let $\|x\|$ denote the encoding length of parameter x . Then $\mathbf{N}^{\text{exp}}(\mathcal{I}) = \sum_{S \in \Theta} \|p_S\| + \sum_{ab} (\|c_1(ab)\| + \|c_2(ab)\|) + \|\lambda\| + |A \cup B|^2$.

Theorem 1. *There exists an algorithm such that given any instance $\mathcal{I}$ of (2STO) where the second-stage distribution is given explicitly solves the the problem in time polynomial in the input size $N^{\text{exp}}(\mathcal{I})$.*

To prove Theorem 1, we show that the two-stage problem (2STO) can be polynomially reduced to a minimum $s - t$ cut problem in a carefully constructed directed graph. In our construction, we introduce and use a new poset defined over stable pairs of a marriage instance, that is, the student-school pairs that are contained in some stable matching. We call this poset the *poset of stable pairs* (see Sect. 2.1). This poset might be of independent interest in approaching stable matching optimization problems⁶.

Implicit Second-Stage Distribution. Next, we consider the more general model where the second-stage distribution $\mathcal{D}$ is given by a sampling oracle. Let $N^{\text{imp}}(\mathcal{I})$ denote the input size⁷ of instance $\mathcal{I}$ of (2STO) in this setting. We show the following hardness result for the two-stage problem (2STO) under this more general model.

Theorem 2. *Unless $P=NP$, there exists no algorithm that given any instance $\mathcal{I}$ of (2STO) where the second-stage distribution is specified implicitly by a sampling oracle, solves⁸ the problem in time and number of calls to the sampling oracle that is polynomial in the input size $N^{\text{imp}}(\mathcal{I})$. This hardness result holds even if the cost parameters λ and $\{c_1(ab)\}_{ab}, \{c_2(ab)\}_{ab}$ are in $\{0, 1\}$.*

Our proof of Theorem 2 relies on a reduction from the problem of counting the number of vertex covers of an undirected graph, which is #P-Hard [9].

On the positive side, we give an arbitrary good additive approximation to (2STO) when the second-stage distribution is specified implicitly. Our algorithm runs in time pseudopolynomial in the input size. In particular, let $M_1^{\mathcal{I}}$ be the optimal solution of an instance $\mathcal{I}$ of (2STO) and let $\text{val}_{\mathcal{I}}(M_1)$ denote the objective value of a first-stage stable matching M_1 . We show the following result.

Theorem 3. *There exists an algorithm that, given an instance $\mathcal{I}$ of (2STO) where the second-stage distribution is specified implicitly by a sampling oracle, and two parameters $\epsilon > 0$ and $\alpha \in (0, 1)$, gives a first-stage stable matching solution M_1 such that,*

$$\mathbb{P}(\text{val}_{\mathcal{I}}(M_1) \leq \text{val}_{\mathcal{I}}(M_1^{\mathcal{I}}) + \epsilon) \geq 1 - \alpha.$$

⁶ Our construction could also be done starting from the poset of rotations [8]; however, our poset (defined over pairs) gives more intuitive constructions and a characterization of stable families of pairs as antichains of the poset, leading to natural reformulations for other problems (e.g., the maximum cardinality of a family of pairs that is stable in two marriage instances can be reformulated using our poset as the maximum independent set in the union of two perfect graphs, studied in [4]).

⁷ $N^{\text{imp}}(\mathcal{I}) = \sum_{ab} (\|c_1(ab)\| + \|c_2(ab)\|) + \|\lambda\| + |A \cup B|^2$.

⁸ In the sense that it gives at least one of the two, the optimal value or the optimal solution. Note that the hardness of finding the optimal value does not exclude the possibility of finding the optimal solution in polynomial time.

The algorithm runs in time polynomial in $N^{\text{imp}}(\mathcal{I})$, $\max_{ab} |c_2(ab)|$, λ , $1/\epsilon$ and $\log(1/\alpha)$.

The algorithm from Theorem 3 runs in polynomial time when λ, c_2 are polynomially bounded (recall that, already in this case, by Theorem 2, no polynomial-time algorithm exists unless $P = NP$). Our algorithm employs the widely used Sample Average Approximation method (e.g. [5, 10, 13]) to approximate an instance of (2STO) with an implicitly specified second-stage distribution by an instance of explicitly specified second-stage distribution. Then we show that a relatively small number of samples (calls to the sampling oracle) is enough to get a good additive approximation with high probability. Whether there exists and FPRAS for (2STO) in this model is an interesting open question.

Numerical Experiments. We illustrate our method on randomly generated instances. We compare the performance of five first-stage matchings: The optimal solution of our problem denoted by M^* , the one-side optimal matchings, the matching minimizing $c_1(\cdot)$, and finally the best stable matching in hindsight denoted by M^{off} , which is the first-stage stable matching that would have been optimal to pick if one knew the realization of the second-stage distribution. Our results suggest that for a wide range of values of the penalty coefficient λ , our solution M^* is a strictly better choice to pick in the first stage and that it has a performance that is not very far from the performance of the best matching in hindsight M^{off} .

Note. Due to space limitation, all proofs are deferred to the full version of the paper.

1.3 Notation and Definitions

Consider a marriage instance $I = (A, B, \succ)$. We say that a pair $ab \in A \times B$ (resp., family of pairs) is *stable* if it is contained in some stable matching of I . An agent is stable if it is not matched to the outside option in some stable matching of I ⁹. Let $\mathcal{S}(I)$ and $\mathcal{F}(I)$ denote the set of stable pairs and stable families of I , respectively. For a student a , let $\mathcal{S}_a(I)$ denote the set of stable pairs containing student a . Let $A^{\text{st}}(I)$ and $B^{\text{st}}(I)$ denote the set of stable students and schools of I respectively. The collection of orders $(\succ_a)_{a \in A}$ induce a natural partial order $>_A$ over the set of pairs $A \times B^+$, such that $ab >_A a'b'$ iff $a = a'$ and $b \succ_a b'$. Given a non-empty finite set Y and a total order $>$ over Y , we denote by $\max_{>} Y$ (resp., $\min_{>} Y$) the maximum (resp., minimum) of Y with respect to $>$.

2 Explicit Second-Stage Distribution

We consider in this section an instance of (2STO) given by a first-stage marriage instance $(A, B, \succ)$, a second-stage distribution $\mathcal{D}$ given explicitly by the list of all

⁹ We recall that if an agent is not matched to the outside option in a stable matching then they will not be matched to the outside option in all stable matchings (see, e.g., [8]).

possible second-stage scenarios Θ and their probabilities of occurrence $\{p_S\}_{S \in \Theta}$, and finally costs c_1, c_2 , and λ . The two-stage stochastic problem (2STO) can be written as follows,

$$\min_{M_1 \in \mathcal{M}_{I_1}} c_1(M_1) + \sum_{S \in \Theta} p_S \left(\min_{M_2^S \in \mathcal{M}_{I_2^S}} c_2(M_2^S) + d_S(M_1, M_2^S) \right),$$

or equivalently,

$$\min_{\substack{M_1 \in \mathcal{M}_{I_1} \\ \{M_2^S \in \mathcal{M}_{I_2^S}\}_{S \in \Theta}}} c_1(M_1) + \sum_{S \in \Theta} p_S c_2(M_2^S) + p_S d_S(M_1, M_2^S). \quad (\text{EXP-2STO})$$

where $I_2^S = (A \cap S, B \cap S, \succ)$ is the second-stage instance under scenario S . We now give a polynomial time algorithm to solve (EXP-2STO). Our algorithm uses a new partial order that we introduce next.

2.1 Poset of Stable Pairs

In this section, we introduce the *poset of stable pairs*. This is a partial order that we define over the set of stable pairs, and which characterizes the stable families of a marriage instance.

Definition 1. (*Poset of stable pairs*) Let $I = (A, B, \succ)$ be a marriage instance. The poset of stable pairs of I denoted by $(\mathcal{S}(I), >)$ is the set of stable pairs $\mathcal{S}(I)$ over which we define the partial order $>$ as follows: for every $ab, a'b' \in \mathcal{S}(I)$, we have $a'b' < ab$ if and only if for every stable matching M such that $ab \in M$ it holds that $M(a') \succ_a b'$. We write $a'b' \leq ab$ if either $ab = a'b'$ or $a'b' < ab$.

In particular, a stable pair $ab \in \mathcal{S}(I)$ is greater than a stable pair $a'b' \in \mathcal{S}(I)$ if whenever the larger pair ab appears in some stable matching M , the student of the smaller pair (i.e., a') is matched in M to a partner they strictly prefer to b' . Note that in particular, $ab' < ab$ when $ab' <_A ab$ or equivalently $b' \prec_a b$. In the reminder of the paper, we use the order $<_A$ when comparing (non-necessarily stable) pairs involving the same student, we use the order $<$ when comparing stable pairs (potentially involving different students), and we use the order $\prec_v$ when comparing agents (and the outside option) from the opposite side of v . The next lemma shows that $(\mathcal{S}(I), >)$ is, as claimed in Definition 1, a poset, and can be constructed efficiently.

Lemma 1. Let I be a marriage instance. Then, $(\mathcal{S}(I), >)$ is a partially ordered set and can be constructed in time polynomial in the number of agents.

The following lemma gives a characterization of the stable families of a marriage instance by the means of its poset of stable pairs.

Lemma 2. Let I be a marriage instance. Then $F \in \mathcal{F}(I)$ if and only if F is an antichain of $(\mathcal{S}(I), >)$.

We now leverage Lemma 2 to construct a capacitated directed graph with vertices $\mathcal{S}(I)$ in which cuts with finite capacity are in a one-to-one correspondence with the maximal antichains of $(\mathcal{S}(I), >)$, hence with stable matchings. We begin with the definition of the *Next-to-Smallest Dominating Stable Pair*.

Definition 2. (*Next-to-Smallest Dominating Stable Pair*) Let $I = (A, B, \succ)$ be a marriage instance. Let $ab \in \mathcal{S}(I)$, $a' \in A^{\text{st}}(I)$ with $a' \neq a$, and $\mathcal{S}_{a'}(I) = \{a'b'_1, a'b'_2, \dots, a'b'_{n_{a'}}\}$ with $b'_1 \prec_{a'} \dots \prec_{a'} b'_{n_{a'}}$. Suppose ab is dominated by an element of $\mathcal{S}_{a'}(I)$, that is, there exists $i \in [n_{a'}]$ such that $ab < a'b'_i$, and let $i^* = \min\{i \in [n_{a'}] \mid ab < a'b'_i\}$. We define the Next-to-Smallest Dominating Stable Pair of ab in $\mathcal{S}_{a'}(I)$ as $\sigma(ab, a') = a'b'_{i^*-1}$.

In order for $\sigma(ab, a')$ to be well-defined in the above definition, b'_{i^*-1} needs to be well-defined (i.e., $i^* \geq 2$). The following lemma shows this is always the case.

Lemma 3. Let $I = (A, B, \succ)$ be a marriage instance. Let $ab \in \mathcal{S}(I)$ and $a' \in A^{\text{st}}(I)$ with $a' \neq a$. Let $\mathcal{S}_{a'}(I) = \{a'b'_1, a'b'_2, \dots, a'b'_{n_{a'}}\}$ with $b'_1 \prec_{a'} \dots \prec_{a'} b'_{n_{a'}}$. If ab is dominated by an element of $\mathcal{S}_{a'}(I)$ then $\min\{i \in [n_{a'}] \mid ab < a'b'_i\} \geq 2$.

We are now ready to define the *Cut Graph* of a marriage instance. This is a capacitated direct graph whose finite cuts are in a one-to-one correspondence with the stable matchings of the instance.

Definition 3. (*Cut Graph*) Consider a marriage instance $I = (A, B, \succ)$. The Cut Graph of I is a capacitated directed graph denoted by $G_{\text{cut}}(I)$. The vertices of $G_{\text{cut}}(I)$ are the stable pairs $\mathcal{S}(I)$, a source vertex s and a sink vertex t . The arcs of $G_{\text{cut}}(I)$ are as follows: (i) for every $a \in A^{\text{st}}(I)$, let $\mathcal{S}_a(I) = \{ab_1, ab_2, \dots, ab_{n_a}\}$ with $b_1 \prec_a \dots \prec_a b_{n_a}$. Add an arc from ab_i to ab_{i+1} for every $i \in [n_a - 1]$ and an arc from s to ab_{n_a} . (ii) for every $ab \in \mathcal{S}(I)$ and $a' \in A^{\text{st}}(I)$ such that $a' \neq a$, if ab is dominated by an element of $\mathcal{S}_{a'}(I)$, add an arc from ab to $\sigma(ab, a')$. All arcs have infinite capacity.

Note that t is isolated in the cut graph. The following lemma shows that there is a one-to-one correspondence between the $s - t$ cuts of finite capacity of $G_{\text{cut}}(I)$ and the stable matchings of I .

Lemma 4. Let $I = (A, B, \succ)$ be a marriage instance. The following is a bijection between $s - t$ cuts of finite capacity in $G_{\text{cut}}(I)$ and stable matchings in I :

- (a) Let $(C, \overline{C})$ be a finite capacity $s - t$ cut of $G_{\text{cut}}(I)$. Then, for every $a \in A^{\text{st}}(I)$, $\mathcal{S}_a(I) \cap C \neq \emptyset$. Moreover, the matching M where every agent $a \in A^{\text{st}}(I)$ is matched to its least preferred pair in C , i.e., the pair $\min_{<_A} \mathcal{S}_a(I) \cap C$, and every other non-stable agent $v \notin A^{\text{st}}(I) \cup B^{\text{st}}(I)$ is matched to the outside option is a stable matching.
- (b) Let M be a stable matching of I . Let

$$C = \{s\} \cup \bigcup_{a \in A^{\text{st}}(I)} \{ab \in \mathcal{S}_a(I) \mid aM(a) \leq_A ab\},$$

which is the union of $\{s\}$ and the set of all stable pairs that are greater than $aM(a)$ for every stable agent $a \in A^{\text{st}}(I)$. Then $(C, \overline{C})$ is an $s - t$ cut of finite capacity of $G_{\text{cut}}(I)$.

2.2 Algorithm

We now present our algorithm for (EXP-2STO). First of all, note that adding a positive constant $c \in \mathbb{Q}_+$ to the cost functions c_1 and c_2 changes the objective value by exactly $c|A^{\text{st}}(I_1)| + c \sum_{S \in \Theta} p_S |A^{\text{st}}(I_2^S)|$ which is a constant and does not depend on the chosen solution $M_1, \{M_2^S\}_{S \in \Theta}$. Hence, without loss of generality, we suppose that the cost functions c_1 and c_2 are non-negative. From now on, when necessary, we add a superscript indicating the instance at hand when referring to an order. For example, we use $>^{I_1}$, to refer to the poset of stable pairs of the first-stage instance I_1 . We construct a capacitated directed graph $\mathcal{N}(\mathcal{I})$ and show that (EXP-2STO) is equivalent to the minimum $s-t$ cut problem over $\mathcal{N}(\mathcal{I})$.

Consider the cut graphs $G_{\text{cut}}(I_1)$ and $\{G_{\text{cut}}(I_2^S)\}_{S \in \Theta}$ constructed following Definition 3 for the first-stage instance I_1 and second-stage instances $\{I_2^S\}_{S \in \Theta}$ respectively. In order to distinguish the vertices of these graphs, we use the notation $[v]$ to denote vertex v when it belongs to $G_{\text{cut}}(I_1)$ and use $[v]_S$ to denote v when it belongs to $G_{\text{cut}}(I_2^S)$. For example, $[s]$ denotes the vertex s of $G_{\text{cut}}(I_1)$ and $[s]_S$ denotes the vertex s of $G_{\text{cut}}(I_2^S)$.

High-Level Idea. Our directed graph $\mathcal{N}(\mathcal{I})$ is constructed in three steps.

- (i) *Bijection between finite capacity cuts and solutions* $(M_1, \{M_2^S\}_{S \in \Theta})$. We consider the cut graphs $G_{\text{cut}}(I_1)$ and $\{G_{\text{cut}}(I_2^S)\}_{S \in \Theta}$ to which we add a super source vertex s (resp., super sink vertex t) that we link to $[s]$, $\{[s]\}_{S \in \Theta}$ (resp., $[t]$, $\{[t]_S\}_{S \in \Theta}$) with infinite capacity arcs in such a way that any $s-t$ cut of finite capacity always includes (resp., precludes) the vertices $[s]$, $\{[s]\}_{S \in \Theta}$ (resp., $[t]$, $\{[t]_S\}_{S \in \Theta}$). This ensures that any $s-t$ cut of finite capacity of $\mathcal{N}(\mathcal{I})$ corresponds to a solution $(M_1, \{M_2^S\}_{S \in \Theta})$, where M_1 (resp., M_2^S) is a first- (resp., second-) stage stable matching. We next add arcs of finite positive capacity to capture the cost of the solution $(M_1, \{M_2^S\}_{S \in \Theta})$.
- (ii) *Cost* $\lambda \cdot \sum_{S \in \Theta} p_S d_S(M_1, M_2^S)$. This is the most complex part of the construction, and the one for which the new poset comes in handy. It is achieved by adding edges between graphs $G_{\text{cut}}(I_1)$ and $G_{\text{cut}}(I_2^S)$ for all $S \in \Theta$. In particular, for every $S \in \Theta$ we add edges between the vertices of $G_{\text{cut}}(I_1)$ and those of $G_{\text{cut}}(I_2^S)$ such that an $s-t$ cut of finite capacity is traversed by a subset of these edges of total capacity $\lambda p_S d_S(M_1, M_2^S)$, where M_1, M_2^S are the stable matchings corresponding to the cut. To do so, for every $a \in A \cap S$, we would like to add an arc of capacity λp_S from the copy of a pair ab in $G_{\text{cut}}(I_2^S)$ to the copy of the pair ab in $G_{\text{cut}}(I_1)$ for every b such that $b \succeq_a \emptyset$. This would ensure that an $s-t$ cut of finite capacity that takes from $\mathcal{S}_a(I_1)$ the pairs that a weakly prefers to ab_1 and from $\mathcal{S}_a(I_2^S)$ the pairs that weakly prefers to ab_2 (in which case a will be matched to b_1 and b_2 in the corresponding first and second stage stable matchings respectively) is traversed by exactly $[R_a(b_2) - R_a(b_1)]^+$ many pairs of capacity λp_S . However, ab is not necessarily stable in I_1 (resp., I_2^S) and hence $[ab]$

- (ii) Consider a second-stage scenario $S \in \Theta$ and let $a \in A \cap S$ be an agent present both in the first- and second-stage. Let $b \in B^+$ such that $b \succeq_a \emptyset$ and let $\Pi(ab) = \{ab' \in \mathcal{S}_a(I_1) \mid ab' \leq_A ab\}$ and $\Pi_S(ab) = \{ab' \in \mathcal{S}_a(I_2^S) \mid ab' \leq_A ab\}$ define,

$$\Psi(ab) = \begin{cases} [s] & \text{if } a \notin A^{\text{st}}(I_1) \\ [t] & \text{else if } \Pi(ab) = \emptyset, \\ [\max_{<_A} \Pi(ab)] & \text{otherwise} \end{cases} \quad \Psi_S(ab) = \begin{cases} [s]_S & \text{if } a \notin A^{\text{st}}(I_2^S) \\ [t]_S & \text{else if } \Pi_S(ab) = \emptyset; \\ [\max_{<_A} \Pi_S(ab)]_S & \text{otherwise} \end{cases}$$

add an arc of capacity λp_S ¹⁰ from $\Psi_S(ab)$ to $\Psi(ab)$.

- (iii) Let $a \in A^{\text{st}}(I_1)$ and let $\mathcal{S}_a(I_1) = \{ab_1, ab_2, \dots, ab_{n_a}\}$ with $b_1 \prec_a \dots \prec_a b_{n_a}$. Add an arc from $[s]$ to $[ab_i]$ of capacity $c_1(ab_{i+1})$ for every $i \in [1, n_a - 1]$. Add an arc from $[ab_i]$ to $[t]$ of capacity $c_1(ab_i)$ for every $i \in [1, n_a]$. Similarly, for every $S \in \Theta$, let $a \in A^{\text{st}}(I_2^S)$ and let $\mathcal{S}_a(I_2^S) = \{ab_1, ab_2, \dots, ab_{n_a}\}$ with $b_1 \prec_a \dots \prec_a b_{n_a}$, add an arc from $[s]_S$ to $[ab_i]_S$ of capacity $p_S c_2(ab_{i+1})$ for every $i \in [1, n_a - 1]$ and an arc from $[ab_i]_S$ to $[t]_S$ of capacity $p_S c_2(ab_i)$ for every $i \in [1, n_a]$.

The following lemma shows that solving (EXP-2STO) is equivalent to solving the minimum $s - t$ cut problem over $\mathcal{N}(\mathcal{I})$. Our first main result Theorem 1 follows immediately from Lemma 5.

Lemma 5. *Let $(C, \bar{C})$ be a minimum $s - t$ cut in $\mathcal{N}(\mathcal{I})$. Let $[C]$ (resp., $[C]_S$) denote the pairs corresponding to the vertices of C from the first-stage (resp., scenario S of the second-stage) cut graph. Let M_1 be the matching where every stable agent in the first stage $a \in A^{\text{st}}(I_1)$ is matched to its least preferred stable pair in $[C]$, i.e., the pair $\min_{<_A} \mathcal{S}_a(I_1) \cap [C]$ (which exists as $\mathcal{S}_a(I_1) \cap [C] \neq \emptyset$ by Lemma 4), and every other non-stable agent $v \notin A^{\text{st}}(I_1) \cup B^{\text{st}}(I_1)$ is matched to the outside option. Similarly, for every scenario $S \in \Theta$, let M_2^S be the matching where every stable agent in the second stage $a \in A^{\text{st}}(I_2^S)$ is matched to its least preferred stable pair in $[C]_S$, i.e., the pair $\min_{<_A} \mathcal{S}_a(I_2^S) \cap [C]_S$ (which exists as $\mathcal{S}_a(I_2^S) \cap [C]_S \neq \emptyset$ by Lemma 4), and every other non-stable agent $v \notin A^{\text{st}}(I_2^S) \cup B^{\text{st}}(I_2^S)$ is matched to the outside option. Then M_1 and $\{M_2^S\}_{S \in \Theta}$ are stable matchings and an optimal solution of (EXP-2STO).*

3 Sampling Oracle Model

We consider in this section the more general model where the second-stage distribution $\mathcal{D}$ of an instance $\mathcal{I}$ is given by a sampling oracle. We show an hardness result for the problem under this more general input model and give an arbitrary good additive approximation.

¹⁰ For the generalized version of the dissatisfaction function in Remark 1, if b is not the most preferred school of a , let $b' = \min_{\succ_a} \{b'' \mid b'' \succ_a b\}$, then use here capacity $(w(ab) - w(ab')) \cdot \lambda p_S \geq 0$ instead.

3.1 Computational Complexity

Theorem 2 follows from Lemma 6 below and the fact that the counting problem of vertex covers is #P-Hard [9], hence a polynomial time algorithm for this problem would imply that #P=FP, which in turns implies that P=NP (see, e.g., [2]).

Lemma 6. *Suppose that there exists an algorithm $\mathcal{A}$ such that given any instance $\mathcal{I}$ of (2STO), algorithm $\mathcal{A}$ solves the two-stage problem over $\mathcal{I}$ (i.e., finds either the optimal solution or the optimal value) in time and number of calls to the sampling oracle that is polynomial in the input size $N^{\text{imp}}(\mathcal{I})$. Then $\mathcal{A}$ can be used to count the number of vertex covers in any undirected graph $G(V, E)$ in time polynomial in $|V|$.*

3.2 Algorithm

We give a near-optimal additive approximation to the two-stage problem (2STO). Our algorithm is sampling based and runs in time pseudopolynomial in the input.

Our starting point is to approximate the expected value in the two-stage stochastic problem (2STO) by a sample average over N i.i.d. samples. This method, known in the stochastic programming literature as the *Sample Average Approximation* (SAA) method has been widely used to approximate two-stage stochastic combinatorial problems [5, 10, 13]. In particular, let $S^1, \dots, S^N$ be N i.i.d. samples drawn from the distribution $\mathcal{D}$. We replace the expected value in the two-stage stochastic problem (2STO) with the sample average taken over $S^1, \dots, S^N$. We get the following sample average minimization problem,

$$\min_{M_1 \in \mathcal{M}_{I_1}} c_1(M_1) + \frac{1}{N} \sum_{j=1}^N \min_{M_2^j \in \mathcal{M}_{I_2^j}} c_2(M_2^j) + d_{S^j}(M_1, M_2^j). \quad (\text{SAA})$$

Note that (SAA) is an instance of (EXP-2STO) where $\Theta = \{S^1, \dots, S^N\}$ and $p_S = \frac{1}{N}$ for all $S \in \Theta$, and can therefore be solved exactly using our algorithm for explicitly specified distributions, in time polynomial in N and the other input parameters. It remains to show that, for any $\epsilon > 0$, a small number of samples N is enough to get an ϵ additive approximation of (SAA) with high probability.

Let $\hat{M}_1$ be an optimal solution of the sample average problem (SAA). The following lemma bounds the (additive) quality of the first-stage stable matching $\hat{M}_1$ in the original problem (SAA) as a function of the number of samples N .

Lemma 7. *Let $\alpha \in (0, 1)$. Then with probability at least $1 - \alpha$ it holds that,*

$$\text{val}_{\mathcal{I}}(\hat{M}_1) \leq \text{val}_{\mathcal{I}}(M_1^{\mathcal{I}}) + |A|(\max_{ab} |c_2(ab)| + \lambda|B|) \sqrt{\frac{\max\{|A|, |B|\} \log(3.88/\alpha)}{N}}.$$

Our third result Theorem 3 follows immediately from Lemma 7. In particular, for any $\epsilon > 0$, a number of samples

$$N = \left(|A|(\max_{ab} |c_2(ab)| + \lambda |B|) \right)^2 \frac{\max\{|A|, |B|\} \log(3.88/\alpha)}{\epsilon^2}$$

is enough to get an ϵ additive approximation with probability at least $1 - \alpha$.

4 Numerical Experiments

Experimental Setup. We consider an instance of (2STO) with $n = 50$ students and schools of uniformly randomly generated preferences. We suppose that each agent leaves the market in the second-stage with probability $p = 0.25$ independently of the other agents. In terms of costs, we set the cost of a pair ab , $c_1(ab)$ and $c_2(ab)$, to be the average rank between the two sides¹¹, and solve the problem for different values of the penalty coefficient λ . We compare the performance of five first-stage matchings: the optimal matching computed by our algorithm denoted by M^* , the one-side optimal matchings M_0 and M_z that are the common practice that is usually used, the matching $\hat{M}$ maximizing c_1 , and the optimal matching in hindsight M^{off} , that is, the first-stage stable matching that would have been optimal to take if one knew the realization of the second-stage distribution.




Fig. 3. Simulation results on an instance with 50 students and 50 schools.

Results. The results are given in Fig. 3. We observe that for a wide range of λ , the optimal stable matching is different from M_0 , M_z , and $\hat{M}$. Moreover, the optimal value $\text{val}(M^*)$ is relatively close to the optimal value in hindsight $\text{val}(M^{\text{off}})$. Observe that for large values of λ (in which case we are more in favor of minimizing the students downgrades in the second-stage than the quality of

¹¹ This is known as an *egalitarian* stable matching (see, e.g., [8]), since it balances between the utilities of the two sides of the market.

the matching), the optimal matching becomes M_z . Our results suggest that for a wide range of λ , it is better to choose M^* over choosing M_0 , M_z , or $\hat{M}$ in the first-stage. They also suggest that M^* will not be a lot worse than what we could have done in hindsight.

Choosing λ . The coefficient λ measures how much we favor the minimization of changes between the two stages to the maximization of the quality of the matchings. In practice, the appropriate λ is problem-dependent and a good coefficient λ can be learned from a few applications of the method. In particular, given some measure of performance (e.g. surveys where students/faculty express how satisfied they are with the matching process), one can choose, for example, λ that maximizes the performance on average.

Acknowledgments. The authors thank the IPCO referees for their useful comments. Yuri Faenza and Chengyue He acknowledge the support of the NSF Grant 2046146 *CAREER: An Algorithmic Theory of Matching Markets*.

References

1. Abdulkadiroğlu, A., Pathak, P.A., Roth, A.E.: The New York City high school match. *Am. Econ. Rev.* **95**(2), 364–367 (2005)
2. Arora, S., Barak, B.: *Computational Complexity: A Modern Approach*. Cambridge University Press, Cambridge (2009)
3. Ashlagi, I., Graur, A., Lo, I., Mentzer, K.: Overbooking with priority-respecting reassignment. In: *Presentation at the 3rd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization (EAAMO’23)* (2023)
4. Chakaravarthy, V.T., Pandit, V., Roy, S., Sabharwal, Y.: Finding independent sets in unions of perfect graphs. In: *IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2010)*, pp. 251–259 (2010)
5. Charikar, M., Chekuri, C., Pál, M.: Sampling bounds for stochastic optimization. In: *International Workshop on Approximation Algorithms for Combinatorial Optimization*, pp. 257–269 (2005)
6. Faenza, Y., Zhang, X.: Legal assignments and fast EADAM with consent via classic theory of stable matchings. *Oper. Res.* **70**(3), 1873–1890 (2022)
7. Gale, D., Shapley, L.S.: College admissions and the stability of marriage. *Am. Math. Mon.* **69**(1), 9–15 (1962)
8. Gusfield, D., Irving, R.W.: *The Stable Marriage Problem: Structure and Algorithms*. MIT Press, Cambridge (1989)
9. Provan, J.S., Ball, M.O.: The complexity of counting cuts and of computing the probability that a graph is connected. *SIAM J. Comput.* **12**(4), 777–788 (1983)
10. Ravi, R., Sinha, A.: Hedging uncertainty: approximation algorithms for stochastic optimization problems. *Math. Program.* **108**, 97–114 (2006)
11. Roth, A.E., Peranson, E.: The redesign of the matching market for American physicians: some engineering aspects of economic design. *Am. Econ. Rev.* **89**(4), 748–780 (1999)
12. Roth, A.E., Sotomayor, M.: Two-sided matching. *Handb. Game Theory Econ. Appl.* **1**, 485–541 (1992)
13. Swamy, C., Shmoys, D.B.: Sampling-based approximation algorithms for multi-stage stochastic optimization. In: *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2005)*, pp. 357–366 (2005)