

UCLA

UCLA Previously Published Works

Title

Learning strategies in decentralized matching markets under uncertain preferences

Permalink

<https://escholarship.org/uc/item/8gn3f4gw>

Authors

Dai, X

Jordan, MI

Publication Date

2021

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

Learning Strategies in Decentralized Matching Markets under Uncertain Preferences

Xiaowu Dai

Department of Economics

University of California, Berkeley, CA 94720-1776, USA

XWDAI@BERKELEY.EDU

Michael I. Jordan

Division of Computer Science and Department of Statistics

University of California, Berkeley, CA 94720-1776, USA

JORDAN@CS.BERKELEY.EDU

Editor: Amos Storkey

Abstract

We study the problem of decision-making in the setting of a scarcity of shared resources when the preferences of agents are unknown a priori and must be learned from data. Taking the two-sided matching market as a running example, we focus on the decentralized setting, where agents do not share their learned preferences with a central authority. Our approach is based on the representation of preferences in a reproducing kernel Hilbert space, and a learning algorithm for preferences that accounts for uncertainty due to the competition among the agents in the market. Under regularity conditions, we show that our estimator of preferences converges at a minimax optimal rate. Given this result, we derive optimal strategies that maximize agents' expected payoffs and we calibrate the uncertain state by taking opportunity costs into account. We also derive an incentive-compatibility property and show that the outcome from the learned strategies has a stability property. Finally, we prove a fairness property that asserts that there exists no justified envy according to the learned strategies.

Keywords: Uncertain preferences, Stability, Fairness, Matching markets, Reproducing kernel Hilbert spaces

1. Introduction

Many real-world decision-making problems can be viewed from both an economic point of view and a statistical point of view. The economic point of view focuses on scarcity of shared resources and the need to coordinate among multiple decision-makers. Thus, decision-makers must assess preferences over outcomes and those preferences need to interact in determining an overall set of outcomes. The statistical point of view recognizes that preferences are often not known a priori, but must be learned from data; moreover, agents' decisions are often influenced by latent state variables whose values must be inferred in order to determine a preferred outcome. Unfortunately, it is uncommon that these two perspectives are brought together in the literature, with economic work rarely addressing the need to learn preferences from data, and statistical machine learning rarely addressing scarcity and its consequences for decision-making. In this paper we aim to bridge this gap, studying a core microeconomic problem—two-sided matching markets—in the setting in which preferences must be learned from data. Moreover, we focus on *decentralized*

matching markets, reflecting several desiderata that are common in the machine-learning literature—that agents are autonomous and private, and that scalability and avoidance of central bottlenecks is a principal concern of an overall system design.

Two-sided matching markets have played an important role in microeconomics for several decades (Roth and Sotomayor, 1990), both as a theoretical topic and as a mainstay of real-world applications. Matching markets are used to allocate indivisible “goods” to multiple decision-making agents based on mutual compatibility as assessed via sets of preferences. Matching markets are often organized in a decentralized way such that each agent makes their decision independently of others’ decisions. Examples include college admissions, decentralized labor markets, and online dating. Matching markets embody a notion of scarcity in which the resources on both sides of the market are limited. The *congestion* is a key issue in decentralized decision-making under scarcity, as participants may not be able to make enough offers and acceptances to clear the market (Roth and Xing, 1997). The uncertainty of participants’ preferences is ubiquitous in real-world decentralized matching markets. For instance, college admissions in the United States face applicants’ uncertain preferences. The admitted students of a college may receive offers from other colleges. Students need to accept one or reject all offers, often within a *short* period. The process provides little opportunity for the college to learn students’ preferences, which depends on colleges’ competition and colleges’ uncertain popularity in the current year. Consequently, the college may end up enrolling too many or too few students relative to its capacity (Avery et al., 2003). It is in the colleges’ interest to decide which applicants to admit, such that the entering class will meet reasonably close to their quotas and be close to the attainable optimum in quality (Gale and Shapley, 1962).

When traditional machine learning methods are applied to problems involving scarcity and competition, they can create problems involving congestion, poor allocations, and lack of fairness (Busoniu et al., 2008; Finocchiaro et al., 2021). These problems need to be addressed by mechanisms that go beyond mere load-balancing or post-processing of the outputs of a pattern recognition system. Indeed, there are interactions between matching and learning that must be addressed jointly. Examples include: (1) The effects of scarcity imply that preferences are non-stationary from the perspective of any individual agent. (2) Agents’ decisions are often influenced by latent state variables whose values must be inferred in practice. This presents calibration challenges that need to take into account capacity limits and opportunity costs. (3) Learning methods should be asked to produce outcomes that are not merely optimal in a pointwise sense, but stable in a game-theoretic sense. Thus, lack of stability needs to inform the learning process.

In this work, we present a statistical framework that explicitly models a decision-making process under uncertainty and scarcity and yields learned strategies that only use local information. We use the classical college admissions market as our running example. In the proposed model, there are a set of agents (for example, colleges), each with limited capacity, and a set of arms (for example, students), each of which can be matched to at most one agent. Agents value two attributes of an arm: a “score” (for example, SAT/ACT score) that is common to all agents and a “fit” (for example, college-specific essay) that is agent-specific and is independent across agents. According to their score and fit, agents rank arms, but they do not observe arms’ preferences, which have no restriction. The model incorporates the arm’s uncertain preference into an acceptance probability, depending on both the

unknown state and agents’ competition. We want to learn arms’ acceptance probabilities from the historical data of arms’ binary choices, where arms’ attributes may vary over time. Various statistical learning algorithms allow efficient learning of the acceptance probability under the proposed model. To fix ideas, we present the penalized log-likelihood method in the reproducing kernel Hilbert space (RKHS) to learn the acceptance probability. Under regularity conditions the estimator converges at a minimax optimal rate. We focus on the single-stage decentralized matching that involves a single stage interaction: agents simultaneously pull sets of arms (for example, colleges offer admissions to students). Each arm accepts one of the agents (if any) that pulled it. We propose an efficient algorithm called calibrated decentralized matching (CDM) for maximizing the agents’ expected payoff. We calibrate the unknown state by perturbing the state and balancing the marginal utility and the marginal penalty. The proposed calibration procedure takes the opportunity costs into account. The CDM can perform the calibration in both the average-case and worst-case scenarios, depending on maximizing the averaged or minimal expected payoff concerning the unknown true state.

The CDM algorithm can serve as a recommendation engine in decentralized platform for decision-making. Indeed, our results add a learning component to help participants decide which participants on the other side of the market are the best to connect to. In particular, even with an unknown state, agents can estimate the probability of successfully pulling an arm using historical data. The prediction of match compatibility is also possible in another direction that arms can learn how much an agent may prefer them. The CDM procedure eases congestion in practice via recommendation and pulling more arms compared to alternative methods. Moreover, we show that asymptotically CDM makes it safe for agents to act straightforwardly on their preferences. That is, CDM ensures incentive-compatibility for agents. We also show that CDM yields a stable outcome for the market. Our notion of stability is similar to that of Liu et al. (2014), which is designed to study decentralized matching with incomplete information. This stability notion extends the classical notion of stability due to Gale and Shapley (1962) that assumes complete information of participants’ preferences. We prove that CDM is asymptotically stable. Moreover, we show that CDM is asymptotically fair for arms, in the sense of no justified envy in the matching process according to CDM.

Beyond the setting of matching markets, there are many other real-world machine learning problems that involve scarcity and competition, and in which decentralized methods of the kind that we explore here are appropriate. For example, decision-making problems in manufacturing, information gathering, and load balancing can often be treated as decentralized problems, where time and capacity are scarce and agents need to build up beliefs about the outcome of their actions from noisy observations (cf. Ooi and Wornell, 1996; Cogill et al., 2006; Seuken and Zilberstein, 2008). Related problems, such as the analysis of social dilemmas (Leibo et al., 2017) and competitive games (Tampuu et al., 2017) all operate in a multi-agent domain in which agents compete for scarce resources. Our model framework serves as a step towards explicitly modeling a decision-making process under uncertainty and scarcity and developing decentralized algorithms for learning from data.

1.1 Related Work

We briefly review some of the related work from multiple literatures, including matching markets, mechanism design, and multi-agent learning.

Matching Markets Most theoretical work on matching markets traces back to the papers of Gale and Shapley (1962) and Shapley and Shubik (1971). Gale and Shapley (1962) formulated a model of two-sided matching without side payments which they called the *marriage problem* for one-to-one matching and the *college admissions problem* for many-to-one matching. Shapley and Shubik (1971) formulated a model of two-sided matching with side payments which they called the *assignment game*, which is also called the *maximum weighted bipartite matching* (Karlin and Peres, 2017). Recently, there has been a surge of literature on the online version of assignment game and its generalizations, due to the important new application domain of Internet advertising (Mehta, 2013). Although there exist striking analogies between some of the results for the college admissions problem and those for the assignment game (Roth and Sotomayor, 1990), the key difference is that in college admissions there are agents and arms on both sides of the assignment who care about the outcome (Roth et al., 1993). This matters for a variety of reasons, not least of which is that if a matching outcome disagrees with the liking of the agents and arms involved, groups of them may be able to disregard the matching outcome by making private matches among themselves. This imposes constraints such as stability and fairness on what matching outcomes can be achieved. The main focus of our paper is the college admissions problem, which has no side payments and no centralized clearinghouse for coordination. Our goal is to design algorithms for college admissions that achieve the constraints such as stability and fairness under *uncertain preferences*.

Decentralized Matching Markets There has been significant work in the economics literature on congestion in decentralized markets, where participants cannot make enough offers and acceptances to clear the market. Roth and Xing (1997) discussed such a decentralized market for graduating clinical psychologists, focusing on the market’s timing aspect, which lasts over a day. They found that such a decentralized but coordinated market exhibits congestion since the interviews that a student could schedule were limited, and the resulting matching was unstable. Das and Kamenica (2005) considered that both sides of the decentralized market have uncertain preferences and presented an empirical study of the resulting matchings. Unlike these works, we provide both an analytical model and optimal learning algorithms for managing congestion. Haeringer and Wooders (2011) considered the decentralized job matching with complete information, where agents and arms are assumed to know the entire sequence of actions employed by their opponents. Coles et al. (2013) showed how introducing a signaling device in a decentralized matching market alleviates congestion. By contrast, we consider incomplete information where agents do not know their opponents’ actions, and focus on the optimal strategy under uncertain preference without signaling.

Bandit Learning in Decentralized Matching Markets A recent thread of research focuses on matching markets with bandit learning. Liu et al. (2020a,b) and Basu et al. (2021) extended the multi-armed bandit framework to a setting in which multiple agents compete with respect to their preferred matches. They proposed an algorithm to achieve low

cumulative regret for decentralized matching. Cen and Shah (2021) further incorporated costs and transfers in this bandit learning setting to faithfully model competition between agents. Sankararaman et al. (2021) studied the decentralized serial dictatorship setting, a two-sided matching market where the agents have heterogeneous valuations over the arms, while the arms have known uniform preference over agents. Unlike these works, we study the optimal strategy for single-stage matching, which involves no accumulated regret and allows arms to have heterogeneous preferences over agents.

College Admissions Another closely related line of work is the algorithmic literature on college admissions (see Gale and Shapley, 1962). The general setting involves multiple colleges competing for students in decentralized markets. Recently, Epple et al. (2006) modeled equilibrium admissions, financial aid, and enrollment. Fu (2014) studied effects of tuition on equilibria, incorporating application costs, admissions, and enrollment. Unlike these works, we emphasize students’ multidimensional abilities and the uncertainty of students’ acceptance probability. Chade and Smith (2006) provided an optimal algorithm for students’ application decisions, characterized as a portfolio choice problem. Hafalir et al. (2018) discussed student efforts instead of colleges’ response to congestion in decentralized college admissions with restricted applications. We instead consider colleges’ optimal strategies. Avery and Levin (2010) studied early admissions when colleges have no enrollment uncertainty and showed that a cutoff strategy is optimal in equilibrium. Chade et al. (2014) developed a decentralized Bayesian model of college admissions for two colleges and a continuum of students under a particular preference structure. By contrast, we study multiple colleges in the face of enrollment uncertainty. Azevedo and Leshno (2016) adopted a continuum model for students in a centralized market. They found a characterization of equilibria in terms of supply and demand. This model is different from the decentralized market studied in our paper. Che and Koh (2016) considered aggregated uncertainty in college admissions and focused on two colleges and a continuum of students. By contrast, our model considers incomplete information with multiple colleges and a finite number of arms. We study a statistical model for learning strategies in the face of enrollment uncertainty using historical data.

Stochastic Knapsack Problems Our optimization problem relates to the stochastic knapsack problems that have applications in advertising, logistics, medical diagnosis and robotics (Ross and Tsang, 1989; Dean et al., 2008; Bhalgat et al., 2011). In the stochastic knapsack problems, the rewards of items are deterministic, but the sizes of items are independent random variables with known, arbitrary distributions. The actual size of an item is known only when it is placed into the knapsack. The objective of a stochastic knapsack problem is to add items sequentially (one by one) into the knapsack of a given capacity to maximize the expected reward of the items that fit into the knapsack. Unlike the stochastic knapsack problems, our optimization problem involves the challenge of unknown and uncertain acceptance distributions due to latent state variables. It requires new algorithms for calibrating states and finding the optimal set of arms based on the historical data.

Multi-Agent Learning Our work also relates to multi-agent learning in noncooperative settings. Seuken and Zilberstein (2008) provided an overview of algorithms that address decentralized learning and coordination of multiple decision makers under uncertainty. See also Busoniu et al. (2008) for a survey of multi-agent learning approaches. Several papers

have attempted to address the non-stationarity in multi-agent environments by conditioning on other agents’ policies or using importance sampling (Tesauro, 2003; Foerster et al., 2017). However, these works require knowing other agents’ policies or actions, a condition which is impractical in decentralized matching markets. By contrast, our model framework is applicable given limited information regarding other agents. Recently, Tampuu et al. (2017) investigated a deep Q -learning approach to train competing Pong agents. Foerster et al. (2018) and Lowe et al. (2017) studied policy gradient with actor-critic for multi-agent learning. Our work differs from these in that it requires less coordination and information flow between agents. Hu and Storkey (2014) presented a prediction market model with multiple agents, showing that the market model solves the dual of a certain machine learning problem. By contrast, our model is motivated by a different class of markets—two-sided matching markets—and the model framework that we present aims at solving machine learning problems involving scarcity and uncertain preferences.

1.2 Contributions and Outline

We develop a framework for learning uncertain preferences that drive decision-making in the context of scarcity of shared resources. We summarize our principal methodological and theoretical contributions as follows.

- We provide a general introduction to the problem of learning under scarcity of shared resources, as exemplified by two-sided matching markets (Section 2). We develop a modeling framework for uncertain preferences that depends on a local state component (Theorem 1). We then construct an estimator for learning preferences under this model. We show that under regularity conditions the estimator converges at a minimax optimal rate (Theorem 2).
- We prove that a version of cutoff strategy approximates an optimal strategy for maximizing an agent’s expected payoff (Section 3). In doing so, we provide near-optimal bounds on the loss of expected payoff (Theorem 3). We propose a new data-driven method, the calibrated decentralized matching (CDM) method, to calibrate latent state variables for the purposes of determining a preferred outcome (Theorems 4 and 5). CDM takes the opportunity cost and penalty for exceeding the quota into account. Moreover, we discuss extensions of CDM to other real-world machine learning problems in which there is a scarcity of shared resources.
- We provide theoretical characterizations of the decisions derived from the CDM (Section 4). Specifically, we show an incentive-compatibility property (Theorem 6), show the stability of the matching outcomes obtained from CDM (Theorem 7), and we show the fairness of the matching process (Theorem 8).
- We conduct simulation studies (Section 5) to assess the properties of CDM and compare the payoffs with alternative methods. We find that arms can be better off under decentralized matching compared to centralized matching. We demonstrate aspects of the theoretical results through experiments (Section 6) using real data from college admissions and simulated graduate school admissions.

We conclude the paper with further research directions in Section 7. All proofs are provided in the Supplementary Appendix.

2. Learning under Scarcity and Uncertain Preferences

In this section, we define a learning model under a scarcity of shared resources and uncertain preferences. Denote a set of m agents by $\mathcal{P} = \{P_1, P_2, \dots, P_m\}$ and a set of n arms by $\mathcal{A} = \{A_1, A_2, \dots, A_n\}$, where \mathcal{P} and \mathcal{A} are participants on the two sides of the market. Each agent can attempt to pull multiple arms and there are no constraints on the overlap among the choices of different agents. When multiple agents select the same arm, only one agent can successfully pull the arm, with the choice of agent made according to the arm's preferences. For example, \mathcal{P} and \mathcal{A} might represent colleges and students in the college admissions market, or firms and workers in the decentralized labor market. The matching markets involve scarcity as the resources on both the agent and arm sides are limited. For example, colleges send admission offers to applicants. When multiple colleges send offers to the same applicant, the applicant can accept at most one offer. Each agent P_i has a limited quota of successfully pulling $q_i \geq 1$ arms, where $q_1 + q_2 + \dots + q_m \leq n$. We denote $[m] \equiv \{1, \dots, m\}$, $[n] \equiv \{1, \dots, n\}$, and $[K] \equiv \{1, \dots, K\}$.

In decentralized matching markets, agents and arms make their decisions independently of the decisions made by others. This feature distinguishes decentralized matching from centralized matching, which makes use of central clearinghouses to coordinate decision-making. Notable examples of centralized matching include the national medical residency matching (Roth, 1984) and public school choice (Abdulkadiroğlu and Sönmez, 2003). While centralized matching has been the major focus of the literature on matching, we see decentralized matching as having a potentially greater range of applications, and as providing a better platform on which to bring machine-learning tools to bear. Indeed, decentralized learning and decentralized matching are natural when (i) arms have uncertain preferences that depend on an unknown state that has a local component, and (ii) agents possess incomplete information on other agents' decisions.

2.1 Running Example

Our running example is the college admissions market, where colleges match with students (see Gale and Shapley, 1962). College admissions in countries such as the United States, Korea, and Japan are organized in a decentralized way, where colleges make admission offers to applicants. The admitted students accept or reject the offers, often within a short period of time. Consider a college $P_i \in \mathcal{P}$ with a quota of enrolling q_i students. It is not satisfactory for college P_i to only make offers to the q_i best-qualified students since some students may reject offers. The enrollment uncertainty for P_i can be attributed to a lack of two types of information. First, P_i has little knowledge of which other colleges admit the applicants admitted by P_i . Second, P_i is uncertain about students' preferences: how each applicant ranks the colleges that she has applied to. Thus colleges do not know their own popularity in the current year, which is an aggregate over the uncertain rankings by each applicant. In the face of such uncertainty, college P_i needs to decide which applicants to admit such that the entering class will meet reasonably close to its quota q_i and be close to the attainable optimum in quality.

2.2 Latent Utility

We assume that agents' deterministic preferences are a function of underlying latent utilities. In particular, we consider the following latent utility model:

$$U_i(A_j) = v_j + e_{ij}, \quad \forall i \in [m] \text{ and } j \in [n], \quad (1)$$

where $v_j \in [0, 1]$ is arm A_j 's *systematic score*, which is available to all agents, and $e_{ij} \in [0, 1]$ is an agent-specific *idiosyncratic fit* available only to agent P_i . For example, in college admissions, v_j can be a function of student A_j 's test score on a nationwide test observed by all colleges. The e_{ij} corresponds to a function of student A_j 's performance on college-specific essays or tests conducted by college P_i . In Appendix B.1, we show that a general utility function with multidimensional scores and fits can be transformed to the model (1) via the ANOVA decomposition. Thus, the separable structure in model (1) is without loss of generality and allows us to characterize the pattern of competition of agents in Section 3. Similar separable structures have been used in the matching markets literature (cf. Choo and Siow, 2006; Menzel, 2015; Chiappori and Salanié, 2016; Ashlagi et al., 2020). The analysis that we present in this paper will also hold if one restricts the range of e_{ij} to $[0, \bar{e}]$ with some $\bar{e} < 1$. This restriction may be necessary for some applications. For instance, suppose the idiosyncratic fit in college admissions is viewed as more important to colleges than the systematic score. In this case, enrolled students may find it unfair to have other enrolled students with significantly lower test scores than theirs. This, in turn, would result in a reputation cost for the college.

2.3 Uncertain Preferences

We now consider the uncertainty of an arm's acceptance of an agent's pulling, which depends on the arm's uncertain preference. The arm's uncertainty mainly consists of two parts. The first part is the *state of the world*, which determines an arm's preference over all agents. Let the parameter $s_i \in [0, 1]$ represent the local state of the world for agent P_i for $i \in [m]$ (Savage, 1972). The second part of the uncertainty is the agents' strategies, which is because an arm would only accept its most preferred agent among those who have pulled it. The main difficulty in decentralized markets is that an agent has incomplete information, which means that the agent lacks information on their opponents' strategies. The following theorem shows that although with incomplete information, there still exists a valid probability function that models an arm's acceptance probability.

Theorem 1 *There exists a probability mass function $\pi_i(s_i, v_j)$ for any $i \in [m], j \in [n]$ that characterizes the probability of arm $A_j \in \mathcal{A}$ accepting agent P_i , given any fixed strategies of other agents $P_{i'}, i' \neq i$. Moreover, the function $\pi_i(s_i, v_j)$ is increasing in s_i , and the expected utility that agent P_i receives from pulling arm $A_j \in \mathcal{A}$ is*

$$\mathbb{E}[\text{utility}] = (v_j + e_{ij}) \cdot \pi_i(s_i, v_j), \quad i \in [m] \text{ and } j \in [n]. \quad (2)$$

The arms' scores v_j are public information, and agents compete for arms with better scores. This competition aspect is incorporated into the acceptance probability $\pi_i(s_i, v_j)$ that depends on the score v_j . Moreover, since $\pi_i(s_i, v_j)$ is increasing in s_i , a large value of s_i

corresponds to the case that the agent P_i is popular. On the other hand, we note that the true state is unknown a priori to P_i . For example in college admissions, the *yield* defined as the rate at which a college's admitted students accept the offers is a proxy for the state s_i (cf. Che and Koh, 2016). However, the yield of the current year is unknown a priori to the college.

Example 1 *We show the explicit form of $\pi_i(s_i, v_j), i = 1, 2$ for a two-agent model with agents P_1 and P_2 . Suppose that P_1 pulls an arm $A_j \in \mathcal{A}$. Denote by $\mu_1(s_1)$ the probability that an arm prefers P_1 under the state s_1 . Let σ_2 be P_2 's strategy, which is defined by $\sigma_2(v_j, e_{2j}) = \mathbf{1}(P_2 \text{ pulls } A_j \text{ with attributes } \{v_j, e_{2j}\})$. Since A_j would be pulled by P_1 with probability $1 - \sigma_2(v_j, e_{2j})$ and pulled by both P_1 and P_2 with probability $\sigma_2(v_j, e_{2j})$, A_j would accept P_1 with probability $1 - \sigma_2(v_j, e_{2j}) + \mu_1(s_1)\sigma_2(v_j, e_{2j})$. Since e_{2j} is unknown to A_j and σ_2 depends on e_{2j} , the expected probability that A_j accepts P_1 is as follows,*

$$\pi_1(s_1, v_j) = 1 - \mathbb{E}[\sigma_2(v_j, e_{2j})] + \mu_1(s_1)\mathbb{E}[\sigma_2(v_j, e_{2j})],$$

where the expectation is taken over e_{2j} . Moreover, we can derive an explicit form of $\pi_2(s_2, v_j)$ by following a similar argument as above.

2.4 Learning Arms' Uncertain Preferences

Agents decide which arms to pull based on the attributes of the arms and historical matches. In general, we do not make the assumption that an arm with the same score and fit appears in the historical data in the matching market. Indeed, no college admissions applicant likely has exactly the same attributes as the applicants in previous years. Even the "repeated applicants" will often change their records in their second applications. The repeated applicants either have no offer or reject all their offers and wait a full year to apply for colleges. The value of learning from the historical matches accrues when the data are used to estimate an untried arm's acceptance probability.

Denote by $\mathcal{A}^t = \{A_1^t, A_2^t, \dots, A_{n_t}^t\}$ the set of arms at $t \in [T] \equiv \{1, \dots, T\}$. Let s_i^t be the state of agent P_i at time t . The state s_i^t is *unknown* until time $t+1$, and it is non-stationary and varies over t . For instance, the yield rate of a college may change over the years. For any arm $A_j^t \in \mathcal{A}^t$, there are an associated pair of score and fit values, (v_j^t, e_{ij}^t) , obtained from Eq. (1), where $i \in [m], j \in [n_t]$. We refer to (v_j^t, e_{ij}^t) as the attributes of arm A_j^t . Define the set

$$\mathcal{B}_i^t = \{j \mid P_i \text{ pulls arm } A_j^t \text{ at time } t \text{ for } 1 \leq j \leq n_t\},$$

where $|\mathcal{B}_i^t| = n_{it}$ and $n_{it} \leq n_t$. For any $j \in \mathcal{B}_i^t$, the outcome that P_i observes is whether an arm A_j^t accepted P_i , that is, $y_{ij}^t = \mathbf{1}\{A_j^t \text{ accepts } P_i\}$. Here, the iid outcome y_{ij}^t has the likelihood $\pi_i(s_i^t, v_j^t)$ that $y_{ij}^t = 1$ and the likelihood $1 - \pi_i(s_i^t, v_j^t)$ that $y_{ij}^t = 0$.

We estimate the function π_i based on historical data. Denote by $\mathcal{D} = \{(s_i^t, v_j^t, e_{ij}^t, y_{ij}^t) : i \in [m]; j \in \mathcal{B}_i^t; t \in [T]\}$ the training data. Let the log odds ratio be

$$f_i(s_i, v) = \log \left(\frac{\pi_i(s_i, v)}{1 - \pi_i(s_i, v)} \right).$$

There exist a variety of methods in supervised learning that can efficiently estimate the log odds ratio (Hastie et al., 2009). For concreteness, we focus on using a penalized log-likelihood method for the log odds ratio estimation in a reproducing kernel Hilbert space

(RKHS) (Wahba, 1999). To this end, we assume that f_i is from RKHS \mathcal{H}_{K_i} with the reproducing kernel K_i . Then we find $f_i \in \mathcal{H}_{K_i}$ to minimize

$$\sum_{t=1}^T \sum_{j \in \mathcal{B}_i^t} [-y_{ij}^t f_i(s_i^t, v_j^t) + \log(1 + \exp(f_i(s_i^t, v_j^t)))] + \frac{1}{2} \sum_{t=1}^T n_{it} \lambda_i \|f_i\|_{\mathcal{H}_{K_i}}^2, \quad (3)$$

where $\|\cdot\|_{\mathcal{H}_{K_i}}$ is a RKHS norm and $\lambda_i \geq 0$ is a tuning parameter. Consider a tensor product structure of the RKHS \mathcal{H}_{K_i} defined $K_i((s_i, v), (s'_i, v')) = K_i^s(s_i, s'_i) K_i^v(v, v')$ based on kernels K_i^s and K_i^v (Wahba et al., 1995). Assume a random feature expansion of the following form:

$$K_i^s(s_i, s'_i) = \mathbb{E}_{w_s}[\phi_i^s(s_i, w_s) \phi_i^s(s'_i, w_s)] \text{ and } K_i^v(v, v') = \mathbb{E}_{w_v}[\phi_i^v(v, w_v) \phi_i^v(v', w_v)],$$

where $\phi_i^s(\cdot, w_s)$ and $\phi_i^v(\cdot, w_v)$ are random features (Rahimi and Recht, 2008). Denote by $\{w_{s1}, w_{s2}, \dots, w_{sp}\}$ and $\{w_{v1}, w_{v2}, \dots, w_{vp}\}$ the sets of p independent copies of w_s and w_v , respectively. Write the feature vector as $\psi_i(s_i, v) \in \mathbb{R}^p$ with the l th entry equal to $\frac{1}{\sqrt{p}} \phi_i^s(s_i, w_{sl}) \phi_i^v(v, w_{vl})$, for $l = 1, \dots, p$. Then $K_i((s_i, v), (s'_i, v'))$ can be approximated by the product $\psi_i(s_i, v)^\top \psi_i(s'_i, v')$. Let the matrix Φ_i have rows $\psi_i(s_i^t, v_j^t)^\top$, where $j \in \mathcal{B}_i^t$ and $t \in [T]$. By the representer theorem in Kimeldorf and Wahba (1971), the solution to Eq. (3) has the form $\widehat{f}_i(s_i, v) = \psi_i(s_i, v)^\top \Phi_i^\top c_i$ for some vector c_i . We only need to find $\theta_i = \Phi_i^\top c_i \in \mathbb{R}^p$ to obtain a solution to Eq. (3):

$$\widehat{f}_i(s_i, v) = \psi_i(s_i, v)^\top \theta_i, \quad \forall i \in [m]. \quad (4)$$

Denote by Y_i the response vector with entries y_{ij}^t , $j \in \mathcal{B}_i^t$ and $t \in [T]$. Applying the Newton-Raphson method to Eq. (3) yields the following iterative updates for vector θ_i :

$$\theta_i^{(\nu+1)} = \left(\Phi_i^\top W_i^{(\nu)} \Phi_i + \sum_{t=1}^T n_{it} \lambda_i \mathbf{I} \right)^{-1} \Phi_i^\top W_i^{(\nu)} \left\{ \Phi_i \theta^{(\nu)} + (W_i^{(\nu)})^{-1} [Y_i - \pi_i^{(\nu)}] \right\}, \quad \nu \geq 1.$$

Here, $\theta_i^{(\nu)}$ is the ν th update of θ_i , and $W_i^{(\nu)} = \text{diag}[\pi_i^{(\nu)}(s_i^t, v_j^t)(1 - \pi_i^{(\nu)}(s_i^t, v_j^t))]_{j \in \mathcal{B}_i^t, 1 \leq t \leq T}$ is a weight matrix with $\pi_i^{(\nu)} = [1 + \exp(-f_i^{(\nu)})]^{-1}$ and $f_i^{(\nu)} = \phi_i^\top \theta_i^{(\nu)}$. The tuning parameter $\lambda_i \geq 0$ can be selected using cross validation or GACV (see Wahba, 1999).

Theorem 2 *The integrated squared error of the estimate in Eq. (4) satisfies the following inequality:*

$$\mathbb{E}[(\widehat{f}_i - f_i)^2] \leq c_f [T(\log T)^{-1}]^{-2r/(2r+1)}, \quad \forall i \in [m],$$

for sufficiently large T , where $\lambda_i \leq c_\lambda [T(\log T)^{-1}]^{-2r/(2r+1)}$ and $p \geq c_p [T(\log T)^{-1}]^{-2r/(2r+1)}$. Here the constants $c_f, c_\lambda, c_p > 0$ are independent of T , and $r \geq 1$ denotes the smoothness of kernels such that $K_i^s(s, \cdot)$ and $K_i^v(v, \cdot)$ have squared integrable r th-order derivatives. Moreover, the estimate in Eq. (4) is minimax rate-optimal.

We make three remarks on the theorem. First, although f_i depends on both covariates s and v , the optimal rate given in the theorem is very close to the minimax rate for the one-dimensional model, with the only difference in the logarithm term. Second, it is possible to

extend the theorem to allow different orders of the smoothness of the kernels K_i^s and K_i^v . Finally, we can predict the acceptance probability for a new arm. Let $\mathcal{A}^{T+1} = \{A_1, \dots, A_n\}$ be the set of arms at a new time point $T + 1$, where each arm A_j has attributes obtained from Eq. (1). Using the estimated log odds ratio \hat{f}_i in Eq. (4) and given the state s_i at time $T + 1$, we have the estimate

$$\hat{\pi}_i(s_i, v_j) = \left[1 + \exp(-\hat{f}_i(s_i, v_j)) \right]^{-1}, \quad (5)$$

which serves as the prediction of the probability that A_j would accept P_i at time $T + 1$.

3. Optimal Strategies in Decentralized Matching

We now study the optimal strategies in decentralized matching at a new time point $T + 1$, where the matching involves a single stage interaction. First, Nature draws a state denoted by s_i^* for agent P_i such that the arms' preferences are realized. Next, arms simultaneously show their interests to all agents. For example, students apply to colleges. Under the assumption that students face negligible application costs, the students' dominant strategy is submitting applications to all colleges as they do not know how colleges evaluate their academic ability or personal essays (Avery and Levin, 2010; Che and Koh, 2016). Then, agents simultaneously decide which arms to pull based on the arm's attributes. Finally, each arm accepts one of the agents (if any) that pulled it.

Since agents simultaneously pull sets of arms once at the new time point $T + 1$ and each arm accepts one of the agents that pulled it, we refer to such matching as the *single-stage decentralized matching*. Each agent's goal is to learn the optimal strategies that maximize agents' expected payoffs based on historical data at $t = 1, \dots, T$. In the repeated game, the number of agents are the same at $t = 1, \dots, T$. However, our model applies to a repeated game with an arbitrary number of agents. See the left plot of Figure 1 for an illustration of the single-stage decentralized matching.

3.1 Algorithm

We propose the calibrated decentralized matching (CDM) procedure for the single-stage decentralized matching. The CDM consists of seven main steps. Step 1 is to predict the arms' acceptance probability. Step 2 is to estimate the distribution of the true state. Step 3 is to construct the cutoffs for a given state and its perturbation. Step 4 is to calculate the arm sets for the given state and its perturbation. Step 5 is to obtain the marginal set for the given state. In addition, Steps 3 to 5 are carried out repeatedly over a number of states. Step 6 is to calibrate the state for maximizing the agent's average-case expected payoff, which calibration is referred to as the mean calibration. The key idea of the mean calibration is to balance the tradeoff between the opportunity cost and the penalty for exceeding the quota. Finally, Step 7 is to produce the final arm set under the calibrated state. We summarize our procedure in Algorithm 1 first, then discuss each step in detail.

3.2 Agent's Expected Payoff

Since the decentralized market has no centralized clearinghouse for coordination, each agent's incentive is to act to maximize the expected payoff. An agent's expected payoff




Figure 1: The left plot shows the process of the single-stage decentralized matching, with no centralized clearinghouse for coordination. The middle plot shows the cutoff $\hat{e}_i(s_i, v)$ of Theorem 3 when Eq. (8) holds. The shaded area represents $\hat{\mathcal{B}}_i(s_i)$. The dotted line represents $b_i^-(s_i) - v$ in verifying Eq. (8). The dashed line represents $\hat{b}_i(s_i) - v$, and if thresholding to $e_i \in [0, 1]$, it yields the cutoff $\hat{e}_i(s_i, v)$ in Eq. (9) and is denoted by the blue solid line segments. The right plot shows the cutoff $\hat{e}_i(s_i, v)$ when Eq. (8) does not hold.

Algorithm 1 Calibrated decentralized matching (CDM) under the mean calibration

- 1: **Input:** Historical data $\mathcal{D} = \{(s_i^t, v_j^t, e_{ij}^t, y_{ij}^t) : i \in [m]; j \in \mathcal{B}_i^t; t \in [T]\}$; New arm set \mathcal{A}^{T+1} with attributes $\{(v_j, e_{ij}) : i \in [m]; j \in [n]\}$ at time $T + 1$; Penalty $\{\gamma_i : i \in [m]\}$ for exceeding the quota.
 - 2: **for** $i = 1$ to m **do**
 - 3: **Step 1:** Predict the acceptance probability $\hat{\pi}_i(s_i, v_j)$ by Eq. (5);
 - 4: **Step 2:** Estimate the probability density function of the state s_i^* by Eq. (11);
 - 5: **for** s_i in a linearly spaced vector in $[0, 1]$ **do**
 - 6: **Step 3:** Construct the cutoffs $\hat{e}_i(s_i, v)$ and $\hat{e}_i(s_i + \delta s_i, v)$ by Eq. (9) with a small $\delta s_i > 0$;
 - 7: **Step 4:** Calculate the arm sets $\hat{\mathcal{B}}_i(s_i)$ and $\hat{\mathcal{B}}_i(s_i + \delta s_i)$ by Eq. (7);
 - 8: **Step 5:** Obtain the marginal set $\partial \hat{\mathcal{B}}_i(s_i) = \{\hat{\mathcal{B}}_i(s_i - \delta s_i) \setminus \hat{\mathcal{B}}_i(s_i)\}$;
 - 9: **end for**
 - 10: **Step 6:** Calibrate the state s_i such that the difference between LHS and RHS of Eq. (12) is below a pre-specified tolerance level;
 - 11: **Step 7:** Calculate the arm set $\hat{\mathcal{B}}_i(s_i)$ by Eq. (7) under the calibrated state s_i .
 - 12: **end for**
 - 13: **Output:** The arm sets $\hat{\mathcal{B}}_1(s_1), \hat{\mathcal{B}}_2(s_2), \dots, \hat{\mathcal{B}}_m(s_m)$ for agents.
-

consists of two parts: the expected utilities of arms that the agent pulls and the penalty for exceeding the quota. Let $\mathcal{B}_i(s_i) \subseteq \mathcal{A}^{T+1}$ be the set of arms that agent P_i pulls at $T + 1$

given the state s_i . We consider the following form of agent P_i 's expected payoff,

$$\mathcal{U}_i[\mathcal{B}_i(s_i)] = \sum_{j \in \mathcal{B}_i(s_i)} (v_j + e_{ij}) \cdot \pi_i(s_i, v_j) - \gamma_i \cdot \max \left\{ \sum_{j \in \mathcal{B}_i(s_i)} \pi_i(s_i, v_j) - q_i, 0 \right\}. \quad (6)$$

Here, $\gamma_i > 0$ denotes the marginal penalty for exceeding the quota, which originates from the scarcity of agents' capacity. The definition of γ_i is motivated by numerous applications of decentralized matching markets. For instance, Princeton reported 1,100 students joining its incoming class in 1995–96, significantly exceeding its capacity. As a result, Princeton had to set up mobile homes and build new dorms to accommodate the students (Avery et al., 2003). Another example is that Korea faces the challenge of controlling yield for colleges, since students apply to a department instead of a college, and each department faces a relatively rigid and low quota. A department that exceeds the quota is subject to a rather harsh penalty by the government in the form of sharply reduced capacities in the subsequent year (Che and Koh, 2016). The choice of γ_i in our framework is user-specified and it depends on the consequences of exceeding the quota. We assume that $\gamma_i > \max_{j \in \mathcal{A}^{T+1}} \{v_j + e_{ij}\}$; that is, the penalty is larger than an arm's latent utility. The expected payoff $\mathcal{U}_i[\mathcal{B}_i(s_i)]$ in Eq. (6) captures agent P_i 's incentive—namely, its benefit from pulling arms from the set $\mathcal{B}_i(s_i)$, holding fixed its opponents' decisions. As a result, in the optimization problem in Eq. (6), taking the optimal action is aligned with the agent's incentive. Eq. (6) excludes the situation when P_i faces extra uncertainty in receiving the reward even if P_i has successfully pulled an arm. For example in the dating market, Eq. (6) models the agent's expected payoff from the dates, instead of a subsequent relationship that may eventually result from a date.

Finding and checking an optimal solution to Eq. (6) is NP-complete; roughly because we need to individually check a significant fraction of the combinations of $\text{card}(\mathcal{A}^{T+1})$ arms to determine the optimal solution for Eq. (6). Since the number of combinations grows exponentially with the number $\text{card}(\mathcal{A}^{T+1})$, the complexity of any systematic algorithm becomes impractically large. We propose a *cutoff strategy* which greedily pulls arms according to the latent utilities and show that it is near-optimal for agents in the following theorem.

Theorem 3 *Agent P_i 's cutoff strategy which pulls arm from the set*

$$\widehat{\mathcal{B}}_i(s_i) = \{j \mid A_j \in \mathcal{A}^{T+1} \text{ whose attributes } (v_j, e_{ij}) \text{ satisfy } e_{ij} \geq \widehat{e}_i(s_i, v_j)\} \quad (7)$$

is near-optimal as its expected payoff satisfies

$$0 \leq \max_{\mathcal{B}_i \subseteq \mathcal{A}^{T+1}} \mathcal{U}_i[\mathcal{B}_i] - \mathcal{U}_i[\widehat{\mathcal{B}}_i(s_i)] \leq UE^\dagger.$$

Here the expected payoff \mathcal{U}_i is defined in Eq. (6), and UE^\dagger is defined by $UE^\dagger \equiv [\min_{j \in \mathcal{B}_i^-(s_i)} (v_j + e_{ij})][q_i - \sum_{j \in \mathcal{B}_i^-(s_i)} \pi_i(s_i, v_j)]$. The cutoff $\widehat{e}_i(s_i, v)$ is chosen according to Eq. (9), where $\widehat{e}_i(s_i, v)$ is decreasing in $v \in [0, 1]$ and satisfies $\partial \widehat{e}_i(s_i, v) / \partial v = -1$ when $\widehat{e}_i(s_i, v) \in (0, 1)$. Moreover, if there is a continuum of arms and $\pi_i(\cdot, v)$ is continuous in v , then $UE^\dagger = 0$.

We now specify the cutoff $\widehat{e}_i(s_i, v)$. Suppose that arms on the cutoff have the latent utility $b_i \geq 0$. The expected number of arms in $\widehat{\mathcal{B}}_i(s_i)$ that would accept P_i is:

$$\Pi_i(b_i) \equiv \sum_{j \in \mathcal{A}^{T+1}} \mathbf{1}(e_{ij} \geq \min\{\max\{b_i - v_j, 0\}, 1\}) \pi_i(s_i, v_j).$$

If there exists some $b_i \geq 0$ such that $\Pi_i(b_i) = q_i$, we let $\widehat{b}_i(s_i) = b_i$ and we have $\widehat{e}_i(s_i, v) = \min\{\max\{\widehat{b}_i(s_i) - v, 0\}, 1\}$. On the other hand, if there is no solution to $\Pi_i(b_i) = q_i$, we let

$$b_i^+(s_i) = \arg \max_{b_i \geq 0} \{\Pi_i(b_i) > q_i\} \quad \text{and} \quad b_i^-(s_i) = \arg \min_{b_i \geq 0} \{\Pi_i(b_i) < q_i\}.$$

To choose between $b_i^+(s_i)$ or $b_i^-(s_i)$, it is necessary to balance the expected utility and the expected penalty for exceeding the quota due to pulling arms on the *boundary*. Define two cutoffs $e_i^+(s_i, v) \equiv \min\{\max\{b_i^+(s_i) - v, 0\}, 1\}$ and $e_i^-(s_i, v) \equiv \min\{\max\{b_i^-(s_i) - v, 0\}, 1\}$, which correspond to arm sets $\mathcal{B}_i^+(s_i) = \{j \mid e_{ij} \geq e_i^+(s_i, v_j)\}$ and $\mathcal{B}_i^-(s_i) = \{j \mid e_{ij} \geq e_i^-(s_i, v_j)\}$, respectively. Then the arms on the boundary are those in the discrete set $\{\mathcal{B}_i^+(s_i) \setminus \mathcal{B}_i^-(s_i)\}$, whose expected utility is larger than the expected penalty if

$$\sum_{j \in \mathcal{B}_i^+(s_i) \setminus \mathcal{B}_i^-(s_i)} (v_j + e_{ij}) \cdot \pi_i(s_i, v_j) \geq \gamma_i \sum_{j \in \mathcal{B}_i^+(s_i)} \pi_i(s_i, v_j) - \gamma_i q_i. \quad (8)$$

Here set difference $\{\mathcal{B}_i^+(s_i) \setminus \mathcal{B}_i^-(s_i)\}$ consists of all elements in $\mathcal{B}_i^+(s_i)$ that are not in $\mathcal{B}_i^-(s_i)$. If Eq. (8) holds, let $\widehat{b}_i(s_i) = b_i^+(s_i)$; otherwise, let $\widehat{b}_i(s_i) = b_i^-(s_i)$. Finally, we define the cutoff

$$\widehat{e}_i(s_i, v) = \min \left\{ \max \left\{ \widehat{b}_i(s_i) - v, 0 \right\}, 1 \right\}, \quad (9)$$

Here we require $\widehat{e}_i(s_i, v) \in [0, 1]$ due to the assumption that e_i in Eq. (1) lies in the interval $[0, 1]$. It is clear that $\partial \widehat{e}_i(s_i, v) / \partial v = -1$ when $\widehat{e}_i(s_i, v) \in (0, 1)$. Thus, P_i prefers the arms with larger latent utilities defined by Eq. (1). Figure 1 illustrates this cutoff strategy. Although arms *independently* accept or reject agents in decentralized markets, the cutoff $\widehat{e}_i(s_i, v)$ in Eq. (9) ensures that the expected number of arms accepting P_i excluding those on the boundary is bounded by the quota q_i . According to Eq. (8), the arm on the boundary is pulled if the expected utility is larger than the expected penalty.

The cutoff strategy in Theorem 3 relates to the straightforward behavior (see Fisman et al., 2006), where agents pull arms that they value more than those they do not pull. Theorem 3 shows that although agents face uncertainty with respect to acceptance of their offers in decentralized markets, the straightforward behavior suffices.

3.3 Calibrated Decentralized Matching (CDM)

Let s_i^* be the true state for agent P_i at time $T+1$. Theorem 3 shows that the cutoff strategy maximizes P_i 's expected payoff if s_i^* is known. However, the true state s_i^* is generally unknown in practice. A natural question is how to calibrate the state s_i in Theorem 3. We propose a calibration method for s_i that maximizes the average-case expected payoff $\mathbb{E}_{s_i^*} \{\mathcal{U}_i[\widehat{\mathcal{B}}_i(s_i)]\}$ over the uncertain true state s_i^* . To formulate the theorem, we introduce some additional notation. Let $\partial \widehat{\mathcal{B}}_i(s_i)$ be the *marginal set*, defined as the change of the discrete set $\widehat{\mathcal{B}}_i(s_i)$ with respect to a perturbation of s_i :

$$\partial \widehat{\mathcal{B}}_i(s_i) \equiv \lim_{\delta s_i \rightarrow 0_+} \left\{ \widehat{\mathcal{B}}_i(s_i - \delta s_i) \setminus \widehat{\mathcal{B}}_i(s_i) \right\}.$$

Theorem 4 *The average-case expected payoff, $\mathbb{E}_{s_i^*}\{\mathcal{U}_i[\widehat{\mathcal{B}}_i(s_i)]\}$, is maximized if $s_i \in (0, 1)$ is chosen as the solution to the following equation,*

$$\begin{aligned} \mathbb{P}(s_i^* \neq s_i) \sum_{j \in \partial \widehat{\mathcal{B}}_i(s_i)} (v_j + e_{ij}) \cdot \mathbb{E}_{s_i^*}[\pi_i(s_i^*, v_j) \mid s_i^* \neq s_i] \\ = \gamma_i \mathbb{P}(s_i < s_i^* \leq 1) \sum_{j \in \partial \widehat{\mathcal{B}}_i(s_i)} \mathbb{E}_{s_i^*}[\pi_i(s_i^*, v_j) \mid s_i < s_i^* \leq 1]. \end{aligned} \quad (10)$$

We refer to the calibration by Theorem 4 as the *mean calibration* since it maximizes the agent's average-case expected payoff. We show in the proof that the average-case expected payoff has a unique maximizer that is given by the solution to Eq. (10). The key idea of the mean calibration is to balance the tradeoff between the opportunity cost and the penalty for exceeding the quota. In particular, the left side of Eq. (10) considers the opportunity cost and the marginal utility, and the right side of Eq. (10) estimates the marginal penalty for exceeding the quota. We make four remarks on Theorem 4. First, the calibrated state in Theorem 4 is different from a naive estimate of the mean $\mathbb{E}[s_i^*]$. We show in the simulation of Section 5 that the latter estimate is inefficient in maximizing the agent's expected payoff compared to the calibration by Theorem 4.

Second, the calibration in Theorem 4 takes agents' competition into account. Note that $\mathbb{E}_{s_i^*}[\pi_i(s_i^*, v_j) \mid s_i < s_i^* \leq 1]$ in Eq. (10) is strictly increasing in s_i due to the monotonicity of $\pi_i(s_i, v_j)$ in s_i . Thus, it is more costly for an agent to pull an arm when the agent is popular. This result is intuitive because when an agent P_i is popular, it is more likely that an arm A_j pulled by multiple agents would accept P_i . Since P_i has a larger probability of exceeding the quota when P_i is popular, it is more costly for P_i to pull A_j compared to the case when P_i is not popular.

Third, we discuss the algorithm for implementing the mean calibration in Theorem 4. A key quantity is the probability density function of s_i^* , which can be estimated based on the historical data of the states $\{s_i^1, \dots, s_i^T\}$. By the kernel density method, we obtain an estimator for the probability density function of s_i^* as follows,

$$\widehat{p}_{s_i^*}(\cdot) = \frac{1}{T} \sum_{t=1}^T K_i^s \left(\frac{\cdot - s_i^t}{h} \right), \quad (11)$$

where h is a bandwidth parameter and K_i^s is the kernel function introduced in Section 2.4. It is well known that the estimator $\widehat{p}_{s_i^*}(\cdot)$ achieves the minimax rate-optimality (Silverman, 1986). Another key quantity for implementing Theorem 4 is the acceptance probability $\pi_i(s_i, v)$, which can also be estimated based on the historical data as shown in Eq. (5), $\widehat{\pi}_i(s_i, v) = [1 + \exp(-\widehat{f}_i(s_i, v))]^{-1}$. The last key quantity is the marginal set $\partial \widehat{\mathcal{B}}_i(s_i)$, which can be computed as follows. For any s_i and a small perturbation δs_i , we choose the cutoffs $\widehat{e}_i(s_i, v)$ and $\widehat{e}_i(s_i + \delta s_i, v)$ according to Eq. (9) and calculate the arm sets $\widehat{\mathcal{B}}_i(s_i)$ and $\widehat{\mathcal{B}}_i(s_i + \delta s_i)$ by Eq. (7), which together give the marginal set $\partial \widehat{\mathcal{B}}_i(s_i) = \{\widehat{\mathcal{B}}_i(s_i - \delta s_i) \setminus \widehat{\mathcal{B}}_i(s_i)\}$. Given the estimators $\widehat{F}_{s_i^*}(s_i)$, $\widehat{\pi}_i(s_i, v)$ and the marginal set $\partial \widehat{\mathcal{B}}_i(s_i)$, we calibrate s_i in Theorem 4

as the solution to the following equation,

$$\begin{aligned} & \sum_{j \in \partial \widehat{\mathcal{B}}_i(s_i)} (v_j + e_{ij}) \cdot \left[\int_0^1 \widehat{\pi}_i(t, v_j) \widehat{p}_{s_i^*}(t) dt - \widehat{\pi}_i(s_i, v_j) \widehat{p}_{s_i^*}(s_i) \right] \\ &= \gamma_i \sum_{j \in \partial \widehat{\mathcal{B}}_i(s_i)} \left[\int_{s_i}^1 \widehat{\pi}_i(t, v_j) \widehat{p}_{s_i^*}(t) dt - \widehat{\pi}_i(s_i, v_j) \widehat{p}_{s_i^*}(s_i) \right]. \end{aligned} \quad (12)$$

We summarize the CDM under the mean calibration in Algorithm 1. It is possible to incorporate side information into the estimation of $F_{s^*}(\cdot)$. For example, one can incorporate the belief that an agent tends to be popular at $T + 1$ by overweighting the popular states in Eq. (11). Moreover, if Eq. (12) has more than one solution, then s_i is chosen as the largest one. If the probability density function of s_i^* has a discrete support, we change the objective in Eq. (12) to choosing the minimal $s_i \in [0, 1]$ such that the left side of Eq. (12) is not less than the right side of Eq. (12). Here the search of s_i starts from the maximum value in the support and decreases to the minimal value.

Finally, besides the average-case expected payoff in Theorem 4, we also consider the worst-case expected payoff concerning s_i^* . In the following theorem, we propose a *maximin calibration*, where the calibration maximizes the minimal expected payoff $\min_{s_i^*} \{\mathcal{U}_i[\widehat{\mathcal{B}}_i(s_i)]\}$ over the uncertain true state s_i^* .

Theorem 5 *The worst-case expected payoff $\min_{s_i^*} \{\mathcal{U}_i[\widehat{\mathcal{B}}_i(s_i)]\}$ is maximized if $s_i \in [0, 1]$ is chosen as the solution to the following equation,*

$$\begin{aligned} & \sum_{j \in \widehat{\mathcal{B}}_i(s_i)} (v_j + e_{ij}) \cdot [\pi_i(1, v_j) - \pi_i(0, v_j)] - \gamma_i \sum_{j \in \widehat{\mathcal{B}}_i(s_i)} \pi_i(1, v_j) + \gamma_i q_i \\ &= \sum_{j \in \widehat{\mathcal{B}}_i(1)} (v_j + e_{ij}) \cdot \pi_i(1, v_j) - \sum_{j \in \widehat{\mathcal{B}}_i(0)} (v_j + e_{ij}) \cdot \pi_i(0, v_j). \end{aligned}$$

We also summarize in Algorithm 2 of Appendix A.2 the CDM algorithm under the maximin calibration.

The calibrations in Theorems 4 and 5 take into account the opportunity cost and a penalty for exceeding the quota in order to ensure maximal payoff under scarcity. These analyses contribute to the emerging calibration literature in machine learning (Kennedy and O’Hagan, 2001; Guo et al., 2017). Beyond decentralized matching markets, the CDM framework has applications to other noncooperative multi-agent learning problems. For example in decision-making problems in manufacturing, information gathering, and load balancing, scarcity often arises from capacity and time constraints (see Ooi and Wornell, 1996; Cogill et al., 2006; Seuken and Zilberstein, 2008). In these applications, agents have partial, heterogeneous information regarding the state of the world and the other agents. Moreover, they need to cope with uncertainty about the outcomes arising from their actions. Related problems, such as the analysis of social dilemmas (Leibo et al., 2017) and competitive games (Tampuu et al., 2017) also require agents to compete for scarce resources. The CDM framework can be extended in several ways to be suitable for these applications. For instance, one may be able to replace the expected payoff in Eq. (6) with other reward

functions. A primary motivation behind CDM is that defining the reward function does not require knowing the observations and actions of other agents. This is not the case for most traditional learning methods in multi-agent domains (Foerster et al., 2018; Lowe et al., 2017). Moreover, it is possible to modify the single-stage payoff in Eq. (6) to represent discounted multi-stage payoffs, a topic that we are actively investigating.

4. Properties of CDM: Incentives, Stability and Fairness

Each participant alone, on both sides of the market, knows their own preference. However, the decentralized market has no centralized system for eliciting the participants’ preferences. The proposed CDM algorithm provides a mechanism to aggregate the historical data of the arms’ choices. It also provides a framework for the learning of the agent strategies. This section addresses the question of whether CDM gives the agents incentives to act according to their true preferences. Such an incentive-compatibility property for agents is desired for the design of matching markets (Roth, 1982).

A related problem to incentives is the stability of the matching outcomes that we obtain. We show that any pair of agent and arm has no incentive to disregard the CDM matching and seek an alternative outcome in the context of incomplete information in decentralized markets.

Finally, a different kind of problem concerning the matching process is fairness. For example, it is crucial to have a matching process that is fair for students in college admissions. We show that the CDM is fair for arms.

4.1 Incentive-Compatibility

We first need to define incentive-compatibility in decentralized matching. The payoffs of a matching problem are defined by the outcome of the matching process. A procedure that gives agents an incentive to act according to true preferences is defined as one that aggregates participants’ actions so that, in the resulting game, it is a dominant strategy for each agent to act on their preferences honestly. That is, agents pull arms according to the arm’s latent utility in Eq. (1). Such a procedure is called incentive-compatible when no matter what strategies other players may play, an agent who deviates from the true preference can achieve no better outcome than if the agent had acted on their preferences honestly.

Theorem 6 *As $T \rightarrow \infty$, the CDM gives agents an incentive to act straightforwardly on their true preferences. That is, agents pull arms according to the arm’s latent utility in Eq. (1).*

Arms pulled by multiple agents self-select based on preference. The CDM algorithm provides a mechanism to aggregate the historical matches from arms’ preference-based sorting in decentralized markets. Theorem 6 affirms an incentive-compatibility property for agents even with arbitrary uncertain preference profiles of arms. This theorem also distinguishes CDM from strategies according to arms’ expected utilities in Eq. (2) (see Das and Kamenica, 2005).

We also provide a discussion on incentive-compatibility for $T < \infty$. According to the definition, there are two parts of incentive-compatibility: (i) acting on the agents’ true

preferences and (ii) it is a dominant strategy in terms of maximizing the payoffs. We discuss these two parts for the CDM for the finite T . First, Theorem 3 asserts that the CDM uses a cutoff strategy that greedily pulls arms according to the latent utilities in Eq. (1). Since the arms' latent utilities determine the agents' true preferences, the CDM allows the agents to act on the true preferences for the finite T . Second, the CDM yields the strategy that approximates the optimal solution to the optimization in Eq. (6), which is computationally intractable and involves unknown quantities. The proposed CDM learns the unknown quantities from historical data with the finite T , and achieves the minimax convergence rate, shown in Theorem 2. Moreover, the CDM calibrates the uncertain state by taking opportunity costs into account. Together, the CDM yields the strategy that well approximates the optimal strategy for the finite T . The numerical examples in Section 5.3 demonstrate that no matter what strategies other players may play, the alternative methods that deviate from the true preferences would achieve worse outcomes than the CDM that allows agents to act on their preferences honestly. In summary, for the finite T , the CDM guarantees approximate incentive-compatibility in that agents can act on their true preferences and achieve the almost optimal payoffs, which are better than alternative methods that deviate from the true preferences.

4.2 Stability

An agent-arm pair (P_i, A_j) is referred to as *matched* if P_i pulls A_j , and A_j accepts P_i . A *blocking pair* is an agent-arm pair that is not matched, but both the agent and arm prefer to be matched together. Gale and Shapley (1962) defined a matching of arms to agents to be *stable* if there is no blocking pair. This definition was motivated by the concern that following a matching process, some agent-arm pairs with both the incentive and the power may deviate together, thus hindering the implementation of the intended matching outcome. This form of stability is widely considered to be a key factor in a successful implementation of a centralized clearinghouse (Roth, 1990).

Gale and Shapley (1962) also proved that the set of stable matching is never empty. Among all stable matchings, the one that is the most preferred by all arms is called an *arm-optimal matching*. The one that is the most preferred by all agents is called an *agent-optimal matching*. The arm- (or agent-) optimal matching is also agents' (or arms') least-preferred matching among all stable matchings (Knuth, 1997).

The stability of a centralized matching relies on the stringent assumption of complete information on the preferences of the participants. In our decentralized matching formulation, it is more natural to assume that agents only have incomplete information on preferences. Thus we require a modification of the stability notion in Gale and Shapley (1962). In particular, it is necessary to specify the uninformed agents' beliefs that might block a candidate stable matching (cf. Liu et al., 2014). We formulate the following notion of *individual rationality* to specify the beliefs of uninformed agents. Given any P_i 's strategy, let \mathcal{N}_i be the expected number of arms that would accept P_i . An additional arm A_j with attributes (v_j, e_{ij}) , which has a smaller latent utility than those currently pulled by P_i , is acceptable to P_i if and only if individual rationality holds:

$$(v_j + e_{ij}) \cdot \pi_i(s_i, v_j) \geq \gamma_i \cdot \max \{ \mathcal{N}_i + \pi_i(s_i, v_j) - q_i, 0 \}. \quad (13)$$

In other words, under P_i 's current strategy, A_j is acceptable to P_i if and only if A_j 's expected utility is at least the expected penalty for exceeding the quota. By Eq. (8), the proposed CDM satisfies the individual rationality condition (13).

A matching in a decentralized market is defined as stable if there is no blocking pair under the individual rationality condition (13). This stability notion is motivated by considering feasibility constraints (for example, single-stage interaction) in decentralized markets. In practice, participants are often compelled to accept the outcomes in such markets. For instance, the procedure in which some high school athletes are matched to colleges involves signing “letters of intent,” which prohibits athletes from further negotiating with other colleges. Suppose that an arm A_j does not satisfy Eq. (13). Then P_i finds that A_j only has a small latent utility but pulling A_j may incur a large penalty due to exceeding the quota. It is individually rational that P_i does not pull A_j in order to receive at least zero payoff. After the matching process, suppose P_i finds that A_j is acceptable. That is, P_i has remaining capacity, and A_j prefers to match to P_i . Still, P_i has no incentive to disregard the matching and seek an alternative outcome due to the feasibility constraints in decentralized markets.

Theorem 7 *As $T \rightarrow \infty$, the CDM procedure yields a stable matching. That is, agent and arm has no incentive to disregard the CDM matching and seek an alternative outcome in decentralized markets.*

There exists a lattice structure of stable matchings yielded by the CDM. If CDM allows each agent to pull all arms, the outcome is the arm-optimal stable matching. If CDM allows each agent to pull arms up to their quota, and those pulled arms are distinct, the outcome is the agent-optimal stable matching (Roth, 2008). In practice, the stable matching yielded by CDM is generally between the agent-optimal and arm-optimal matchings due to the competition of agents and a lack of coordination. Indeed, agents are incentivized to pull more arms than their quotas in decentralized markets in order to hedge against the uncertainty of whether the pulled arms will accept them. On the other hand, arms may benefit from this uncertainty. For example, suppose the college admissions market uses a clearinghouse to form a centralized market. All students apply to top colleges. In that case, the deferred acceptance algorithm yields the college-optimal matching (Gale and Shapley, 1962; Knuth, 1997). Compared to the CDM outcomes, students are worse off in the centralized market.

In contrast to CDM, the strategies of pulling arms according to the expected utilities in Eq. (2) may result in an unstable matching. For instance, if the agent P_i decides according to A_j 's expected utility, $\pi_i(s_i, v_j)(v_j + e_{ij})$, then P_i will not pull A_j if the acceptance probability $\pi_i(s_i, v_j)$ is small enough such that the expected utility of A_j is the smallest among all arms, while the latent utility $(v_j + e_{ij})$ of A_j is the largest. As a result, (P_i, A_j) will form a blocking pair if A_j is unmatched, and the matching is unstable.

We also provide a discussions on stability for $T < \infty$. There are two cases that an agent-arm pair (P_i, A_j) forms a blocking pair: (i) P_i prefers A_j to some of its matched arms and (ii) P_i has unfilled quota, and A_j is unmatched. We discuss these two cases of a blocking pair for the CDM for the finite T . For the first case, we note that Theorem 3 asserts that P_i must have pulled A_j according to the cutoff strategy of CDM. However, P_i must have been subsequently rejected by A_j in favor of some agent that A_j liked better. Hence, A_j must prefer its currently matched agent to P_i , and there is no instability. This

holds for finite T . For the second case, we note that P_i did not pull A_j since otherwise, A_j would have accepted P_i . This happens when the individual rationality in Eq. (13) suggests that P_i would find A_j unacceptable. For finite T , since the estimate of the population acceptance probability could be imperfect, the second case is possible to happen. However, these neglected arms are on the boundary of acceptance. In summary, for finite T , the CDM achieves the approximate stability in that there is no agent-arm pair such that the agent prefers the arm to some of its matched arms, and the instability occurs only when an agent missed some arms on the boundary of acceptance and those arms end up unmatched. However, we argue that in practice, the agent has limited incentive to deviate the intended matching outcome for those arms on the boundary of acceptance. Moreover, the numerical examples in Sections 5.1 and 5.2 demonstrate that for finite T , the CDM is likely to pull the arms are on the boundary of acceptance, and there is no instability.

4.3 Fairness

We consider fairness in terms of “no justified envy” (Balinski and Sönmez, 1999; Abdulkadiroğlu and Sönmez, 2003). An arm A_j has justified envy if there exists an agent P_i such that (P_i, A_j) form a blocking pair. A matching process in decentralized markets is fair if no arm has justified envy.

Eliminating justified envy is mathematically the same as stability in college admissions under a centralized clearinghouse (Gale and Shapley, 1962). However, there exists a significant difference between stability and eliminating justified envy in decentralized markets. For example, each school is a strategic agent in college admissions, so stability has strategic implications, whereas eliminating justified envy is motivated by fairness considerations. Another difference is that stability in our setting of decentralized markets requires the additional individual rationality condition (13).

Theorem 8 *As $T \rightarrow \infty$, the CDM is fair for arms. That is, no arm has justified envy in the matching process according to CDM.*

We now compare CDM with the oracle arm set, where the latter maximizes agents’ average-case expected payoff by assuming complete information on opponent agents’ strategies. For any arm set $\mathcal{B}_i \subseteq \mathcal{A}^{T+1}$, let $O_{\mathcal{B}_i}$ be the set of true states under which agent P_i ’s strategy of pulling \mathcal{B}_i results in the over-enrollment:

$$O_{\mathcal{B}_i} \equiv \left\{ s_i^* \in [0, 1] \mid \sum_{j \in \mathcal{B}_i} \pi_i(s_i^*, v_j) > q_i \right\}, \quad \forall i \in [m].$$

Under the true state s_i^* , agent P_i ’s expected payoff of pulling \mathcal{B}_i is

$$\mathcal{U}_i[\mathcal{B}_i] = \sum_{j \in \mathcal{B}_i} (v_j + e_{ij}) \cdot \pi_i(s_i^*, v_j) - \gamma_i \cdot \max \left\{ \sum_{j \in \mathcal{B}_i} \pi_i(s_i^*, v_j) - q_i, 0 \right\}.$$

Then P_i ’s average-case expected payoff $\mathbb{E}_{s_i^*} \{\mathcal{U}_i[\mathcal{B}_i]\}$ is

$$\sum_{j \in \mathcal{B}_i} (v_j + e_{ij}) \cdot \mathbb{E}_{s_i^*} [\pi_i(s_i^*, v_j)] - \gamma_i \cdot \mathbb{P}(s_i^* \in O_{\mathcal{B}_i}) \left\{ \sum_{j \in \mathcal{B}_i} \mathbb{E}_{s_i^*} [\pi_i(s_i^*, v_j) \mid s_i^* \in O_{\mathcal{B}_i}] - q_i \right\}.$$

Hence, the oracle arm set \mathcal{B}_i^* for maximizing $\mathbb{E}_{s_i^*}\{\mathcal{U}_i[\mathcal{B}_i]\}$ becomes

$$\mathcal{B}_i^* = \left\{ j \in \mathcal{A}^{T+1} \mid v_j + e_{ij} \geq \gamma_i \cdot \mathbb{P}(s_i^* \in O_{\mathcal{B}_i^*}) \frac{\mathbb{E}_{s_i^*}[\pi_i(s_i^*, v_j) \mid s_i^* \in O_{\mathcal{B}_i^*}]}{\mathbb{E}_{s_i^*}[\pi_i(s_i^*, v_j)]} \right\}. \quad (14)$$

Che and Koh (2016) discussed a particular case of the oracle arm set in a two-agent model. Eq. (14) generalizes the oracle set to multiple agents. Although \mathcal{B}_i^* has a closed-form expression in Eq. (14), \mathcal{B}_i^* is not estimable from training data. This is because the probability $\mathbb{P}(s_i^* \in O_{\mathcal{B}_i^*})$ requires the knowledge of opponent agents' strategies, which are unknown in decentralized markets. Moreover, we show that the matching process according to the strategy of P_i pulling arms from \mathcal{B}_i^* is unfair for arms.

Theorem 9 *The strategy corresponding to the oracle set in Eq. (14) is unfair if for at least one agent P_i , there exists an interval $(v', v'') \subset [0, 1]$ such that*

$$\mathbb{E} \left[\frac{\partial \pi_i(s_i^*, v)}{\partial v} \mid s_i^* \notin O_{\mathcal{B}_i^*} \right] < \mathbb{E} \left[\frac{\partial \pi_i(s_i^*, v)}{\partial v} \mid s_i^* \in O_{\mathcal{B}_i^*} \right] < 0, \quad \forall v \in (v', v''). \quad (15)$$

Here, Eq. (15) means that the probability of arms with score $v \in (v', v'')$ accepting P_i decreases less when P_i is popular (i.e., $s_i^* \in O_{\mathcal{B}_i^*}$) than that when P_i is not popular (i.e., $s_i^* \notin O_{\mathcal{B}_i^*}$). This condition holds for many decentralized matching market models, including the two-agent model of Example 1 in Section 2.3.

Example 2 *Consider the two-agent model of Example 1. Suppose that agent P_2 adopts a cutoff strategy $\sigma_2(v_j, e_{2j}) = \mathbf{1}(e_{2j} \geq \tilde{e}_2(s_1^*, v_j))$, where $\tilde{e}_2(s_1^*, v) \in [0, 1]$ is a cutoff curve under the true state s_1^* . Assume that $\partial \tilde{e}_2(s_1^*, v_j) / \partial v < 0$ for $v_j \in (v', v'')$, which means that the requirement on the fit e_{2j} decreases as the score $v_j \in (v', v'')$ increases. This is a natural assumption in college admissions and holds for many cutoffs, including the cutoff given by Eq. (9). If e_{2j} follows a uniform distribution on $[0, 1]$, then $\mathbb{E}[\sigma_2(v_j, e_{2j})] = 1 - \tilde{e}_2(s_1^*, v_j)$. Recall that Example 1 gives that $\pi_1(s_1^*, v_j) = 1 - \mathbb{E}[\sigma_2(v_j, e_{2j})] + \mu_1(s_1^*)\mathbb{E}[\sigma_2(v_j, e_{2j})]$. In the current example, it is equivalent to*

$$\pi_1(s_1^*, v_j) = \tilde{e}_2(s_1^*, v_j) + \mu_1(s_1^*)(1 - \tilde{e}_2(s_1^*, v_j)).$$

Since $O_{\mathcal{B}_1^*}$ is the set of true states corresponding to the over-enrollments,

$$\mathbb{E}[\mu_1(s_1^*) \mid s_1^* \notin O_{\mathcal{B}_1^*}] < \mathbb{E}[\mu_1(s_1^*) \mid s_1^* \in O_{\mathcal{B}_1^*}].$$

Together with $\partial \tilde{e}_2(s_1^*, v_j) / \partial v < 0$ for $v_j \in (v', v'')$, we have that for any $v_j \in (v', v'')$,

$$\mathbb{E} \left[\frac{\partial \tilde{e}_2(s_1^*, v_j)}{\partial v} (1 - \mu_1(s_1^*)) \mid s_1^* \notin O_{\mathcal{B}_1} \right] < \mathbb{E} \left[\frac{\partial \tilde{e}_2(s_1^*, v_j)}{\partial v} (1 - \mu_1(s_1^*)) \mid s_1^* \in O_{\mathcal{B}_1} \right] < 0,$$

which directly implies Eq. (15) for agent P_1 .

Figure 2 illustrates the unfairness of the strategy corresponding to the oracle set. The proof of Theorem 9 shows that the slope of \mathcal{B}_i^* 's cutoff curve is in the interval $(-1, 0)$, for any $v \in (v', v'')$. Hence, there are arms not selected to \mathcal{B}_i^* (for example, those in $\mathcal{B}_i^{(2)}$).

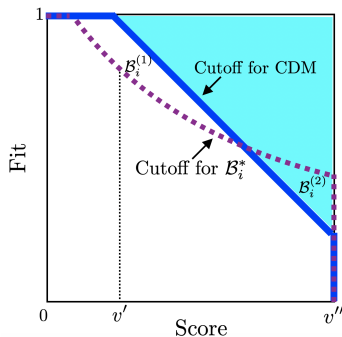


Figure 2: Fairness of CDM compared to the unfairness of oracle arm set \mathcal{B}_i^* . The dotted curve represents the cutoff for \mathcal{B}_i^* in Eq. (14). The solid line segments denote the cutoff for CDM. The arms in $\mathcal{B}_i^{(2)}$ have justified envy towards arms in $\mathcal{B}_i^{(1)}$.

However, they rank higher than some selected arms (for example, those in $\mathcal{B}_i^{(1)}$) according to agent P_i 's true preference. Thus, arms in $\mathcal{B}_i^{(2)}$ have justified envy towards arms in $\mathcal{B}_i^{(1)}$. On the contrary, Theorem 3 shows that the slope of CDM's cutoff equals -1 and hence the agent prefers arms with larger latent utilities.

The intuition behind the results on fairness is as follows. Using the oracle arm set in Eq. (14), agents would respond to congestion strategically to avoid head-on competition under the uncertain state. Specifically, agents strategically favor arms that rank highly in fits against those that rank highly in scores, for such arms are likely to be overlooked by competitors and their acceptance rates are more stable under the uncertain state. Strategic targeting makes the matching process unfair, in the sense of creating justified envy as illustrated in Figure 2. In contrast, using the proposed CDM, agents would firstly calibrate the uncertain state. Under a calibrated state, arms' acceptance rates are predictable, and agents can greedily pull arms according to their true preferences. This straightforward behavior makes the matching process fair.

5. Empirical Study

In this section, we provide a numerical investigation of the fairness and stability properties of CDM. We also study the payoffs achieved by CDM compared to alternative methods.

5.1 Stability and Fairness of CDM

Suppose there are three agents, P_1, P_2, P_3 , $m = 3$, and three arms, A_1, A_2, A_3 , $n = 3$. The latent utilities and the arms' true preferences are given in Table 1. Arms have scores and fits as follows: $v_1 = v_2 = v_3 = 2$, and $e_{11} = e_{21} = e_{32} = 0, e_{13} = e_{22} = e_{31} = 0.5, e_{12} = e_{23} = e_{33} = 1$, respectively. Each agent has quota $q = 1$ and the penalty for exceeding the quota is $\gamma = 10$. Agents have to make decisions on which arms to pull without knowing the arms' true preferences. We compare the CDM procedure with the *greedy action*, which

chooses arms with maximum expected utilities under the total expected acceptance up to the quota (Das and Kamenica, 2005).

| (a) Latent utility | | | | (b) Arm's preference | | | (c) Expected utility | | | | |
|--------------------|-------|-------|-------|----------------------|-------|-------|----------------------|-------|-------|-------|-------|
| | A_1 | A_2 | A_3 | | P_1 | P_2 | P_3 | | A_1 | A_2 | A_3 |
| P_1 | 2 | 3 | 2.5 | A_1 | 3 | 2 | 1 | P_1 | 0.52 | 1.99 | 2.5 |
| P_2 | 2 | 2.5 | 3 | A_2 | 2 | 3 | 1 | P_2 | 0.67 | 0 | 0 |
| P_3 | 2.5 | 2 | 3 | A_3 | 1 | 3 | 2 | P_3 | 2.5 | 2 | 1.05 |

Table 1: (a) Arm's latent utilities for each agent, which corresponds to Eq. (1). For example, P_1 receives utility 2.5 when it successfully pulls A_3 . (b) Arms' preferences with the number indicating the arms' ranking of agents. For example, A_1 ranks P_3 first, P_2 second, P_1 third. These preferences are unknown to agents. (c) Expected utilities, which corresponds to Eq. (2). For example, P_1 expects to receive utility 0.52 if it pulls A_1 .

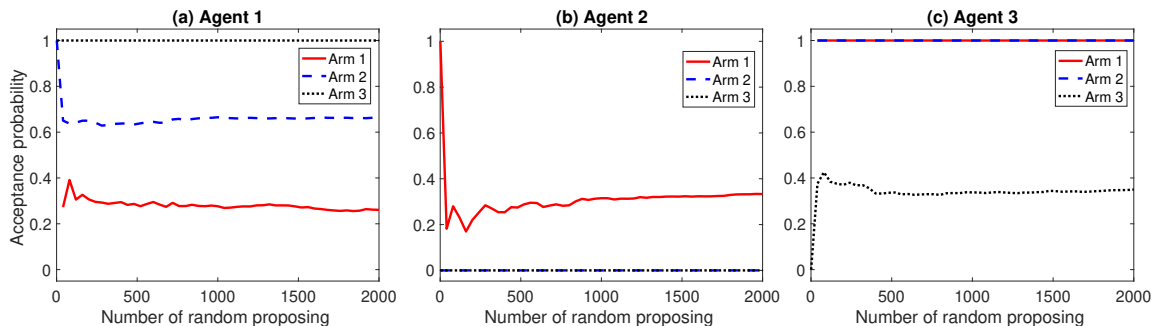


Figure 3: Arms' acceptance probabilities for three agents. (a) P_1 , (b) P_2 , (c) P_3 .

The training data are simulated by having each agent pull a random number of arms according to its latent utilities. Figure 3 shows arms' acceptance probabilities $\pi_i(s_i^*, v_j)$ based on a total of 2000 rounds of random proposals. That is, the sample size is $T = 2000$. There is a unique true state. We implement CDM according to Algorithm 1. The CDM procedure suggests that P_1 pulls A_2 , P_2 pulls A_1, A_2, A_3 , and P_3 pulls A_3 , and it gives the stable matching $(A_1, P_2), (A_2, P_1), (A_3, P_3)$. On the contrary, the greedy action suggests that P_1 pulls A_3 , P_2 pulls A_1, A_2, A_3 , and P_3 pulls A_1 , and it yields the following matching: $(A_1, P_3), (A_2, P_2), (A_3, P_1)$. We note three differences between the two matchings. First, the greedy action is *unfair* since A_2 has justified envy towards A_3 in the sense that A_2 prefers P_1 to P_2 . However, P_1 pulls A_3 that ranks below A_2 according to the true preference of P_1 . Second, the greedy action also yields an *unstable* matching since (A_2, P_1) is a blocking pair. On the other hand, CDM is fair and yields a stable matching. The stable matching is both agent-optimal and arm-optimal in this example. Finally, the total payoff that agents

receive using CDM equals $3 + 2 + 3 = 8$, which is larger than the total payoff that agents receive from the greedy action, $2.5 + 2.5 + 2.5 = 7.5$.

5.2 Lattice Structure for the Stability of CDM

We consider the decentralized matching with four different preference structures: S1, S2, S3, S4. The market consists of three agents, $m = 3$, and three arms, $n = 3$. Each agent has the same quota $q = 1$ and penalty $\gamma = 5$. Table 2 gives arms' latent utilities and true preferences. The training data are simulated by having agents pull random numbers of arms according to their latent utilities. The last column of Table 2 gives the estimates of arms' acceptance probabilities $\pi_i(s_i^*, v_j)$ after 2000 rounds of random proposals under each of the structures S1—S4. That is, the sample size is $T = 2000$ for each of the structures S1—S4. These acceptance probabilities are evaluated at convergence.

We find that CDM in Algorithm 1 gives stable matchings under all of the structures S1—S4. However, these stable matchings have different optimality in terms of agents' and arms' welfare. In particular, the matching in S1 is $(A_1, P_1), (A_2, P_3), (A_3, P_2)$, which is both agent-optimal and arm-optimal; the matching in S2 is $(A_1, P_2), (A_2, P_1), (A_3, P_3)$, which is arm-optimal but not agent-optimal; the matching in S3 is $(A_1, P_2), (A_2, P_3), (A_3, P_1)$, which is not agent-optimal or arm-optimal; the matching in S4 is $(A_1, P_1), (A_2, P_2), (A_3, P_1)$, which is not agent-optimal or arm-optimal. This lattice structure corroborates the results in Section 4.2. We further make three remarks on these matchings. First, some arms benefit from the decentralized matching compared with the centralized matching using the agent-proposing DA algorithm. For example, the CDM matching in S2 is arm-optimal. The agent-proposing DA would give the agent-optimal matching, that is, the arms' least-preferred stable matching.

Second, no strategy in decentralized markets guarantees yielding the agent-optimal matching, mainly due to the competition of agents and a lack of coordination. Consider the S3 structure as an example, where CDM suggests that P_1 pulls A_1, A_2 , P_2 pulls A_2, A_3 and P_3 pulls A_1, A_3 . In this example, the strategy, according to the CDM, is a subgame perfect equilibrium since each agent's action is the best response against other agents' actions. For instance, if P_1 changes to the strategy by only pulling A_2 while other agents' strategies do not change, this gives the following matching: $(A_1, P_2), (A_2, P_3), (A_3, P_2)$, where P_1 is worse off due to the unfilled quota. If P_1 changes to pull all arms A_1, A_2, A_3 while other agents' strategies do not change, then the resulting matching is $(A_1, P_1), (A_2, P_3), (A_3, P_1)$, where P_1 is also worse off due to exceeding the quota. A similar argument applies to P_2 and P_3 . Hence, no agent has an incentive to change its strategy. In contrast, the agent-proposing DA permits each agent to pull one arm each time according to its latent utility. Such coordination results in the agent-optimal matching. However, if there is no coordination, the competition of agents generates uncertainty on the acceptance of the pulled arms. Hence, agents have the incentive to pull more arms than their quotas to combat the uncertainty.

Third, we discuss individual rationality in Eq. (13). Without this condition, any stable matching can be supported by a subgame perfect equilibrium, and only stable matchings can arise in equilibrium (Alcalde and Romero-Medina, 2000). However, under Eq. (13), the set of stable matchings is enlarged. It includes (but may not coincide with) the set of subgame perfect equilibria. Consider S4 as an example, where CDM suggests that P_1

| Latent utility of S1 | | | | Arm's preference of S1 | | | | Acceptance probability of S1 | | | |
|----------------------|-------|-------|-------|------------------------|-------|-------|-------|------------------------------|-------|-------|-------|
| | A_1 | A_2 | A_3 | | P_1 | P_2 | P_3 | | A_1 | A_2 | A_3 |
| P_1 | 2.5 | 3 | 2 | A_1 | 1 | 2 | 3 | P_1 | 1 | 0.34 | 1 |
| P_2 | 3 | 2.5 | 2 | A_2 | 2 | 3 | 1 | P_2 | 0.35 | 0 | 0.65 |
| P_3 | 3 | 2.5 | 2 | A_3 | 1 | 2 | 3 | P_3 | 0 | 1 | 0.45 |

| Latent utility of S2 | | | | Arm's preference of S2 | | | | Acceptance probability of S2 | | | |
|----------------------|-------|-------|-------|------------------------|-------|-------|-------|------------------------------|-------|-------|-------|
| | A_1 | A_2 | A_3 | | P_1 | P_2 | P_3 | | A_1 | A_2 | A_3 |
| P_1 | 3 | 2.5 | 2 | A_1 | 3 | 1 | 2 | P_1 | 0.10 | 1 | 0 |
| P_2 | 2.5 | 3 | 2 | A_2 | 1 | 2 | 3 | P_2 | 1 | 0.35 | 0 |
| P_3 | 2.5 | 2 | 3 | A_3 | 2 | 3 | 1 | P_3 | 0.32 | 0 | 1 |

| Latent utility of S3 | | | | Arm's preference of S3 | | | | Acceptance probability of S3 | | | |
|----------------------|-------|-------|-------|------------------------|-------|-------|-------|------------------------------|-------|-------|-------|
| | A_1 | A_2 | A_3 | | P_1 | P_2 | P_3 | | A_1 | A_2 | A_3 |
| P_1 | 2 | 3 | 2.5 | A_1 | 1 | 2 | 3 | P_1 | 1 | 0.22 | 0.67 |
| P_2 | 2.5 | 2 | 3 | A_2 | 3 | 1 | 2 | P_2 | 0.66 | 1 | 0.24 |
| P_3 | 3 | 2.5 | 2 | A_3 | 2 | 3 | 1 | P_3 | 0.22 | 0.66 | 1 |

| Latent utility of S4 | | | | Arm's preference of S4 | | | | Acceptance probability of S4 | | | |
|----------------------|-------|-------|-------|------------------------|-------|-------|-------|------------------------------|-------|-------|-------|
| | A_1 | A_2 | A_3 | | P_1 | P_2 | P_3 | | A_1 | A_2 | A_3 |
| P_1 | 3 | 2 | 2.5 | A_1 | 3 | 2 | 1 | P_1 | 0.43 | 0.29 | 1 |
| P_2 | 2 | 2.5 | 3 | A_2 | 2 | 1 | 3 | P_2 | 0.69 | 1 | 0 |
| P_3 | 2 | 2.5 | 3 | A_3 | 1 | 3 | 2 | P_3 | 1 | 0.22 | 0.35 |

Table 2: The left column shows arms' latent utilities for each agent. The middle column shows arms' preferences, where the number indicates the arms' ranking of agents. The right column shows arms' acceptance probabilities.

pulls A_1, A_3 , and P_2 pulls A_2, A_3 , and P_3 pulls A_2, A_3 . The strategy given by CDM is not a subgame perfect equilibrium. However, the resulting matching is stable under Eq. (13). For instance, P_3 finds A_1 having a larger expected penalty for exceeding the quota than its expected utility. That is, P_3 finds A_1 unacceptable. Hence, there is no blocking pair.

5.3 Agents' Payoffs

Consider ten agents $m = 10$ and varying numbers of arms, $n = \{50, 70, 90, 110, 130, 150\}$. Each arm has a score v_j and fits e_{ij} drawn uniformly from $[0, 1]$. The agents' preferences are determined by arms' latent utilities as in Eq. (1). Each agent has the same quota

$q = 5$ and the same penalty γ chosen from $\{2, 2.5, 3\}$. The simulation generates random arm preferences with 10 different states from $\{s_1, \dots, s_{10}\} \subset [0, 1]$. The training data are simulated by having agents pull random numbers of arms according to their latent utilities. The training data consists of 20 rounds of random proposals under each of the arms' preference structures. In total, the sample size is $T = 200$. The testing data draws a random state from $\{s_1, \dots, s_{10}\}$ and generates the corresponding arms' preferences. This example simulates top colleges competing for top students. Students' preferences are uncertain and depend on colleges' reputation and popularity in the current year.

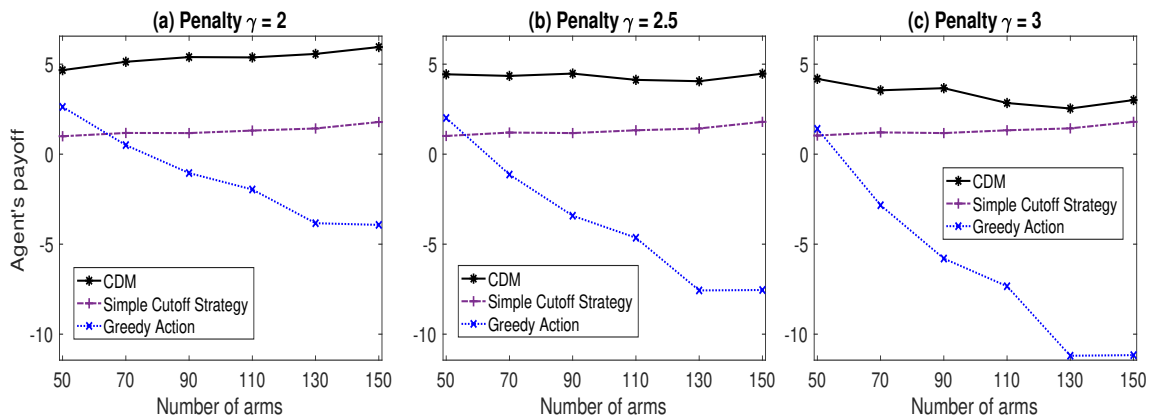


Figure 4: Performance of three strategies with varying numbers of arms. The results are averaged over 500 data replications. Penalty levels (a) $\gamma = 2$, (b) $\gamma = 2.5$, (c) $\gamma = 3$.

We compare the agent's payoff achieved by CDM in Algorithm 1 with that by other methods. In particular, we consider two alternative methods: (i) the *simple cutoff strategy* where the agent chooses q best arms according to the latent utility; (ii) the *greedy action* where the agent chooses arms with maximum expected utilities and a total expected acceptance up to q . Figure 4 reports the agent P_1 's averaged payoffs over 500 data replications. Here, all agents other than P_1 use the CDM with mean calibration and P_1 uses one of the three methods: CDM, simple cutoff strategy, and greedy action. It is seen that CDM gives the largest average payoffs compared to alternative methods, and the advantage of CDM is robust to different numbers of arms and penalty levels. We make two further remarks. First, the state's calibration is useful in improving the agent's utility under uncertain preferences of the arms. For example, CDM outperforms the simple cutoff strategy that has no calibration on the uncertain state. Moreover, CDM helps to ease congestion as it pulls more arms compared to the simple cutoff strategy. Second, CDM performs particularly well if the matching market has intense competition. In this simulated market where arms' preferences are random, a smaller number of arms corresponds to a higher competition level. It is seen that CDM is significantly better than other methods in the regime of small numbers of students. On the other hand, the simple cutoff strategy does not work well with small

numbers of students. The reason is that under an intense competition among the agents, arms reject most offers sent according to the simple cutoff strategy.

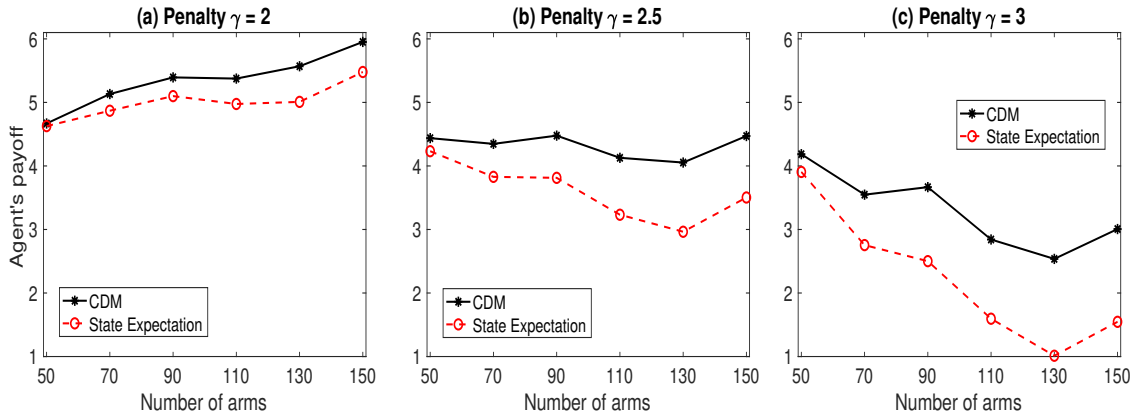


Figure 5: Performance of two calibration methods with varying numbers of arms. The results are averaged over 500 data replications. Penalty levels (a) $\gamma = 2$, (b) $\gamma = 2.5$, (c) $\gamma = 3$.

We also compare two different calibration methods: CDM and *state expectation*. The latter calibrates the unknown state using the naive mean estimate of states, $\mathbb{E}[s_i^*]$, which is discussed in Section 3.3. Figure 5 shows P_1 's averaged payoffs over 500 data replications. Here, all agents other than P_1 use the CDM with mean calibration and P_1 uses one of the two methods: CDM and state expectation. We observe that CDM is adaptive to different levels of penalty γ . However, the state expectation method degrades quickly as γ increases. This difference between the two methods is because CDM balances the marginal payoff and the marginal penalty for exceeding the quota. Hence CDM is sensitive to the penalty γ , while the state expectation method is not.

6. Application to the Study of College Admissions

We demonstrate aspects of the theoretical predictions through a real data from college admissions and a simulated example on graduate school admissions.

6.1 U.S. College Admissions

In this section we present an analysis of admissions data from the *New York Times* “The Choice” blog (<https://thechoice.blogs.nytimes.com/category/admissions-data>). The data set contains admission yields in 2015-17 for 37 U.S. colleges, including liberal arts colleges, national universities, and other undergraduate programs. That is, $m = 37$ and $T = 3$. As we discussed in Section 2.3, a college’s yield is a proxy for the state s_i as an indicator of the college’s popularity. We excluded two colleges, Harvard and Yale, from the sample due to a significant proportion of missing values. We test if the yields of regular admissions

changed over 2015-17. The null hypothesis is that the state is the same over 2015-17. We use a simultaneous chi-squared test for all colleges under false-discovery rate (FDR) control at the .05 significance level (Benjamini and Hochberg, 1995). There are 13 colleges tested with significant p -values. Figure 6 shows that colleges with large numbers of admitted students are more likely to have significantly varied yields compared to colleges with small numbers of admitted students. Moreover, top-ranked national universities and liberal arts colleges are likely to have significantly varied yields. This result corroborates the students' uncertain preferences in college admissions.

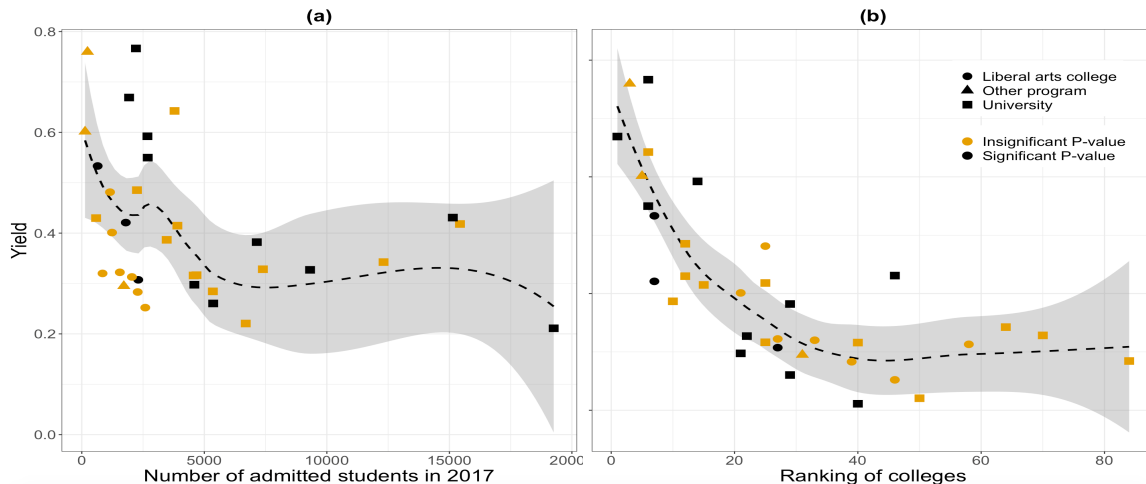


Figure 6: Regression of the yield on (a) the size of admitted class and (b) the ranking of colleges. We fit the dashed curves and the 95% confidence bands using smoothing splines. The ranking is according to *U.S. News and World Report*, where two (or more) colleges might tie, and liberal arts colleges, national universities, and other undergraduate programs are ranked separately within their categories.

6.2 Simulated Graduate School Admissions

We study a simulated graduate school admission problem, where each graduate has a limited quota. Suppose there are a total of $m = 50$ graduate schools from three tiers of colleges: five top colleges $\{P_1, \dots, P_5\}$, ten good colleges $\{P_6, \dots, P_{15}\}$, and 35 other colleges $\{P_{16}, \dots, P_{50}\}$. Each has the same quota $q = 5$ and penalty $\gamma = 2.5$. The simulation generates students' preferences with ten different states $\{s_1, \dots, s_{10}\} \subset [0, 1]$. For any state, students' preferences for colleges from the same tier are random. However, students always prefer top colleges to good colleges, and then the other colleges. The random preferences depend on the state due to colleges' uncertain reputation and popularity in the current year. We consider varying numbers of students $n = \{250, 260, 270, 280, 290, 300\}$. For each number of students, there are ten students having a score v_j chosen uniformly and independently from $[0.9, 1]$ and 100 students having score v_j drawn uniformly and independently from $[0.7, 0.9)$. The rest of the students have score v_j chosen uniformly and indepen-




Figure 7: Performance of three strategies with varying numbers of students. The results are averaged over 500 data replications. (a1) and (a2): College P_1 from tier 1. (b1) and (b2): College P_6 from tier 2. (c1) and (c2): College P_{16} from tier 3.

dently from $[0, 0.7)$. The fits e_{ij} for all college-student pairs are drawn uniformly from $[0, 1]$. The colleges' preferences are determined by students' latent utilities according to Eq. (1). We consider a simple case that students face negligible application costs. Since students do not know how colleges evaluate their fits (e.g., personal essays), submitting applications to all colleges is students' dominant strategy.

We compare the colleges' expected payoffs achieved by three methods: CDM, simple cutoff strategy, and greedy action. The latter two methods are described in Section 5.3. We implement the CDM by Algorithm 1. The training data are simulated from colleges' random proposing, where colleges admit random numbers of students according to their latent utilities. The training data consists of 20 draws of random proposals under each of the students' preference structures. In total, the sample size $T = 200$. The testing data draws a random state from $\{s_1, \dots, s_{10}\}$, which gives the corresponding student preferences. Figure 7 reports the averaged payoffs of three colleges P_1, P_6 and P_{16} over 500 data replications. Here colleges P_1, P_6 and P_{16} belong to the three different tiers, respectively. In Figure

7(a1)—(a2), all colleges other than P_1 use the CDM with mean calibration, while P_1 uses one of the three methods. The same setup applies to Figure 7(b1)—(b2) and Figure 7(c1)—(c2). It is seen that the CDM gives the largest payoffs for P_1, P_6 and P_{16} . The CDM performs particularly well for tier 2 and tier 3 colleges compared to the simple cutoff strategy. Moreover, the CDM outperforms the greedy action, especially for tier 1 and 2 colleges. The reason is that the greedy action does not calibrate the state and enrolls too many students.

7. Discussion

We have presented a statistical framework for decision-making under uncertainty, in settings when preferences are unknown a priori and must be learned from data, and when the output decisions involve the allocation of shared resources. Our approach is decentralized, and thus applicable to agents who may not wish to communicate their preferences or share their decisions. In the model, arms have uncertain preferences that depend on the unknown state of the world and the arms' acceptance probabilities also depend on the competition among the agents. We construct a minimax optimal estimator for the learning of acceptance probabilities under this model. We propose a specific algorithm, the CDM algorithm, for learning optimal strategies from historical data. The CDM algorithm calibrates the unknown state by perturbing the state and balancing the marginal utility and the marginal penalty for exceeding the capacity. This calibration procedure takes opportunity costs into account. We show that CDM makes it safe for agents to act straightforwardly on their preferences. The CDM procedure achieves stability under incomplete information, where the formulation of the individual rationality condition models agents' beliefs. Moreover, CDM is fair for arms in the sense that arms have no justified envy.

It is possible to extend the CDM framework to other competition structures under a scarcity of shared resources. It is of interest to study multi-stage decentralized matching, which has applications on college admissions with a waiting list. However, multi-stage matching involves new challenges. For example, it requires a different model to incorporate the new hierarchical structure arising from the multi-stage decision-making, where arms available at subsequent stages are likely to be worse than those available at the current stage. Hence, the optimal strategy and economic implications in multi-stage matching will be different from those in single-stage matching, and the analysis would be more technical and tedious. We are currently working on that extension. Another possible extension is to consider algorithmic strategies where agents' preferences exhibit complementarities in decentralized matching. For instance, some firms demand workers that complement one another in terms of their skills and roles. Another interesting extension is to consider decentralized markets with indifferent preferences. For example, many applicants may be indistinguishable for a college. Still, it is necessary to break ties since the college may have insufficient capacity to admit all applicants in the same indifference class.

Acknowledgments

The authors thank Robert M. Anderson, two anonymous reviewers, and the Editor for their invaluable feedback on an earlier version of this work.

References

- Atila Abdulkadiroğlu and Tayfun Sönmez. School choice: A mechanism design approach. *American Economic Review*, 93(3):729–747, 2003.
- José Alcalde and Antonio Romero-Medina. Simple mechanisms to implement the core of college admissions problems. *Games and Economic Behavior*, 31(2):294–302, 2000.
- Itai Ashlagi, Mark Braverman, Yash Kanoria, and Peng Shi. Clearing matching markets efficiently: Informative signals and match recommendations. *Management Science*, 66(5):2163–2193, 2020.
- Christopher Avery and Jonathan Levin. Early admissions at selective colleges. *American Economic Review*, 100(5):2125–56, 2010.
- Christopher Avery, Andrew Fairbanks, and Richard J Zeckhauser. *The Early Admissions Game: Joining the Elite*. Harvard University Press, Cambridge, MA, 2003.
- Eduardo M Azevedo and Jacob D Leshno. A supply and demand framework for two-sided matching markets. *Journal of Political Economy*, 124(5):1235–1268, 2016.
- Michel Balinski and Tayfun Sönmez. A tale of two mechanisms: Student placement. *Journal of Economic Theory*, 84(1):73–94, 1999.
- Soumya Basu, Karthik Abinav Sankararaman, and Abishek Sankararaman. Beyond $\log^2(t)$ regret for decentralized bandits in matching markets. *arXiv preprint arXiv:2103.07501*, 2021.
- Yoav Benjamini and Yosef Hochberg. Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society: Series B (Methodological)*, 57(1):289–300, 1995.
- Anand Bhalgat, Ashish Goel, and Sanjeev Khanna. Improved approximation results for stochastic knapsack problems. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1647–1665. SIAM, 2011.
- Lucian Busoniu, Robert Babuska, and Bart De Schutter. A comprehensive survey of multiagent reinforcement learning. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 38(2):156–172, 2008.
- Sarah H Cen and Devavrat Shah. Regret, stability, and fairness in matching markets with bandit learners. *arXiv preprint arXiv:2102.06246*, 2021.
- Hector Chade and Lones Smith. Simultaneous search. *Econometrica*, 74(5):1293–1307, 2006.
- Hector Chade, Gregory Lewis, and Lones Smith. Student portfolios and the college admissions problem. *Review of Economic Studies*, 81(3):971–1002, 2014.
- Yeon-Koo Che and Youngwoo Koh. Decentralized college admissions. *Journal of Political Economy*, 124(5):1295–1338, 2016.

- Pierre-André Chiappori and Bernard Salanié. The econometrics of matching models. *Journal of Economic Literature*, 54(3):832–61, 2016.
- Eugene Choo and Aloysius Siow. Who marries whom and why. *Journal of Political Economy*, 114(1):175–201, 2006.
- Randy Cogill, Michael Rotkowitz, Benjamin Van Roy, and Sanjay Lall. An approximate dynamic programming approach to decentralized control of stochastic systems. In *Control of Uncertain Systems: Modelling, Approximation, and Design*, pages 243–256. Springer, 2006.
- Peter Coles, Alexey Kushnir, and Muriel Niederle. Preference signaling in matching markets. *American Economic Journal: Microeconomics*, 5(2):99–134, 2013.
- Sanmay Das and Emir Kamenica. Two-sided bandits and the dating market. In *Proceedings of the 19th International Joint Conference on Artificial Intelligence*, volume 5, page 19, New York, 2005. AAAI Press.
- Brian C Dean, Michel X Goemans, and Jan Vondrák. Approximating the stochastic knapsack problem: The benefit of adaptivity. *Mathematics of Operations Research*, 33(4):945–964, 2008.
- Dennis Epple, Richard Romano, and Holger Sieg. Admission, tuition, and financial aid policies in the market for higher education. *Econometrica*, 74(4):885–928, 2006.
- Jessie Finocchiaro, Roland Maio, Faidra Monachou, Gourab K Patro, Manish Raghavan, Ana-Andreea Stoica, and Stratis Tsirtsis. Bridging machine learning and mechanism design towards algorithmic fairness. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, pages 489–503, 2021.
- Raymond Fisman, Sheena S Iyengar, Emir Kamenica, and Itamar Simonson. Gender differences in mate selection: Evidence from a speed dating experiment. *The Quarterly Journal of Economics*, 121(2):673–697, 2006.
- Jakob Foerster, Nantas Nardelli, Gregory Farquhar, Triantafyllos Afouras, Philip HS Torr, Pushmeet Kohli, and Shimon Whiteson. Stabilising experience replay for deep multi-agent reinforcement learning. In *International Conference on Machine Learning*, pages 1146–1155. PMLR, 2017.
- Jakob Foerster, Gregory Farquhar, Triantafyllos Afouras, Nantas Nardelli, and Shimon Whiteson. Counterfactual multi-agent policy gradients. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2018.
- Chao Fu. Equilibrium tuition, applications, admissions, and enrollment in the college market. *Journal of Political Economy*, 122(2):225–281, 2014.
- David Gale and Lloyd S Shapley. College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):9–15, 1962.

- Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. On calibration of modern neural networks. In *International Conference on Machine Learning*, pages 1321–1330. PMLR, 2017.
- Guillaume Haeringer and Myrna Wooders. Decentralized job matching. *International Journal of Game Theory*, 40(1):1–28, 2011.
- Isa E Hafalir, Rustamdjan Hakimov, Dorothea Kübler, and Morimitsu Kurino. College admissions with entrance exams: Centralized versus decentralized. *Journal of Economic Theory*, 176:886–934, 2018.
- Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer Science & Business Media, New York, 2009.
- Jinli Hu and Amos Storkey. Multi-period trading prediction markets with connections to machine learning. In *International Conference on Machine Learning*, pages 1773–1781. PMLR, 2014.
- Anna R Karlin and Yuval Peres. *Game theory, alive*, volume 101. American Mathematical Soc., 2017.
- Marc C Kennedy and Anthony O’Hagan. Bayesian calibration of computer models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 63(3):425–464, 2001.
- George Kimeldorf and Grace Wahba. Some results on tchebycheffian spline functions. *Journal of Mathematical Analysis and Applications*, 33(1):82–95, 1971.
- Donald Ervin Knuth. *Stable Marriage and Its Relation to Other Combinatorial Problems: An Introduction to the Mathematical Analysis of Algorithms*, volume 10. CRM Proceedings and Lecture Notes. American Mathematical Society., 1997.
- SangMok Lee. Incentive compatibility of large centralized matching markets. *The Review of Economic Studies*, 84(1):444–463, 2016.
- Joel Z Leibo, Vinicius Zambaldi, Marc Lanctot, Janusz Marecki, and Thore Graepel. Multi-agent reinforcement learning in sequential social dilemmas. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*, pages 464–473, 2017.
- Yi Lin. *Tensor Product Space ANOVA Models in Multivariate Function Estimation*. PhD thesis, University of Pennsylvania, Philadelphia, Pennsylvania, 1998.
- Lydia T Liu, Horia Mania, and Michael I Jordan. Competing bandits in matching markets. In *International Conference on Artificial Intelligence and Statistics*, pages 1618–1628. PMLR, 2020a.
- Lydia T Liu, Feng Ruan, Horia Mania, and Michael I Jordan. Bandit learning in decentralized matching markets. *arXiv preprint arXiv:2012.07348*, 2020b.
- Qingmin Liu, George J Mailath, Andrew Postlewaite, and Larry Samuelson. Stable matching with incomplete information. *Econometrica*, 82(2):541–587, 2014.

- Ryan Lowe, Yi Wu, Aviv Tamar, Jean Harb, Pieter Abbeel, and Igor Mordatch. Multi-agent actor-critic for mixed cooperative-competitive environments. In *Advances in Neural Information Processing Systems*, volume 30. PMLR, 2017.
- Aranyak Mehta. Online matching and ad allocation. *Foundations and Trends in Theoretical Computer Science*, 8(4):265–368, 2013.
- Konrad Menzel. Large matching markets as two-sided demand systems. *Econometrica*, 83(3):897–941, 2015.
- James M Ooi and Gregory W Wornell. Decentralized control of a multiple access broadcast channel: Performance bounds. In *Proceedings of 35th IEEE Conference on Decision and Control*, volume 1, pages 293–298. IEEE, 1996.
- Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. In *Advances in Neural Information Processing Systems*, pages 1177–1184. PMLR, 2008.
- Keith W Ross and Danny HK Tsang. The stochastic knapsack problem. *IEEE Transactions on Communications*, 37(7):740–747, 1989.
- Alvin E Roth. The economics of matching: Stability and incentives. *Mathematics of Operations Research*, 7(4):617–628, 1982.
- Alvin E Roth. The evolution of the labor market for medical interns and residents: A case study in game theory. *Journal of Political Economy*, 92(6):991–1016, 1984.
- Alvin E Roth. New physicians: a natural experiment in market organization. *Science*, 250(4987):1524–1528, 1990.
- Alvin E Roth. Deferred acceptance algorithms: History, theory, practice, and open questions. *International Journal of Game Theory*, 36(3-4):537–569, 2008.
- Alvin E Roth and Marilda Sotomayor. *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*, volume 18. Econometric Society Monographs, Cambridge University Press, Cambridge, 1990.
- Alvin E Roth and Xiaolin Xing. Turnaround time and bottlenecks in market clearing: Decentralized matching in the market for clinical psychologists. *Journal of Political Economy*, 105(2):284–329, 1997.
- Alvin E Roth, Uriel G Rothblum, and John H Vande Vate. Stable matchings, optimal assignments, and linear programming. *Mathematics of Operations Research*, 18(4):803–828, 1993.
- Alessandro Rudi and Lorenzo Rosasco. Generalization properties of learning with random features. In *Advances in Neural Information Processing Systems*, pages 3215–3225. PMLR, 2017.
- Abishek Sankararaman, Soumya Basu, and Karthik Abinav Sankararaman. Dominate or delete: Decentralized competing bandits in serial dictatorship. In *International Conference on Artificial Intelligence and Statistics*, pages 1252–1260. PMLR, 2021.

- Leonard J Savage. *The Foundations of Statistics*. Dover Publications, Inc., New York, 1972.
- Sven Seuken and Shlomo Zilberstein. Formal models and algorithms for decentralized decision making under uncertainty. *Autonomous Agents and Multi-Agent Systems*, 17(2): 190–250, 2008.
- Lloyd S Shapley and Martin Shubik. The assignment game i: The core. *International Journal of Game Theory*, 1(1):111–130, 1971.
- Bernard W Silverman. *Density Estimation for Statistics and Data Analysis*, volume 26. Chapman and Hall, London, 1986.
- Ardi Tampuu, Tambet Matiisen, Dorian Kodelja, Ilya Kuzovkin, Kristjan Korjus, Juhan Aru, Jaan Aru, and Raul Vicente. Multiagent cooperation and competition with deep reinforcement learning. *PloS One*, 12(4):e0172395, 2017.
- Gerald Tesauro. Extending q-learning to general adaptive multi-agent systems. In *Advances in Neural Information Processing Systems*, page None. Citeseer, 2003.
- Grace Wahba. Support vector machines, reproducing kernel hilbert spaces and the randomized gacv. *Advances in Kernel Methods-Support Vector Learning*, pages 69–87, 1999.
- Grace Wahba, Yuedong Wang, Chong Gu, Ronald Klein, and Barbaras Klein. Smoothing spline anova for exponential families, with application to the wisconsin epidemiological study of diabetic retinopathy. *Annals of Statistics*, 23(6):1865–1895, 1995.

Appendix A. Additional Example, Algorithm and Simulation

A.1 Three-Agent Model

We give an additional example of the three-agent model and show the explicit form of the acceptance probability $\pi_i(s_i, v_j), i = 1, 2, 3$ for this model.

Example 3 Consider a three-agent model with agents P_1, P_2 and P_3 . Suppose that P_1 pulls an arm $A_j \in \mathcal{A}$. Denote by $\mu_i(s_1)$ the probability that an arm prefers P_i under the state $s_1, i = 1, 2, 3$, where $\mu_1(s_1) + \mu_2(s_1) + \mu_3(s_1) = 1$. Let σ_i be P_i 's strategy, which is defined by $\sigma_i(v_j, e_{2j}) = \mathbf{1}(P_i \text{ pulls } A_j \text{ with attributes } \{v_j, e_{ij}\}), i = 2, 3$. Note that A_j would be pulled by P_1 with probability $p_1 \equiv [1 - \sigma_2(v_j, e_{2j})][1 - \sigma_3(v_j, e_{3j})]$, and pulled by both P_1 and P_2 with probability $p_2 \equiv \sigma_2(v_j, e_{2j})[1 - \sigma_3(v_j, e_{3j})]$, and pulled by both P_1 and P_3 with probability $p_3 \equiv [1 - \sigma_2(v_j, e_{2j})]\sigma_3(v_j, e_{3j})$, and pulled by all of P_1, P_2 and P_3 with probability $p_4 \equiv \sigma_2(v_j, e_{2j})\sigma_3(v_j, e_{3j})$. Hence A_j would accept P_1 with probability

$$p_1 + \frac{\mu_1(s_1)}{\mu_1(s_1) + \mu_2(s_1)}p_2 + \frac{\mu_1(s_1)}{\mu_1(s_1) + \mu_3(s_1)}p_3 + \mu_1(s_1)p_4.$$

Since e_{ij} are independent and unknown to A_j , and σ_i depends on $e_{ij}, i = 2, 3$, the expected probability that A_j accepts P_1 is as follows,

$$\begin{aligned} \pi_1(s_1, v_j) &= \mathbb{E} \left[p_1 + \frac{\mu_1(s_1)}{\mu_1(s_1) + \mu_2(s_1)}p_2 + \frac{\mu_1(s_1)}{\mu_1(s_1) + \mu_3(s_1)}p_3 + \mu_1(s_1)p_4 \right], \\ &= [1 - \mathbb{E}_{e_{2j}}[\sigma_2(v_j, e_{2j})]] [1 - \mathbb{E}_{e_{3j}}[\sigma_3(v_j, e_{3j})]] \\ &\quad + \frac{\mu_1(s_1)}{\mu_1(s_1) + \mu_2(s_1)} \mathbb{E}_{e_{2j}}[\sigma_2(v_j, e_{2j})] [1 - \mathbb{E}_{e_{3j}}[\sigma_3(v_j, e_{3j})]] \\ &\quad + \frac{\mu_1(s_1)}{\mu_1(s_1) + \mu_3(s_1)} [1 - \mathbb{E}_{e_{2j}}[\sigma_2(v_j, e_{2j})]] \mathbb{E}_{e_{3j}}[\sigma_3(v_j, e_{3j})] \\ &\quad + \mu_1(s_1) \mathbb{E}_{e_{2j}}[\sigma_2(v_j, e_{2j})] \mathbb{E}_{e_{3j}}[\sigma_3(v_j, e_{3j})]. \end{aligned}$$

Moreover, we can derive an explicit form of $\pi_i(s_i, v_j), i = 2, 3$ by following a similar argument as above.

Compared to Example 1 in Section 2.3, the acceptance probability $\pi_1(s_1, v_j)$ in the three-agent model is different from that in the two-agent model, unless $\mu_3(s_1) = 0$. Hence the acceptance probability would change with the number of agents competing for the arm.

A.2 Algorithm of CDM under the Maximin Calibration

We provide the algorithm for the calibrated decentralized matching (CDM) procedure under the maximin calibration in Theorem 5. The CDM under the maximin calibration consists of five main steps. Step 1 is to estimate the acceptance probability based on the historical data, where we obtain the estimate $\hat{\pi}_i(s_i, v) = [1 + \exp(-\hat{f}_i(s_i, v))]^{-1}$ by Eq. (5). Step 2 is to construct the cutoff $\hat{e}_i(s_i, v)$ for a given state s_i according to Eq. (9). Step 3 is to calculate the arm set $\hat{\mathcal{B}}_i(s_i)$ for the given s_i according to Eq. (7). In addition, Steps 2 to 3 are carried out repeatedly over a number of states. Step 4 is to calibrate the state for

maximizing the agent’s worst-case expected payoff, which calibration is referred to as the maximin calibration. Given the estimator $\hat{\pi}_i(s_i, v)$ and the arm set $\hat{\mathcal{B}}_i(s_i)$, the calibration in Theorem 5 can be written the solution to the following equation,

$$\begin{aligned} & \sum_{j \in \hat{\mathcal{B}}_i(s_i)} (v_j + e_{ij}) \cdot [\hat{\pi}_i(1, v_j) - \hat{\pi}_i(0, v_j)] - \gamma_i \sum_{j \in \hat{\mathcal{B}}_i(s_i)} \hat{\pi}_i(1, v_j) + \gamma_i q_i \\ &= \sum_{j \in \hat{\mathcal{B}}_i(1)} (v_j + e_{ij}) \cdot \hat{\pi}_i(1, v_j) - \sum_{j \in \hat{\mathcal{B}}_i(0)} (v_j + e_{ij}) \cdot \hat{\pi}_i(0, v_j). \end{aligned} \quad (16)$$

Finally, Step 5 is to produce the final arm set under the calibrated state. We summarize the CDM under the maximin calibration in Algorithm 2.

Algorithm 2 Calibrated decentralized matching (CDM) under the maximin calibration

- 1: **Input:** Historical data $\mathcal{D} = \{(s_i^t, v_j^t, e_{ij}^t, y_{ij}^t) : i \in [m]; j \in \mathcal{B}_i^t; t \in [T]\}$; New arm set \mathcal{A}^{T+1} with attributes $\{(v_j, e_{ij}) : i \in [m]; j \in [n]\}$ at time $T + 1$; Penalty $\{\gamma_i : i \in [m]\}$ for exceeding the quota.
 - 2: **for** $i = 1, 2, \dots, m$ **do**
 - 3: **Step 1:** Predict the acceptance $\hat{\pi}_i(s_i, v_j)$ by Eq. (5);
 - 4: **for** s_i in a linearly spaced vector in $[0, 1]$ **do**
 - 5: **Step 2:** Construct the cutoff $\hat{e}_i(s_i, v)$ by Eq. (9);
 - 6: **Step 3:** Calculate the arm set $\hat{\mathcal{B}}_i(s_i)$ by Eq. (7);
 - 7: **end for**
 - 8: **Step 4:** Calibrate the state s_i such that the difference between LHS and RHS of Eq. (16) is below a pre-specified tolerance level;
 - 9: **Step 5:** Calculate the arm set $\hat{\mathcal{B}}_i(s_i)$ by Eq. (7) with the calibrated state s_i ;
 - 10: **end for**
 - 11: **Output:** The arm sets $\hat{\mathcal{B}}_1(s_1), \hat{\mathcal{B}}_2(s_2), \dots, \hat{\mathcal{B}}_m(s_m)$ for agents.
-

A.3 Sensitivity Analysis

We next carry out a sensitivity analysis of the CDM with respect to the sample size of training data. We adopt the same simulation setup as in Section 5.3 with $m = 10$ agents and $n = 150$ arms. Each arm has a score v_j and fits e_{ij} drawn uniformly from $[0, 1]$. The agents’ preferences are determined by arms’ latent utilities as in Eq. (1). Each agent has the same quota $q = 5$ and the same penalty γ chosen from $\{2, 2.5, 3\}$. The simulation generates random arm preferences with 10 different states from $\{s_1, \dots, s_{10}\} \subset [0, 1]$. The training data are simulated by having agents pull random numbers of arms according to their latent utilities. We vary the sample size of training data $T = \{100, 150, 200, 250, 300\}$, where the training data consists of $T/10$ rounds of random proposals under each of the arms’ preference structures. The testing data draws a random state from $\{s_1, \dots, s_{10}\}$ and generates the corresponding arms’ preferences. We compare the CDM with alternative methods: (i) the *simple cutoff strategy* where the agent chooses q best arms according to the latent utility; (ii) the *greedy action* where the agent chooses arms with maximum expected utilities and a total expected acceptance up to q ; (iii) the *state expectation* that calibrates the unknown state using the naive mean estimate of states.




Figure 8: Performance of three strategies with varying numbers of sample size. The results are averaged over 500 data replications. Penalty levels (a) $\gamma = 2$, (b) $\gamma = 2.5$, (c) $\gamma = 3$.

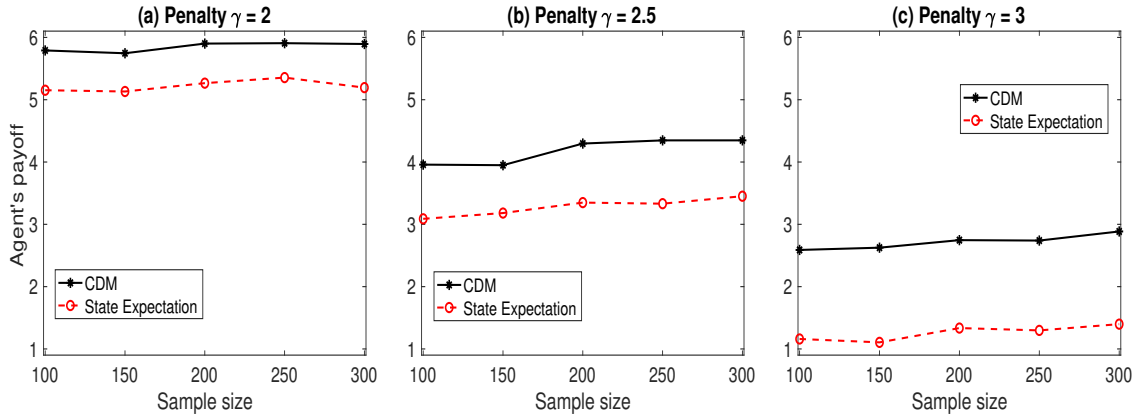


Figure 9: Performance of two calibration methods with varying numbers of sample size. The results are averaged over 500 data replications. Penalty levels (a) $\gamma = 2$, (b) $\gamma = 2.5$, (c) $\gamma = 3$.

Figures 8 and 9 report the agent P_1 's averaged payoffs with the varying sample sizes T based on 500 data replications. Here, all agents other than P_1 use the CDM with mean calibration in Algorithm 1. On the other hand, P_1 uses one of the four methods: CDM, simple cutoff strategy, greedy action, and state expectation. It is seen that the CDM consistently outperforms alternative methods. Moreover, the CDM performs relatively robust with respect to the sample size T .

Appendix B. Technical Proofs

B.1 ANOVA Decomposition for General Utility Function

Suppose that each arm $A_j \in \mathcal{A}$ is described by x_j and $(\epsilon_{1j}^\top, \epsilon_{2j}^\top, \dots, \epsilon_{mj}^\top)^\top$, where x_j is a multidimensional vector available to all agents and ϵ_{ij} is a multidimensional vector that is only available to agent P_i . In college admissions, x_j can be student A_j 's high school record and test score on a nationwide exam such as SAT/ACT, and ϵ_{ij} represents student A_j 's college-specific essays or test scores.

Proposition 10 *For any agent-specific latent utility function U_i , the ANOVA decomposition allows the separable form*

$$U_i(A_j) = v_j + e_{ij}, \quad \forall i \in [m] \text{ and } j \in [n].$$

Here, $v_j \in \mathbb{R}$ is a function of x_j and is common to all agents. The agent-specific $e_{ij} \in \mathbb{R}$ is a function of x_j and ϵ_{ij} and is considered only by agent P_i . Thus, the separable structure of the utility function in Eq. (1) can be assumed without loss of generality.

Proof Denote the utility function $U_i(A_j) \equiv U_i(x_j, \epsilon_{ij}) \in \mathbb{R}$. By the analysis of variance (ANOVA) decomposition, we have that for $P_i \in \mathcal{P}, A_j \in \mathcal{A}$,

$$\begin{aligned} v_j &\equiv \frac{1}{m} \sum_{i=1}^m \mathbb{E}_{\epsilon_{ij}} [U_i(x_j, \epsilon_{ij})] \quad \text{and} \\ e_{ij}^\dagger &\equiv \mathbb{E}_{\epsilon_{ij}} [U_i(x_j, \epsilon_{ij})] - \frac{1}{m} \sum_{i=1}^m \mathbb{E}_{\epsilon_{ij}} [U_i(x_j, \epsilon_{ij})], \quad e_{ij}^\ddagger \equiv U_i(x_j, \epsilon_{ij}) - \mathbb{E}_{\epsilon_{ij}} [U_i(x_j, \epsilon_{ij})]. \end{aligned} \tag{17}$$

Here, v_j represents the average utility of x_j and is common to all agents. The e_{ij}^\dagger is agent P_i 's adjustment for the utility of x_j . The e_{ij}^\ddagger is the utility of ϵ_{ij} received by agent P_i . Thus, e_{ij}^\dagger and e_{ij}^\ddagger are agent-specific and they are only known to agent P_i . Moreover, letting $e_{ij} \equiv e_{ij}^\dagger + e_{ij}^\ddagger$, then Eq. (17) implies the desired result that $U_i(x_j, \epsilon_{ij}) = v_j + e_{ij}$. We refer to Figure 10 for an illustration on the ANOVA decomposition.

The utility function $U_i(x_j, \epsilon_{ij})$ is generally assumed to be strictly increasing in x_j and nondecreasing in ϵ_{ij} (Che and Koh, 2016; Lee, 2016). Then v_j in Eq. (17) is strictly increasing in x_j , and e_{ij} is nondecreasing in ϵ_{ij} , for any $i \in [m]$. \blacksquare

We demonstrate the ANOVA decomposition in Eq. (17) through the college admission example. The score v_j in Eq. (17) represents the ‘‘public valuation’’ of a high school’s quality, high school GPA, and SAT/ACT score. Although most colleges place the highest importance on academic achievement in evaluating applications, each factor’s specific weight can differ from college to college. Moreover, ‘‘holistic admission’’—such that a high SAT/ACT score and a high GPA is no guarantee of admission—is not rare in college admissions, especially for top colleges. The decomposition in Eq. (17) incorporates the college-specific weight for students’ academic performance and extracurricular activity to terms e_{ij}^\dagger and e_{ij}^\ddagger . The term e_{ij}^\dagger in Eq. (17) can represent college-specific adjustment, which is ‘‘private valuation’’ of a student’s high school record and SAT/ACT score, in addition to special talent,

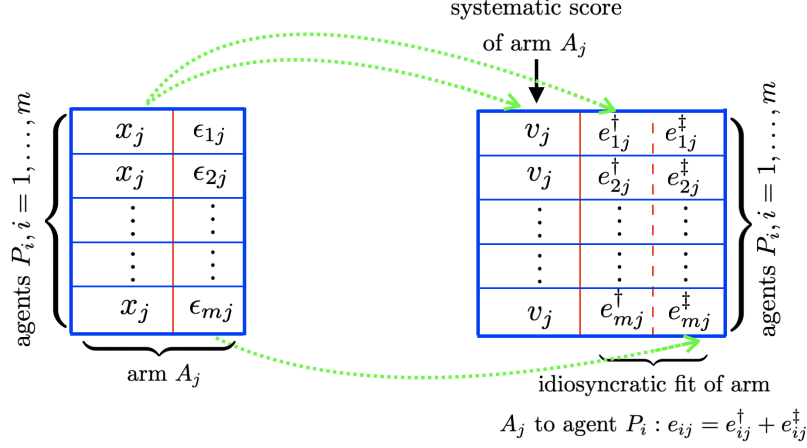


Figure 10: ANOVA decomposition of the utility function corresponding to Eq. (17).

grades of challenging college preparatory curriculum and work experience. For example, some colleges may place a larger weight on the SAT/ACT score than other colleges. The term e_{ij}^\dagger in Eq. (17) is a college-specific “private valuation” of students’ writing skills and compelling personal stories.

B.2 Proof of Theorem 1: Acceptance Probability

Proof The uncertainty of the arm’s acceptance mainly consists of two parts: the state of the world and the agents’ strategies. We employ a four-step construction approach to show that there exists a probability mass function (PMF) $\pi_i(s_i, v_j)$ characterizing the probability of A_j accepting P_i . First, Nature draws the state ω and arms’ preferences for agents that are characterized by ω . Then agent-specific states are determined as a function of ω : $s_i = s_i(\omega)$ for $i \in [m]$.

Second, we derive an arm’s acceptance probability. Suppose now agent P_i pulls arm A_j . Since A_j would accept its most preferred agent among those who have pulled it, A_j ’s acceptance of P_i depends on other agents’ strategies. Let $I \subseteq [n]$ and $P_I \equiv \{P_k, k \in I\}$, which is the set of agents that pull arm A_j . We define

$$\mu_{k, I \cup \{i\}}(\omega, v_j, \mathbf{e}_j) \equiv \mathbb{P}(A_j \text{ accepts } P_k \mid P_{I \cup \{i\}} \text{ pulls } A_j), \quad k \in I \cup \{i\}.$$

That is, $\mu_{k, I \cup \{i\}}(\omega, v_j, \mathbf{e}_j)$ is the PMF that an arm with the same score and fits as A_j would accept P_k conditional on the event that agents $P_{I \cup \{i\}}$ have pulled the arm. The $\mu_{k, I \cup \{i\}}(\omega, v_j, \mathbf{e}_j)$ depends on arms’ preferences as characterized by ω , and also on agents’ strategies determined by the score v_j and the vector of fits $\mathbf{e}_j = (e_{1j}, \dots, e_{mj})$. Moreover, the PMF $\mu_{k, I \cup \{i\}}(\omega, v_j, \mathbf{e}_j)$ satisfies that,

$$\sum_{k \in I \cup \{i\}} \mu_{k, I \cup \{i\}}(\omega, v_j, \mathbf{e}_j) = 1, \quad \forall j \in [n]. \quad (18)$$

Third, we derive an arm's acceptance probability from an agent P_i 's perspective. Suppose again that agent P_i pulls arm A_j . Define the PMF $\pi_i(s_i, v_j)$ as,

$$\pi_i(s_i, v_j) \equiv \mathbb{P}(A_j \text{ accepts } P_i \mid P_i \text{ pulls } A_j).$$

Let $\mathcal{I}_{-i} \equiv \{I : I \subseteq [n] \setminus \{i\}\}$ be the family of subsets of $[n] \setminus \{i\}$. Then by Tonelli's theorem, we can write

$$\begin{aligned} \pi_i(s_i, v_j) &= \mathbb{E} \left[\mathbb{P}(A_j \text{ accepts } P_i \mid P_{I \cup \{i\}} \text{ pulls } A_j, P_{[n] \setminus (I \cup \{i\})} \text{ do not pull } A_j) \right] \\ &= \mathbb{E}_{I \in \mathcal{I}_{-i}} \left[\mathbb{E}_{\mathbf{e}_{-i,j}} [\mu_{i, I \cup \{i\}}(\omega, v_j, \mathbf{e}_j) \cdot \mathbf{1}(P_I \text{ pulls } A_j) \cdot \mathbf{1}(P_{[n] \setminus (I \cup \{i\})} \text{ do not pull } A_j)] \right], \end{aligned} \quad (19)$$

where $\mathbf{e}_{-i,j} = (e_{1j}, \dots, e_{i-1,j}, e_{i+1,j}, \dots, e_{nj})$. When conditional on v_j , $\mathbf{1}(P_I \text{ pulls } A_j)$ only depends on $\{e_k, k \in I\}$, and similarly, $\mathbf{1}(P_{[n] \setminus (I \cup \{i\})} \text{ do not pull } A_j)$ only depends on $\{e_k, k \in [n] \setminus (I \cup \{i\})\}$. Then the defined PMF $\pi_i(s_i, v_j)$ represents the PMF $\mu_{i, I \cup \{i\}}(\omega, v_j, \mathbf{e}_j)$ being averaged over agents P_I ' strategies. Since π_i is defined by conditioning on P_i 's strategy (i.e., P_i pulls A_j), the probability $\pi_i(s_i, v_j)$ does not depend on the fit e_{ij} . Moreover, if knowing that agents $P_{I \cup \{i\}}$ pull arm A_j , then by Eqs. (18) and (19), the PMF $\pi_i(s_i, v_j)$ satisfies that,

$$\sum_{k \in I \cup \{i\}} \pi_k(s_k, v_j) = 1, \quad \forall j \in [n].$$

Fourth, we note that agent-specific states $s_i = s_i(\omega)$ is chosen by agent P_i for $i \in [m]$. Hence without loss of generality, let s_i is parameterized such that the PMF $\pi_i(s_i, v_j)$ is increasing in s_i .

Finally, the expected utility that agent P_i receives from pulling arm $A_j \in \mathcal{A}$ is

$$\begin{aligned} \mathbb{E}[\text{utility}] &= \mathbb{E}[\text{utility} \mid \text{successful pulling}] \cdot \mathbb{P}(\text{successful pulling}) \\ &= (v_j + e_{ij}) \cdot \pi_i(s_i, v_j). \end{aligned}$$

This completes the proof. ■

Remark 11 *We now compare the functions $\mu_{i, I \cup \{i\}}(\omega, v_j, \mathbf{e}_j)$ and $\pi_i(s_i, v_j)$ from a learning perspective. It is shown in the above proof that the function $\mu_{i, I \cup \{i\}}(\omega, v_j, \mathbf{e}_j)$ captures the distribution of arms' preferences. However, $\mu_{i, I \cup \{i\}}(\omega, v_j, \mathbf{e}_j)$ is imperfect for learning using data due to two reasons. First, if the agent P_i wants to estimate $\mu_{i, I \cup \{i\}}(\omega, v_j, \mathbf{e}_j)$, P_i has to identify other agents who are also pulling A_j , that is, to identify the set I . However, thanks to the nature of decentralized markets, there is little communication among agents, and P_i cannot learn about other agents' strategies. Second, note that $\mu_{i, I \cup \{i\}}(\omega, v_j, \mathbf{e}_j)$ depends on the fit vector $\mathbf{e}_{-i,j} = (e_{1j}, \dots, e_{i-1,j}, e_{i+1,j}, \dots, e_{nj})$. However, since the fits are private information, the vector $\mathbf{e}_{-i,j}$ is also unknown to agent P_i .*

In contrast, the learning of function $\pi_i(s_i, v_j)$ does not require P_i knowing other agents' strategies. Moreover, $\pi_i(s_i, v_j)$ does not depend on the fit vector $\mathbf{e}_{-i,j}$. Therefore, the function $\pi_i(s_i, v_j)$ is favorable than the function $\mu_{i, I \cup \{i\}}(\omega, v_j, \mathbf{e}_j)$ in the sense that agent P_i can use historical data to learn $\pi_i(s_i, v_j)$.

B.3 Proof of Theorem 2: Optimal Estimation of Acceptance

Proof Let \tilde{f}_i be the minimizer of Eq. (3), that is,

$$\tilde{f}_i = \arg \min_{f_i \in \mathcal{H}_{K_i}} \left\{ \sum_{t=1}^T \sum_{j \in \mathcal{B}_i^t} [-y_{ij}^t f_i(s_i^t, v_j^t) + \log(1 + \exp(f_i(s_i^t, v_j^t)))] + \frac{1}{2} \sum_{t=1}^T n_{it} \lambda_i \|f_i\|_{\mathcal{H}_{K_i}}^2 \right\}.$$

By the results in Chapters 5 of Lin (1998), we obtain that

$$\mathbb{E}_{s_i, v}[(\tilde{f}_i - f_i)^2] \leq c_1 [T(\log T)^{-1}]^{-2r/(2r+1)} \quad \text{as } T \rightarrow \infty,$$

where $\lambda_i \leq c_\lambda [T(\log T)^{-1}]^{-2r/(2r+1)}$. Here, $c_\lambda, c_1 > 0$ are constants independent of T . Moreover, the estimate \tilde{f}_i is minimax rate-optimal. By the generalization properties of random features (Rudi and Rosasco, 2017), it is known that if the number of random features satisfies

$$p \geq c_p [T(\log T)^{-1}]^{-2r/(2r+1)},$$

then there exists some constant $c_2 > 0$ such that

$$\mathbb{E}_{s_i, v}[(\hat{f}_i - \tilde{f}_i)^2] \leq c_2 [T(\log T)^{-1}]^{-2r/(2r+1)} \quad \text{as } T \rightarrow \infty.$$

By the triangle inequality, there exists some constant $c_f > 0$ such that

$$\begin{aligned} \mathbb{E}_{s_i, v}[(\hat{f}_i - f_i)^2] &\leq \mathbb{E}_{s_i, v}[(\hat{f}_i - \tilde{f}_i)^2] + \mathbb{E}_{s_i, v}[(\tilde{f}_i - f_i)^2] \\ &\leq c_f [T(\log T)^{-1}]^{-2r/(2r+1)} \quad \text{as } T \rightarrow \infty. \end{aligned}$$

Therefore, the estimate \hat{f}_i in Eq. (4) is minimax rate-optimal. ■

B.4 Proof of Theorem 3: Cutoff Strategy

Proof We divide the proof of this theorem into five steps.

Step 1. We show that the cutoff strategy with respect to the fits is optimal. Suppose that arms $A_{j_1}, A_{j_2} \in \mathcal{A}^{T+1}$ have the same score $v_{j_1} = v_{j_2}$, but A_{j_1} has a worse fit than A_{j_2} for agent P_i . Now assume that A_{j_1} was pulled by P_i but A_{j_2} was not, that is, $A_{j_1} \in \hat{\mathcal{B}}_i(s_i), A_{j_2} \notin \hat{\mathcal{B}}_i(s_i)$. Then the expected number of arms accepting P_i is unchanged if P_i replaces A_{j_1} with A_{j_2} in $\hat{\mathcal{B}}_i(s_i)$. On the other hand, since the P_i 's expected payoff in Eq. (6) is increasing in fit e_{ij} , P_i has a strictly larger expected payoff if P_i replaces A_{j_1} with A_{j_2} . Hence, P_i should pull A_{j_2} instead A_{j_1} . This argument holds regardless of strategies of other agents.

Step 2. We prove that the cutoff $\hat{e}_i(s_i, v)$ in Eq. (9) is well-defined. If the boundary $\{\mathcal{B}_i^+(s_i) \setminus \mathcal{B}_i^-(s_i)\}$ is not empty, the expected penalty due to exceeding the quota is

$$\gamma_i \sum_{j \in \mathcal{B}_i^+(s_i)} \pi_i(s_i, v_j) - \gamma_i q_i.$$

The expected utility of pulling arms on the boundary is

$$\sum_{j \in \mathcal{B}_i^+(s_i) \setminus \mathcal{B}_i^-(s_i)} (v_j + e_{ij}) \cdot \pi_i(s_i, v_j).$$

Agent P_i pulls an arm if the expected utility is at least the expected penalty, which justifies the condition specified by Eq. (8). Moreover, since $\widehat{c}_i(s_i, v) \in [0, 1]$, we conclude that the cutoff is well-defined.

Step 3. We show that the cutoff strategy of pulling arms from the set $\widehat{\mathcal{B}}_i(s_i)$ is near-optimal. Let $\widetilde{\mathcal{B}}_i(s_i) \subseteq \mathcal{A}^{T+1}$ be any other arm set. Define the following mixed strategy:

$$\sigma_{i,k}(s_{i,k}, v, e_i; t) \equiv t \cdot \mathbf{1}\{(v, e_i) \in \widetilde{\mathcal{B}}_i(s_i)\} + (1-t) \cdot \mathbf{1}\{(v, e_i) \in \widehat{\mathcal{B}}_i(s_i)\}, \quad \text{for } t \in [0, 1].$$

The expected payoff of using the mixed strategy σ_i is

$$\begin{aligned} \mathbb{U}_i(t) \equiv & \sum_{j \in \mathcal{A}^{T+1}} (v_j + e_{ij}) \cdot \pi_i(s_i, v_j) \cdot \sigma_i(s_i, v_j, e_{ij}; t) \\ & - \gamma_i \cdot \max \left\{ \sum_{j \in \mathcal{A}^{T+1}} \pi_i(s_i, v_j) \cdot \sigma_i(s_i, v_j, e_{ij}; t) - q_i, 0 \right\}. \end{aligned}$$

It is clear that $\mathbb{U}_i(t)$ is concave in t . We discuss the local change $d\mathbb{U}_i(0)/dt$ in three cases.

Case (I). Consider removing a single arm from $\widehat{\mathcal{B}}_i(s_i)$. If the arm is from the non-empty boundary $\{\mathcal{B}_i^+(s_i) \setminus \mathcal{B}_i^-(s_i)\}$, the condition specified in Eq. (8) implies that P_i 's expected payoff will decrease if not pulling the arm. Moreover, since removing any other arms from $\widehat{\mathcal{B}}_i(s_i)$ results in a greater loss of P_i 's expected payoff compared to removing the arms on the non-empty boundary, we have $d\mathbb{U}_i(0)/dt < 0$ in this case. By the concavity of $\mathbb{U}_i(t)$ in t , we obtain

$$\mathbb{U}_i(1) = \mathbb{U}_i(0) + \frac{d\mathbb{U}_i(0)}{dt}(1-0) < \mathbb{U}_i(0),$$

Case (II). Consider adding a new arm to $\widehat{\mathcal{B}}_i(s_i)$, where the new arm has attributes $\{v_{j'}, e_{ij'}\}$ and it is not from the set $\mathcal{B}_i^+(s_i)$. Denote by $\mathcal{B}'_i(s_i)$ the new arm set with the added arm. Note that P_i pulls a new arm only if the arm has a larger expected utility than the expected penalty due to exceeding the quota, that is,

$$(v_{j'} + e_{ij'}) \cdot \pi(s_i, v_{j'}) \geq \gamma_i \sum_{j \in \mathcal{B}'_i(s_i)} \pi_i(s_i, v_j) - \gamma_i q_i. \quad (20)$$

Since the new arm is not in $\mathcal{B}_i^+(s_i)$, we have

$$\begin{aligned} & \sum_{j \in \mathcal{B}'_i(s_i)} \pi_i(s_i, v_j) - q_i \\ = & \sum_{j \in \mathcal{B}'_i(s_i)} \pi_i(s_i, v_j) - \sum_{j \in \mathcal{B}_i^+(s_i)} \pi_i(s_i, v_j) + \sum_{j \in \mathcal{B}_i^+(s_i)} \pi_i(s_i, v_j) - q_i \\ \geq & \sum_{j \in \mathcal{B}'_i(s_i)} \pi_i(s_i, v_j) - \sum_{j \in \mathcal{B}_i^+(s_i)} \pi_i(s_i, v_j) \\ \geq & \pi_i(s_i, v_{j'}). \end{aligned} \quad (21)$$

Because that $\gamma_i > \sup_{j \in \mathcal{A}^{T+1}} \{v_j + e_{ij}\}$, Eq. (21) is contradictory to Eq. (20). Hence, adding a new arm to $\widehat{\mathcal{B}}_i(s_i)$ induces a loss in P_i 's expected payoff. Hence, $d\mathbb{U}_i(0)/dt < 0$ in this case. By the concavity of $\mathbb{U}_i(t)$ in t , we obtain

$$\mathbb{U}_i(1) = \mathbb{U}_i(0) + \frac{d\mathbb{U}_i(0)}{dt}(1-0) < \mathbb{U}_i(0),$$

Case (III): Consider removing an arm with attributes (v_j, e_{ij}) from $\widehat{\mathcal{B}}_i(s_i)$ and simultaneously adding new arms to $\widehat{\mathcal{B}}_i(s_i)$. Suppose that the new arms are from $\mathcal{B}_i''(s_i)$ with attributes $(v_{j''}, e_{ij''})$. If $\widehat{\mathcal{B}}_i(s_i) = \mathcal{B}_i^-(s_i)$, then the new arms are not in $\mathcal{B}_i^-(s_i)$. By definition of $\mathcal{B}_i^-(s_i)$, we have $v_{j''} + e_{ij''} \leq \min_{j \in \mathcal{B}_i^-(s_i)} (v_j + e_{ij})$. Hence,

$$\begin{aligned} \mathbb{U}_i(1) - \mathcal{U}_i[\mathcal{B}_i^-(s_i)] &\leq \sum_{j'' \in \mathcal{B}_i''(s_i)} (v_{j''} + e_{ij''}) \pi_i(s_i, v_{j''}) \\ &\leq \left[\min_{j \in \mathcal{B}_i^-(s_i)} (v_j + e_{ij}) \right] \cdot \left[q_i - \sum_{j \in \mathcal{B}_{i,k}^-(s_i,k)} \pi_{i,k}(s_{i,k}, v_j) \right] \\ &= \text{UE}^\dagger. \end{aligned} \tag{22}$$

If $\widehat{\mathcal{B}}_{i,k}(s_{i,k}) = \mathcal{B}_{i,k}^+(s_{i,k})$, by definition of $\mathcal{B}_{i,k}^+(s_{i,k})$

$$\mathbb{U}_i(1) - \mathcal{U}_i[\mathcal{B}_i^+(s_i)] \leq \mathbb{U}_i(1) - \mathcal{U}_i[\mathcal{B}_i^-(s_i)] \leq \text{UE}^\dagger.$$

where the last inequality is by Eq. (22). Hence,

$$\mathbb{U}_i(1) - \mathbb{U}_i(0) \leq \text{UE}^\dagger.$$


Therefore, exchanging an arm in $\widehat{\mathcal{B}}_{i,k}(s_{i,k})$ with arms not in $\widehat{\mathcal{B}}_{i,k}(s_{i,k})$ could result in a decrease in the expected payoff $\mathbb{U}_i(t)$ by at most UE^\dagger . Combining the cases (I), (II), (III), we obtain that

$$\max_{\mathcal{B}_i \subseteq \mathcal{A}^{T+1}} \mathcal{U}_i[\mathcal{B}_i] - \mathcal{U}_i[\widehat{\mathcal{B}}_i(s_i)] \leq \text{UE}^\dagger.$$

Step 4. We prove the other direction of the inequality. Since $\widehat{\mathcal{B}}_i(s_i) \subseteq \mathcal{A}^{T+1}$,

$$\max_{\mathcal{B}_i \subseteq \mathcal{A}^{T+1}} \mathcal{U}_i[\mathcal{B}_i] - \mathcal{U}_i[\widehat{\mathcal{B}}_i(s_i)] \geq 0.$$

Step 5. If there is a continuum of arms and $\pi_i(\cdot, v)$ is continuous in v , then there exists $b_i \geq 0$ such that $\Pi_{i,k}(b_{i,k}) = q_i$, where $\Pi_{i,k}(b_{i,k})$ is defined in Section 3.2. By definition of $\widehat{\mathcal{B}}_i(s_i)$ we have, $\widehat{\mathcal{B}}_i(s_i) = \mathcal{B}_i^+(s_i) = \mathcal{B}_i^-(s_i)$, and $\sum_{j \in \mathcal{B}_i^-(s_i)} \pi_i(s_i, v_j) = q_i$. Hence $\text{UE}^\dagger = 0$. This completes the proof. \blacksquare


 Figure 11: Cost of a strategy in the face of uncertain true state s_i^* .

B.5 Proof of Theorem 4: Mean Calibration for CDM

Proof By the proof of Theorem 3, $\widehat{\mathcal{B}}_i(s_i) \subseteq \widehat{\mathcal{B}}_i(s_i - \delta s_i)$ for any $\delta s_i \in (0, s_i)$. Hence, the marginal set $\partial \widehat{\mathcal{B}}_i(s_i)$ is well-defined. Let $V_i(s_i^*, \widehat{\mathcal{B}}_i(s_i))$ be the expected utility that P_i receives by pulling arms from $\widehat{\mathcal{B}}_i(s_i)$ and under the true state s_i^* . That is,

$$V_i(s_i^*, \widehat{\mathcal{B}}_i(s_i)) \equiv \sum_{j \in \widehat{\mathcal{B}}_i(s_i)} (v_j + e_{ij}) \cdot \pi_i(s_i^*, v_j).$$

Let $\mathcal{N}_i(s_i^*, \widehat{\mathcal{B}}_i(s_i))$ be the expected number of arms in $\widehat{\mathcal{B}}_i(s_i)$ accepting P_i under s_i^* . That is,

$$\mathcal{N}_i(s_i^*, \widehat{\mathcal{B}}_i(s_i)) \equiv \sum_{j \in \widehat{\mathcal{B}}_i(s_i)} \pi_i(s_i^*, v_j).$$

Similarly, we define $V_i(s_i^*, \partial \widehat{\mathcal{B}}_i(s_i))$ and $\mathcal{N}_i(s_i^*, \partial \widehat{\mathcal{B}}_i(s_i))$ for the marginal set $\partial \widehat{\mathcal{B}}_i(s_i)$. Let the marginal utility be

$$u_i(s_i^*, \partial \widehat{\mathcal{B}}_i(s_i)) \equiv \frac{V_i(s_i^*, \partial \widehat{\mathcal{B}}_i(s_i))}{\mathcal{N}_i(s_i^*, \partial \widehat{\mathcal{B}}_i(s_i))}.$$

Since the true state s_i^* is unknown, the cost of pulling the arm set $\widehat{\mathcal{B}}_i(s_i)$ consists of two parts. See an illustration in Figure 11. The first is the *over-enrollment* cost (OE), which occurs if the calibration parameter $s_i < s_i^*$. Then the realized number of arms in $\widehat{\mathcal{B}}_i(s_i)$ accepting P_i will be greater than the expected number of arms in $\widehat{\mathcal{B}}_i(s_i)$ accepting P_i . That is, $\mathcal{N}_i(s_i^*, \widehat{\mathcal{B}}_i(s_i)) > \mathcal{N}_i(s_i, \widehat{\mathcal{B}}_i(s_i)) = \mathcal{N}_i(s_i^*, \widehat{\mathcal{B}}_i(s_i^*)) = q_i$. Thus, OE depends on s_i and can be written as

$$\begin{aligned} & \text{OE}(s_i) \\ & \equiv \mathbb{E}_{s_i^*} \left[\gamma_i \left\{ \mathcal{N}_i(s_i^*, \widehat{\mathcal{B}}_i(s_i)) - \mathcal{N}_i(s_i^*, \widehat{\mathcal{B}}_i(s_i^*)) \right\} - \left\{ V_i(s_i^*, \widehat{\mathcal{B}}_i(s_i)) - V_i(s_i^*, \widehat{\mathcal{B}}_i(s_i^*)) \right\} \mid s_i < s_i^* \leq 1 \right] \\ & = \mathbb{E}_{s_i^*} \left[\int_{t=s_i}^{s_i^*} [\gamma_i - u_i(s_i^*, \partial \widehat{\mathcal{B}}_i(t))] \cdot \mathcal{N}_i(s_i^*, \partial \widehat{\mathcal{B}}_i(t)) dt \mid s_i < s_i^* \leq 1 \right]. \end{aligned}$$

Here, $\text{OE}(s_i)$ equals the penalty of the arms in $\widehat{\mathcal{B}}_i(s_i)$ which would accept P_i exceeding the quota, and deducts the utility of these arms.

The second part of the cost is *under-enrollment* (UE), which occurs if the calibration parameter $s_i > s_i^*$. Then $\mathcal{N}_i(s_i^*, \widehat{\mathcal{B}}_i(s_i)) < \mathcal{N}_i(s_i^*, \widehat{\mathcal{B}}_i(s_i^*)) = q_i$. The $\text{UE}(s_i)$ equals the

opportunity cost in the sennse that P_i could have successfully pulled more arms:

$$\begin{aligned} \text{UE}(s_i) &\equiv \mathbb{E}_{s_i^*} [V_i(s_i^*, \widehat{\mathcal{B}}_i(s_i^*)) - V_i(s_i^*, \widehat{\mathcal{B}}_i(s_i)) \mid 0 \leq s_i^* < s_i] \\ &= \mathbb{E}_{s_i^*} \left[\int_{t=s_i^*}^{s_i} u_i(s_i^*, \partial \widehat{\mathcal{B}}_i(t)) \cdot \mathcal{N}_i(s_i^*, \partial \widehat{\mathcal{B}}_i(t)) dt \mid 0 \leq s_i^* < s_i \right]. \end{aligned}$$

Therefore, the goal of finding s_i to maximize the P_i 's average-case expected payoff can be written as:

$$\arg \max_{s_i \in (0,1)} \left\{ \mathbb{E}_{s_i^*} \left[V_i(s_i^*, \widehat{\mathcal{B}}_i(s_i)) - \gamma_i \max\{\mathcal{N}_i(s_i^*, \widehat{\mathcal{B}}_i(s_i)) - q_i, 0\} \right] \right\}.$$

This goal is equivalent to finding s_i to minimize the weighted sum of $\text{OE}(s_i)$ and $\text{UE}(s_i)$ with the occurrence probabilities as the weights:

$$\arg \min_{s_i \in (0,1)} \left\{ (1 - F_{s_i^*}(s_i)) \text{OE}(s_i) + (F_{s_i^*}(s_i) - \mathbb{P}(s_i^* = s_i)) \text{UE}(s_i) \right\}.$$

By the first-order condition, the minimizer $s_i \in (0, 1)$ satisfies

$$\begin{aligned} &[1 - \mathbb{P}(s_i^* = s_i)] \mathbb{E}_{s_i^*} [V_i(s_i^*, \partial \widehat{\mathcal{B}}_i(s_i)) \mid s_i^* \neq s_i] \\ &= \gamma_i (1 - F_{s_i^*}(s_i)) \mathbb{E}_{s_i^*} [\mathcal{N}_i(s_i^*, \partial \widehat{\mathcal{B}}_i(s_i)) \mid s_i < s_i^* \leq 1], \end{aligned} \tag{23}$$

where $F_{s_i^*}(\cdot)$ is the cumulative distribution function of $s_i^* \in [0, 1]$. This result proves Eq. (10). Note that there always exists a solution to Eq. (10) since when $s_i \rightarrow 0_+$, $F_{s_i^*}(s_i) \rightarrow 0$, $\gamma_i > u_i(s_i^*, \partial \widehat{\mathcal{B}}_i(s_i))$; and when $s_i \rightarrow 1_-$, $F_{s_i^*}(s_i) \rightarrow 1$, $V_i(s_i^*, \partial \widehat{\mathcal{B}}_i(s_i)) > 0$; and the right side of Eq. (10) is strictly decreasing in s_i . Moreover, the solution to Eq. (10) is the maximizer of the thee average case expected payoff $\mathbb{E}_{s_i^*} \{\mathcal{U}_i[\widehat{\mathcal{B}}_i(s_i)]\}$. This is because the strategy corresponding to the solution of Eq. (10) has a larger value of $\mathbb{E}_{s_i^*} \{\mathcal{U}_i[\widehat{\mathcal{B}}_i(s_i)]\}$ than the strategy corresponding to the boundary $s_i = 1$ or $s_i = 0$.

Here, we assume the calibration parameter $s_i \in (0, 1)$ in the definitions of $\text{OE}(s_i)$ and $\text{UE}(s_i)$. We prove that this assumption is without less of generality by showing that if $s_i = 1$, P_i can pull more arms to obtain a larger expected payoff, and if $s_i = 0$, P_i can pull less arms to obtain a larger expected payoff. Consider that if $s_i = 1$ and P_i pulls an additional arm A which is not pulled currently, that is, $A \notin \widehat{\mathcal{B}}_i(1)$. Then the expected number of arms that accept agent P_i is

$$\mathcal{N}_i(s_i^*, \widehat{\mathcal{B}}_i(1) \cup \{A\}) = \mathcal{N}_i(s_i^*, \widehat{\mathcal{B}}_i(1)) + \mathcal{N}_i(s_i^*, \{A\})$$

Let \tilde{s}_i satisfy $\mathcal{N}_i(\tilde{s}_i, \widehat{\mathcal{B}}_i(1) \cup \{A\}) = q_i$. Since $\mathcal{N}_i(s_i^*, \widehat{\mathcal{B}}_i(1) \cup \{A\}) > \mathcal{N}_i(s_i^*, \widehat{\mathcal{B}}_i(1))$, we have $\tilde{s}_i < 1$. Let A be the arm such that $\tilde{s}_i > 1 - \epsilon_s$ for some sufficiently small $\epsilon_s > 0$. Then the difference of average-case expected payoffs from pulling two arm sets $\widehat{\mathcal{B}}_i(1) \cup \{A\}$ and $\widehat{\mathcal{B}}_i(1)$

is

$$\begin{aligned}
 & \mathbb{E}_{s_i^*} [V_i(s_i^*, \{A\})] - \gamma_i \mathbb{E}_{s_i^*} \left[q_i - \mathcal{N}_i(s_i^*, \widehat{\mathcal{B}}_i(1) \cup \{A\}) \mid \widetilde{s}_i < s_i^* \leq 1 \right] \\
 &= \mathbb{E}_{s_i^*} [V_i(s_i^*, \{A\})] - \gamma_i \mathbb{E}_{s_i^*} [\mathcal{N}_i(s_i^*, \{A\}) \mid \widetilde{s}_i < s_i^* \leq 1] \\
 & \quad + \gamma_i \mathbb{E}_{s_i^*} \left[q_i - \mathcal{N}_i(s_i^*, \widehat{\mathcal{B}}_i(1)) \mid \widetilde{s}_i < s_i^* \leq 1 \right] \\
 &> \mathbb{E}_{s_i^*} [V_i(s_i^*, \{A\})] - \gamma_i \mathbb{E}_{s_i^*} [\mathcal{N}_i(s_i^*, \{A\}) \mid \widetilde{s}_i < s_i^* \leq 1] \\
 &= U_i(A) \mathbb{E}_{s_i^*} [\mathcal{N}_i(s_i^*, \{A\}) \mid 0 \leq s_i^* \leq \widetilde{s}_i] - [\gamma_i - U_i(A)] \mathbb{E}_{s_i^*} [\mathcal{N}_i(s_i^*, \{A\}) \mid \widetilde{s}_i < s_i^* \leq 1] \\
 &> 0,
 \end{aligned}$$

where $U_i(A)$ is the latent utility of arm A defined in Eq. (1). The last step holds for sufficiently small $\epsilon_s > 0$. Similarly, if $s_i = 0$, P_i can benefit by pulling less arms. Thus, the assumption that $s_i \in (0, 1)$ is without loss of generality.

If $F_{s_i^*}(\cdot)$ has discrete support, we require that $\text{UE}(s_i)$ is at least $\text{OE}(s_i)$. By the first-order condition similar to Eq. (23), we find the minimal $s_i \in [0, 1]$ such that

$$\begin{aligned}
 & [1 - \mathbb{P}(s_i^* = s_i)] \mathbb{E}_{s_i^*} [V_i(s_i^*, \partial \widehat{\mathcal{B}}_i(s_i)) \mid s_i^* \neq s_i] \\
 & \geq \gamma_i (1 - F_{s_i^*}(s_i)) \mathbb{E}_{s_i^*} [\mathcal{N}_i(s_i^*, \partial \widehat{\mathcal{B}}_i(s_i)) \mid s_i < s_i^* \leq 1],
 \end{aligned} \tag{24}$$

where the search of s_i starts from the maximum value in the support to the minimal value. We note that the calibration in Eq. (24) is a *conservative* counterpart as compared with the calibration such that $\text{OE}(s_i)$ is at least $\text{UE}(s_i)$:

$$\begin{aligned}
 & [1 - \mathbb{P}(s_i^* = s_i)] \mathbb{E}_{s_i^*} [V_i(s_i^*, \partial \widehat{\mathcal{B}}_i(s_i)) \mid s_i^* \neq s_i] \\
 & \leq \gamma_i (1 - F_{s_i^*}(s_i)) \mathbb{E}_{s_i^*} [\mathcal{N}_i(s_i^*, \partial \widehat{\mathcal{B}}_i(s_i)) \mid s_i < s_i^* \leq 1].
 \end{aligned} \tag{25}$$

The calibration in Eq. (24) is preferred to that in Eq. (25) since we want the calibration to be sensitive to the penalty γ_i . This completes the proof. \blacksquare

B.6 Proof of Theorem 5: Maximin Calibration for CDM

Proof We use the notations $V_i(s_i^*, \widehat{\mathcal{B}}_i(s_i))$ and $\mathcal{N}_i(s_i^*, \widehat{\mathcal{B}}_i(s_i))$ defined in Appendix B.5. The maximum over-enrollment cost for any $s_i \in [0, 1]$ is

$$\max_{s_i^* \in [0, 1]} \{\text{OE}(s_i)\} = \gamma_i \{\mathcal{N}_i(1, \widehat{\mathcal{B}}_i(s_i)) - \mathcal{N}_i(1, \widehat{\mathcal{B}}_i(1))\} - \{V_i(1, \widehat{\mathcal{B}}_i(s_i)) - V_i(1, \widehat{\mathcal{B}}_i(1))\}.$$

Since $\gamma_i > \sup_{j \in \mathcal{A}^{T+1}} \{v_j + e_{ij}\}$, $\max_{s_i^*} \{\text{OE}(s_i)\}$ is strictly decreasing in s_i . The maximum under-enrollment cost for any $s_i \in [0, 1]$ is

$$\max_{s_i^* \in [0, 1]} \{\text{UE}(s_i)\} = V_i(0, \widehat{\mathcal{B}}_i(0)) - V_i(0, \widehat{\mathcal{B}}_i(s_i)),$$

which is strictly increasing in s_i . Hence, the goal of maximizing the minimal total expected payoff $\max_{s_i} \min_{s_i^*} \{U_i^S(\widehat{\mathcal{B}}_i(s_i))\}$ is equivalent to minimizing the larger one between $\text{OE}(s_i)$ and $\text{UE}(s_i)$:

$$\min_{s_i \in [0, 1]} \max \left\{ \max_{s_i^*} \{\text{OE}(s_i)\}, \max_{s_i^*} \{\text{UE}(s_i)\} \right\}.$$

This objective amounts to finding s_i such that

$$\max_{s_i^* \in [0,1]} \{\text{OE}(s_i)\} = \max_{s_i^* \in [0,1]} \{\text{UE}(s_i)\}. \quad (26)$$

This proves Theorem 5. Moreover, there exists a unique solution to Eq. (26) since when $s_i = 0$, $\max_{s_i^*} \{\text{OE}(0)\} > \max_{s_i^*} \{\text{UE}(0)\} = 0$, and when $s_i = 1$, $\max_{s_i^*} \{\text{UE}(1)\} > 0 = \max_{s_i^*} \{\text{OE}(1)\}$, and together with the fact that $\max_{s_i^*} \{\text{OE}(s_i)\}$ and $\max_{s_i^*} \{\text{UE}(s_i)\}$ are monotonic continuous functions of s_i .

If $F_{s_i^*}(\cdot)$ has discrete support, we requires that the maximal $\text{UE}(s_i)$ is at least the maximal $\text{OE}(s_i)$. Hence, we need to change the goal in Eq. (26) to finding the minimal $s_i \in [0, 1]$ such that $\max_{s_i^*} \{\text{UE}(s_i)\} \geq \max_{s_i^*} \{\text{OE}(s_i)\}$. This completes the proof. ■

B.7 Proof of Theorem 6: Incentives of Agents

Proof First, by Theorem 3, the CDM takes the cutoff strategy with respect to the arms' latent utilities given by Eq. (1). Since the agents' true preferences are determined by the arms' latent utilities, the agents would act according to their true preferences under the CDM.

Second, Theorem 2 proves the consistency of the estimated acceptance probability using historical data. Theorems 4 and 5 show that the CDM maximizes the agent's expected payoff, given the population acceptance probability, where the expected payoff is measured in either average-case or worst-case given the uncertain true state. Thus, it is a dominant strategy for each agent to act according to the CDM.

Combining these two observations, we conclude that as $T \rightarrow \infty$, the CDM is a procedure that gives agents the incentives to act according to true preferences. ■

B.8 Proof of Theorem 7: Stability of CDM

Proof Suppose that an agent-arm pair (P_i, A_j) forms a blocking pair. Then it implies one of the following two cases:

- (I). P_i prefers A_j to some of its matched arms.
- (II). P_i has unfilled quota and A_j is unmatched.

For case (I), P_i must have pulled A_j according to the cutoff strategy of CDM by Theorem 3. However, P_i must have been subsequently rejected by A_j in favor of some agent that A_j liked better. Hence, A_j must prefer its currently matched agent to P_i and there is no instability.

For case (II), P_i did not pull A_j since otherwise, A_j would have accepted P_i . Note that Theorem 2 proves the consistency of acceptance probability estimate using historical data. The individual rationality in Eq. (13) implies that P_i would find A_j unacceptable in the decentralized market, given the population acceptance probability. Thus, there is no instability.

Combining these two cases, we conclude that as $T \rightarrow \infty$, CDM yields a stable matching in decentralized markets. \blacksquare

B.9 Proof of Theorem 8: Fairness of CDM

Proof By Theorem 6, the CDM gives agents the incentives to act on their true preferences, as $T \rightarrow \infty$. Hence, each agent pulls arms according to its true preference for arms. If an arm A_j prefers an agent $P_{i'}$ to another agent P_i that pulls A_j , then all arms pulled by $P_{i'}$ must rank above A_j according to the true preference of $P_{i'}$. By definition, we conclude that the matching process according to CDM is fair. \blacksquare

B.10 Proof of Theorem 9: Unfairness of the Oracle Set

Proof Denote by $e_i^*(v)$ the cutoff curve corresponding to the oracle arm set \mathcal{B}_i^* in Eq. (14):

$$e_i^*(v_j) = \min \left\{ \max \left\{ \gamma_i \cdot \mathbb{P}(s_i^* \in O_{\mathcal{B}_i^*}) \frac{\mathbb{E}_{s_i^*}[\pi_i(s_i^*, v_j) \mid s_i^* \in O_{\mathcal{B}_i^*}]}{\mathbb{E}_{s_i^*}[\pi_i(s_i^*, v_j)]} - v_j, 0 \right\}, 1 \right\}, \quad \forall j.$$

The average-case expected utility of an arm from the cutoff curve is

$$\mathcal{U}_i(v_j, e_i^*(v_j)) = (v_j + e_i^*(v_j)) \cdot \mathbb{E}_{s_i^*}[\pi_i(s_i^*, v_j)] - \gamma_i \cdot \mathbb{P}(s_i^* \in O_{\mathcal{B}_i^*}) \mathbb{E}_{s_i^*}[\pi_i(s_i^*, v) \mid s_i^* \in O_{\mathcal{B}_i^*}].$$

Hence, for $e_i^*(v_j) \in (0, 1)$,

$$\mathcal{U}_i(v_j, e_i^*(v_j)) = 0.$$

Since the acceptance probability $\pi_i(s_i, v)$ is increasing in s_i , $\mathbb{E}_{s_i^*}[\pi_i(s_i^*, v) \mid s_i^* \in O_{\mathcal{B}_i^*}] > \mathbb{E}_{s_i^*}[\pi_i(s_i^*, v)]$. Thus, for $e_i^*(v_j) \in (0, 1)$,

$$\begin{aligned} v_j + e_i^*(v_j) &= \gamma_i \cdot \mathbb{P}(s_i^* \in O_{\mathcal{B}_i^*}) \frac{\mathbb{E}_{s_i^*}[\pi_i(s_i^*, v_j) \mid s_i^* \in O_{\mathcal{B}_i^*}]}{\mathbb{E}_{s_i^*}[\pi_i(s_i^*, v_j)]} \\ &> \gamma_i \cdot \mathbb{P}(s_i^* \in O_{\mathcal{B}_i^*}). \end{aligned} \quad (27)$$

By Eq. (15), we can derive that

$$\mathbb{E}_{s_i^*} \left[\frac{\partial \pi_i(s_i^*, v_j)}{\partial v} \right] < \mathbb{E}_{s_i^*} \left[\frac{\partial \pi_i(s_i^*, v_j)}{\partial v} \mid s_i^* \in O_{\mathcal{B}_i^*} \right] < 0, \quad \forall v_j \in (v', v'').$$

This inequality together with Eq. (27) yield that for $v_j \in (v', v'')$ and $e_i^*(v_j) \in (0, 1)$,

$$\begin{aligned} \frac{\partial \mathcal{U}_i(v, e_i^*(v))}{\partial v} \Big|_{v=v_j} &= \mathbb{E}_{s_i^*}[\pi_i(s_i^*, v_j)] + (v_j + e_i^*(v_j)) \mathbb{E}_{s_i^*} \left[\frac{\partial \pi_i(s_i^*, v_j)}{\partial v} \right] \\ &\quad - \gamma_i \cdot \mathbb{P}(s_i^* \in O_{\mathcal{B}_i^*}) \mathbb{E}_{s_i^*} \left[\frac{\partial \pi_i(s_i^*, v_j)}{\partial v} \mid s_i^* \in O_{\mathcal{B}_i^*} \right] \\ &< \mathbb{E}_{s_i^*}[\pi_i(s_i^*, v_j)]. \end{aligned}$$

Thus, by the implicit function theorem,

$$\left| \frac{de_i^*(v_j)}{dv} \right| = \left| \frac{\partial \mathcal{U}_i(v_j, e_i^*(v_j)) / \partial v}{\partial \mathcal{U}_i(v_j, e_i^*(v_j)) / \partial e_i^*} \right| < \frac{\mathbb{E}_{s_i^*}[\pi_i(s_i^*, v_j)]}{\mathbb{E}_{s_i^*}[\pi_i(s_i^*, v_j)]} = 1, \quad \forall v_j \in (v', v''),$$

Therefore, we can find two arm sets $\mathcal{B}_i^{(1)}, \mathcal{B}_i^{(2)} \subseteq (v', v'') \times [0, 1]$ and a constant c_0 such that, for all $(v_j, e_{ij}) \in \mathcal{B}_i^{(1)}$, $e_{ij} > e_i^*(v_j)$ and $v_j + e_{ij} < c_0$, and for all $(v_j, e_{ij}) \in \mathcal{B}_i^{(2)}$, $e_{ij} < e_i^*(v_j)$ and $v_j + e_{ij} > c_0$. Hence, the arms in $\mathcal{B}_i^{(2)}$ have justified envy toward arms in $\mathcal{B}_i^{(1)}$. We refer to Figure 2 for an illustration. By definition of fairness, the strategy corresponding to the oracle set in Eq. (14) is unfair. \blacksquare