

Incentive Design in Peer Review: Rating and Repeated Endogenous Matching

Yuanzhang Xiao, Florian Dörfler, and Mihaela van der Schaar

◆

Abstract—Peer review (e.g., grading assignments in Massive Open Online Courses (MOOCs), academic paper review) is an effective and scalable method to evaluate the products (e.g., assignments, papers) of a large number of agents when the number of dedicated reviewing experts (e.g., teaching assistants, editors) is limited. Peer review poses two key challenges: 1) identifying the reviewers’ intrinsic capabilities (i.e., *adverse selection*) and 2) incentivizing the reviewers to exert high effort (i.e., *moral hazard*). Some works in mechanism design address *pure* adverse selection using *one-shot* matching rules, and *pure* moral hazard was addressed in repeated games with *exogenously given* and *fixed* matching rules. However, in peer review systems exhibiting both adverse selection and moral hazard, one-shot or exogenous matching rules do not link agents’ current behavior with future matches and future payoffs, and as we prove, will induce myopic behavior (i.e., exerting the lowest effort) resulting in the lowest review quality.

In this paper, we propose for the first time a solution that *simultaneously* solves adverse selection and moral hazard. Our solution exploits the *repeated* interactions of agents, utilizes *ratings* to summarize agents’ past review quality, and designs matching rules that *endogenously* depend on agents’ ratings. Our proposed matching rules are easy to implement and require no knowledge about agents’ private information (e.g., their benefit and cost functions). Yet, they are effective in guiding the system to an equilibrium where the agents are incentivized to exert high effort and receive ratings that precisely reflect their review quality. Using several illustrative examples, we quantify the significant performance gains obtained by our proposed mechanism as compared to existing one-shot or exogenous matching rules.

1 INTRODUCTION

Peer review serves as an effective and scalable method for performance evaluation in systems where the products to evaluate significantly outnumber the dedicated reviewing experts. One example of such systems is Massive Open Online Courses (MOOCs), where the number of students enrolled in a course is in the order of tens of thousands and by far exceeds the number of teaching assistants [2]–[5]. Another example is academic paper review, where the number of papers submitted to a journal by far exceeds the number of (associate) editors. Since the proposed mechanism can be applied to general peer review settings, we keep the discussion in this paper general.¹

Xiao and van der Schaar are with Department of Electrical Engineering, UCLA. Emails: yxiao@seas.ucla.edu and mihaela@ee.ucla.edu
Dörfler is with Automatic Control Laboratory, ETH Zurich, Switzerland. Email: dorfler@ethz.ch

1. We studied explicitly the academic paper review system in the preliminary conference version [1] of this paper.

Peer review systems pose two key challenges. First, the reviewers have different *intrinsic capabilities* (e.g., their review quality functions, benefit and cost functions), which are *unknown*. Hence, one challenge is how to identify their *unknown* intrinsic capabilities; this is known in the game theory literature as the adverse selection problem. Second, the reviewers can choose to exert different levels of (costly) *effort* (e.g., time and energy spent in reviewing), which is *unobservable*. Hence, the other challenge is how to incentivize reviewers to exert high effort; this is known in the game theory literature as the moral hazard problem. A reviewer’s ultimate review quality is determined by her intrinsic capabilities *and* effort. If the capabilities are unknown but the effort is observable (i.e. pure adverse selection), there is hope to identify their capabilities through mechanism design. If the effort is unobservable but the capabilities are known (i.e. pure adverse selection), there is hope to incentivize high effort through social norms. However, in the presence of both adverse selection and moral hazard, the problem becomes significantly more challenging. In fact, no existing work has addressed this problem systematically.

A natural candidate for solving the pure adverse selection problem is to use matching mechanisms [6]–[8]. Matching mechanisms aim to efficiently allocate resources (e.g., hospitals, or reviewers in our setting) to agents (e.g., medical students) or their products (e.g., assignments or papers in our setting). Existing works on matching mechanisms assume that the quality of resources depend only on the types of the agents who provide and receive the resource, but not on the providers’ effort. In other words, there is no moral hazard problem. As a result, they focus on *one-shot* interactions and design *one-shot* matching rules (i.e., each agent is matched only once).² However, their assumption does not hold in peer review systems, where the review quality depends crucially on the reviewers’ effort. We prove that under one-shot matching rules, agents will behave myopically by choosing the lowest effort (i.e., free-riding), because their current effort does not affect future matches and their future payoffs. Hence, the system performance (in terms of the total review quality) is the worst since one-shot matching does not address the moral hazard problem.

One way to address the pure moral hazard problem is

2. There are works on dynamic matching (see representative work [9]). However, matching is called dynamic due to the dynamic arrival and departure of the agents. Each agent is still matched only once.

to use social norms [10], where a central agency assigns the agents with ratings that summarize their past behavior and recommends a “norm” (i.e., desired behavior) that rewards agents with good ratings and punishes those with bad ratings. In this way, the agents are incentivized to conform with the social norm (e.g., exert high effort in our setting), even when they are randomly matched to each other based on some *exogenous* matching rule. However, existing works on social norms assume that the agents are *homogenous*. This assumption does not hold in peer review systems, because different reviewers have different intrinsic capabilities. Ideally, the central agency should recommend different norms to agents of different capabilities; in practice, it cannot do this since the capabilities are unknown. In summary, existing works in social norms [10]–[13] do not deal with the adverse selection problem that is present in peer review.

This paper proposes the *first* mechanism to *simultaneously* solve the adverse selection and moral hazard problems in peer review. Our proposed mechanism exploits the repeated interaction among agents (i.e., by submitting multiple assignments or papers over time), and assigns the agents with ratings, which are summaries of their past review quality. Unlike the works on social norms [10]–[13], we do not recommend desired behavior (i.e., a social norm) to the agents, because they have unknown, different capabilities and thus computing a recommended social norm is impossible. Instead, we propose rules for *repeated* matching that *endogenously* depend on agents’ ratings. Unlike existing one-shot [6]–[8] or exogenous [10]–[13] matching rules, our proposed repeated endogenous matching rules provide strong incentives for agents to exert high effort, because the agents’ behaviors affect their future ratings and hence, their future matches and future payoffs.

We provide design guidelines for endogenous matching rules that are easy to implement without knowledge of agents’ private information (e.g., their benefit, review quality, and cost functions), yet powerful enough to guide the system to desirable equilibria. In particular, in the equilibrium the agents find it in their self-interest to exert high effort, and receive ratings that truly reflect their capabilities. We also provide case studies on specific matching rules with different reward/punishment schemes. We show that different reward/punishment schemes lead to different optimal matching rules, which stresses the importance of tailoring matching rules to reward/punishment schemes. Simulation results demonstrate large performance improvement over existing matching rules.

In the following, we discuss related works in Section 2. Then we describe the model and formulate the design problem in Section 3. We study general matching rules in Section 4, and the baseline matching rule and its extensions in Section 5. Section 6 demonstrates the efficiency of our proposed mechanisms. Finally, Section 7 concludes the paper.

2 RELATED WORKS

2.1 Pure Adverse Selection

The pure adverse selection problem is the focus of a huge literature on matching in resource allocation (e.g., allocation of schools to applicants [6][7]) and exchange (e.g., kidney exchange [8]). These works ignore the moral hazard problem.

In particular, they do not consider “effort”. Once an agent (e.g., an applicant) is matched to another (e.g., a school), the benefit (obtained by this applicant) is fixed. In contrast, in our work, the review quality depends crucially on the reviewer’s effort. This additional moral hazard problem, when ignored, will significantly degrade the system performance (in terms of the total review quality).

Since these works [6]–[8] ignore moral hazard, their matching rules are one-shot (i.e., match each agent only once). In contrast, our matching rules are repeated and changing over time based on agents’ ratings. Hence, our matching rules can incentivize agents to exert high effort to obtain better ratings and thus, favorable future matches.

2.2 Pure Moral Hazard

The pure moral hazard problem has been studied in repeated game theory, where anonymous agents are randomly matched to interact with each other [10]–[13], as in our work. However, these works [10]–[13] focus on the *pure* moral hazard problem, and ignore the adverse selection problem by assuming *homogeneous* agents. In this work, we assume *heterogeneous* agents and deal with both the moral hazard and adverse selection problems. In [10]–[13], due to the homogeneity of agents, binary ratings are usually sufficient to identify whether a player has behaved well or badly. In contrast, in this work the rating is continuous, such that the rating mechanism can identify not only whether a player has behaved well or badly, but also its review quality.

Another key difference is that we *design* matching rules that *endogenously* depend on agents’ ratings and directly affect their incentives, while the matching rules in [10]–[13] are *fixed* and *exogenously given*. In our setting, we will prove that the latter type of matching rules will result in the lowest review quality in the equilibrium.

2.3 Other Works on Peer Review

There are other works on peer review systems but with different problems to solve. For example, there are works focusing on how to aggregate reviewers’ scores/ratings to obtain a final score/rating that accurately reflects the true quality of the assignments (in MOOCs [3]–[5]), the proposals (in NSF proposal reviewing [14]), or the papers (in academic peer review [15]). In contrast, our focus is to incentivize reviewers to exert high effort levels.

3 MODEL

3.1 Basic Setup

Consider a peer review system with a set $\mathcal{N} = \{1, \dots, N\}$ of N agents. Each agent has its products reviewed by the other agents. An agent benefits from the review by its reviewer, and exerts effort in reviewing others’ products. A designer (e.g., the instructor in MOOCs) aims to design a mechanism that incentivizes the reviewers to produce high-quality reviews. The mechanism includes two parts: (i) the rating mechanism that assigns and updates a rating $\theta_i \in \mathbb{R}_+$ for each agent i , and (ii) the matching rule that matches agents with reviewers (possibly based on the ratings). In the following, we write the rating profile, namely the ratings of every agent, as $\theta = (\theta_1, \dots, \theta_N)$. The rating profile is known

only to the designer. We define the **rating distribution**, denoted by a vector $d(\theta)$, as the ordered (from high to low) list of all the ratings. The rating distribution $d(\theta)$ does not count multiple agents with the same rating. For example, if the rating profile is $\theta = (3, 5, 5, 3)$, the rating distribution will be $d(\theta) = (5, 3)$. Write K as the number of distinct ratings in θ (i.e., the dimension of the vector $d(\theta)$), and k_i as agent i 's ranking (i.e., ordered position) in the rating distribution. In the above example, we have $K = 2$, $k_1 = k_4 = 2$, $k_2 = k_3 = 1$. Although K and k_i depend on θ , we write them simply as K and k_i for notational simplicity without causing confusion. Denote the k_i th element of the rating distribution by $d(\theta)_{k_i}$. Then we have $\theta_i = d(\theta)_{k_i}$. Finally, notice that the rating distribution does not disclose any information about the identities of the agents.

Note, importantly, that an agent's rating indicates its review quality, not the quality of its product.

Time is slotted into $t = 0, 1, 2, \dots$. In each time slot t , the entities in the system moves in the following order:³

- The designer publishes the rating distribution $d(\theta)$, and informs agent i of its rating θ_i and its ranking k_i .
- Each agent submits its product to review.
- The designer matches each agent i 's product to other agent(s) for review based on a probabilistic **matching rule** $m_{k_i k_j} : (d(\theta)_{k_i}, d(\theta)_{k_j}) \mapsto [0, 1]$. The matching rule determines the probability $m_{k_i k_j}(d(\theta)_{k_i}, d(\theta)_{k_j})$ that the agent with the k_i th highest rating is matched to the reviewer with the k_j th highest rating. From the definition we can see that the matching does not depend on agents' identities.
- Each reviewer j exerts an **effort level** $e_j^t \in [0, e_j^{\max}]$, where e_j^{\max} is j 's maximum effort level. Reviewer j 's **review quality** then depends on its effort as $q_j(e_j^t)$, where $q_j : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is the review quality function.
- Each agent i receives **benefit** $b_i(q_j(e_j^t))$ from reviewer j 's review, where $b_i : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is i 's benefit function, and incurs a **cost** of $c_i(e_i^t)$ for reviewing a product, where $c_i : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is i 's cost function.
- Each agent i sends a **report** r_i^t about the reviewer to the designer. We assume that the report accurately reflects the reviewer's review quality, namely $r_i^t = q_j(e_j^t)$. For examples, in MOOCs the report can be made accurate by comparing the grading with the true answers to the assignments (posted after the submission of the assignments).
- The designer updates the agents' ratings according to the **rating update rule** $\pi : (\theta_j^t, r_i^t) \mapsto \theta_j^{t+1}$. For fairness, the rating update rule is identical for all reviewers, and is given by a convex combination of the reviewer's old rating and the report about its review quality with a constant step size $\mu \in (0, 1)$:⁴

$$\pi(\theta_j^t, r_i^t) = \begin{cases} (1 - \mu) \cdot \theta_j^t + \mu \cdot r_i^t, & j \text{ has reviewed} \\ \theta_j^t, & \text{otherwise} \end{cases} \quad (1)$$

3. Throughout the paper, the superscript $(\cdot)'$ on a function refers to the derivative, and the superscript $(\cdot)^t$ refers to the variable under consideration at time point $t \in \mathbb{Z}_+$.

4. Note that under the assumption that j 's review quality $q_j(e_j)$ is perfectly observed, it does not matter how many reports about j 's review quality are received.

We make the following remarks on the agents' ratings. Each agent i has a maximum effort level e_i^{\max} , and hence has a maximum review quality $q_i(e_i^{\max})$. Since the new rating is the convex combination of the old rating and the review quality, given any initial rating θ_i^0 , agent i 's rating can only be in the interval $[0, \theta_i^{\max}]$, where $\theta_i^{\max} \triangleq \max\{\theta_i^0, q_i(e_i^{\max})\}$. In other words, the possible ratings of each agent i are contained in the compact set $[0, \theta_i^{\max}]$.

Throughout the paper, we make the following reasonable and standard assumptions on the monotonicity, convexity, concavity, and differentiability of our functions.

Assumption 1 (Cost, Review Quality, and Benefit). Each agent i 's cost function $c_i(\cdot)$, review quality function $q_i(\cdot)$, and benefit function $b_i(\cdot)$ satisfy the following:

- The cost $c_i(\cdot)$ is strictly convex, strictly increasing, and twice continuously differentiable in effort e_i . In addition, $c_i'(0) = 0$.
- The review quality $q_i(\cdot)$ is concave, strictly increasing, and twice continuously differentiable in effort e_i . In addition, $q_i'(0)$ exists and is bounded.
- The benefit $b_i(\cdot)$ is strictly increasing, concave, and continuously differentiable in review quality q_j . In addition, $b_i'(0)$ exists and is bounded.
- We normalize $c_i(0) = 0$, $q_i(0) = 0$, and $b_i(0) = 0$.

3.2 Information – Who Knows What

3.2.1 The designer

The designer receives reports r_i of review quality, and keeps the rating θ_i^t for each agent i at each time slot t . Hence, the designer knows the identity of the agent at the k_i th position of the rating distribution. However, it does not know the review quality functions $q_i(\cdot)$, the benefit functions $b_i(\cdot)$, or the cost functions $c_i(\cdot)$.

3.2.2 Each agent i

Each agent i knows its own review quality function $q_i(\cdot)$, benefit function $b_i(\cdot)$, and cost function $c_i(\cdot)$, but does not know the above functions of the other agents. It knows the matching rule m and the rating update rule π . It also knows its own rating θ_i^t , the rating distribution $d(\theta)^t$, and its position in the rating distribution k_i^t , but does not know the others' ratings or the identity of its reviewer.

3.3 Payoffs and Equilibrium

In each time slot t , agent i 's expected payoff is its expected benefit from the reviewing of its product minus the expected cost of reviewing other agents' products. We write agent i 's expected payoff as $u_i(m, \theta_i, d(\theta), e)$, which depends on the matching rule m , its own rating θ_i , the rating distribution $d(\theta)$, and all the agents' effort levels $e \triangleq (e_1, \dots, e_N)$. The *expected payoff* can be calculated as the expected benefit minus the expected cost:

$$u_i(m, \theta_i, d(\theta), e) = \sum_{k_j \neq k_i} m_{k_i k_j}(d(\theta)_{k_i}, d(\theta)_{k_j}) \cdot b_i(q_j(e_j)) - \left[\sum_{k_j \neq k_i} m_{k_j k_i}(d(\theta)_{k_j}, d(\theta)_{k_i}) \right] \cdot c_i(e_i). \quad (2)$$

Each agent i aims to choose a sequence of effort levels over time to maximize the *discounted average of expected payoffs*, i.e., to solve the *dynamic optimization problem* below:

$$\max_{\{e_i^t \in [0, e_i^{\max}]\}_{t=0}^{\infty}} \mathbb{E} \left\{ (1 - \delta_i) \sum_{t=0}^{\infty} \delta_i^t u_i(m, \theta_i^t, d(\theta^t), e_i^t, e_{-i}^t) \right\}, \quad (3)$$

where e_{-i}^t is the effort levels chosen by all the agents other than i at time t , and $\delta_i \in [0, 1)$ is agent i 's *discount factor*. An agent's discount factor reflects its patience. We take the expectation $\mathbb{E}\{\cdot\}$ because the rating update is random, namely an agent's rating is either updated or kept the same depending on whether it has reviewed a product.

Note that the optimization problem (3) is very hard, if not impossible, to solve. The difficulty lies in the couplings of one agent's decisions over time and with other agents' decisions. First, the agent's current decision (i.e., effort level) affects not only its current payoff (through the cost), but also its future ratings and hence future payoffs. Second, the agent's payoff is affected by the others' decisions (through the benefit). However, since an agent has no knowledge about the others, it cannot predict the others' decisions and the evolution of rating distributions. In summary, an agent cannot solve the optimization problem (3) due to computational complexity and lack of knowledge.

We propose a realistic behavioral model for the agents. To choose the optimal effort level at each time t , each agent i holds a *conjecture* that its future value $\mathbb{E}\{(1 - \delta_i) \sum_{\tau=t}^{\infty} \delta_i^{\tau-t} u_i(m, \theta_i^\tau, d^\tau, e_i^\tau, e_{-i}^\tau)\}$ (i.e., its discounted average payoff after time t) is the following:

$$f_i(\alpha_i, \beta_i^t, \theta_i^t, d(\theta^t), e_i^t) \triangleq \alpha_i \cdot \bar{b}_i(\theta_i^t, d(\theta^t), e_i^t) + \beta_i^t, \quad (4)$$

where $\bar{b}_i(\theta_i^t, d(\theta^t), e_i^t)$ is the conjectured expected benefit of agent i in time $t+1$, assuming that the others' ratings remain the same. We can calculate $\bar{b}_i(\theta_i^t, d(\theta^t), e_i^t)$ as

$$\begin{aligned} & \bar{b}_i(\theta_i^t, d(\theta^t), e_i^t) \\ &= \sum_{k_j \neq k_i^+} \left[m_{k_i^+ k_j} \left(\underbrace{\pi(\theta_i^t, q_i(e_i^t))}_{\triangleq \theta_i^{t+1}}, d(\theta_i^{t+1}, \theta_{-i}^t)_{k_j} \right) \right. \\ & \quad \left. \cdot b_i(d(\theta_i^{t+1}, \theta_{-i}^t)_{k_j}) \right], \end{aligned} \quad (5)$$

where $d(\theta_i^{t+1}, \theta_{-i}^t)$ is the new rating distribution when i 's rating is updated to $\theta_i^{t+1} \triangleq \pi(\theta_i^t, q_i(e_i^t))$ and the others' ratings remain the same, and k_i^+ is i 's new ranking of its new rating θ_i^{t+1} in the new rating distribution $d(\theta_i^{t+1}, \theta_{-i}^t)$. Note that i can compute $d(\theta_i^{t+1}, \theta_{-i}^t)$ based on θ_i^{t+1} and $d(\theta^t)$, without knowing θ_{-i}^t .

Each agent i holds the conjecture (4) for two reasons. First, it cannot predict the others' effort levels or future ratings. Hence, it holds a conjecture that *the others' ratings remain the same*, and that the others' ratings precisely reflect their review quality, namely $d(\theta^t)_{k_j} = q_j(e_j^t)$. Second, it conjectures that its future value is an *affine function of its expected benefit*. For consistency, both of the above conjectures are required to be true in the equilibrium to be defined later.

The coefficient α_i reflects how "optimistic" an agent is about the rating mechanism. An agent with a larger α_i "believes in" the rating mechanism more, because it anticipates a higher future value given the expected benefit. The coefficient β_i^t is updated in each time slot by agent i ,

such that the conjectured future value converges to the true future value in the equilibrium.

Then at each time t , each agent i simply solves the following static problem for its optimal effort level e_i^t :

$$e_i^t = \arg \max_{e_i \in [0, e_i^{\max}]} (1 - \delta_i) \cdot u_i(m, \theta_i^t, d(\theta^t), e_i, e_{-i}) + \delta_i \cdot f_i(\alpha_i, \beta_i^t, \theta_i^t, d(\theta^t), e_i). \quad (6)$$

Note that the others' current effort levels e_{-i} only affect the benefit term in the current payoff $u_i(m, \theta_i^t, d(\theta^t), e_i, e_{-i})$, which does not depend on i 's effort e_i and can be considered as a constant. Hence, each agent i has all the information needed to solve the above static optimization problem.

Definition 1 (Conjectural Equilibrium [16]). Given any matching rule m and any rating update rule π , a conjectural equilibrium (CE) is a triple $\{\theta_i^*, e_i^*, \beta_i^*\}_{i \in \mathcal{N}}$ that satisfies:

- Incentive compatibility constraints: for all $i \in \mathcal{N}$,
$$e_i^* = \arg \max_{e_i \in [0, e_i^{\max}]} (1 - \delta_i) \cdot u_i(m, \theta_i^*, d(\theta^*), e_i, e_{-i}^*) + \delta_i \cdot f_i(\alpha_i, \beta_i^*, \theta_i^*, d(\theta^*), e_i), \quad (7)$$
- Stable and correct ratings: for all $i \in \mathcal{N}$, $\theta_i^* = q_i(e_i^*)$,
- Consistent conjectures: for all $i \in \mathcal{N}$,
$$f_i(\alpha_i, \beta_i^*, \theta_i^*, d(\theta^*), e_i^*) = u_i(m, \theta_i^*, d(\theta^*), e_i^*). \quad (8)$$

In the above definition, the incentive compatibility constraints ensure that the effort level e_i^* is the best response of each agent i . In other words, it will be in agent i 's self-interest to choose e_i^* . A CE also requires that each agent's rating truly reflects its review quality at the equilibrium effort level e_i^* , and hence each agent's rating is stable, namely $\pi(\theta_i^*, q_i(e_i^*)) = \theta_i^*$. Finally, a CE requires that each agent's conjecture about its future value is correct.

There may be many CEs. As a designer, it is desirable that the system will converge to a CE from any initial rating profile. The convergence is important, because the designer can distinguish the true review quality of the reviewers at the equilibrium. The choice of the matching rule plays an important role in ensuring the convergence to a CE. Aside from convergence, certain CEs are more desirable than others, as discussed in the next paragraphs.

3.4 The Design Problem Formulation

The designer's problem is to maximize the equilibrium review quality. We write the designer's objective as a function of the equilibrium review quality $W(q_1(e_1^*), \dots, q_N(e_N^*))$. Then the designer problem can be defined as

$$\begin{aligned} & \max_{m, \pi} W(q_1(e_1^*), \dots, q_N(e_N^*)) \\ & \text{s.t.} \quad \{\theta_i^*, e_i^*, \beta_i^*\}_{i \in \mathcal{N}} \text{ is a CE under } m, \pi. \end{aligned} \quad (9)$$

Note that the designer does *not* maximize the social welfare (i.e., the total benefit minus cost of the agents), because it is more natural from the designer's perspective to maximize the total review quality. The designer of the peer review system may not care about the cost of reviewing; in fact, it would like to elicit more effort from the reviewers, resulting in higher-quality reviews but higher costs.

4 CONVERGENCE TO CONJECTURAL EQUILIBRIA

In this section, we consider general matching rules, and provide important guidelines for designing the matching rules. As discussed before, we would like to have a matching rule under which the system will converge to a CE from any initial rating profile under the best response dynamics. In this way, the designer can distinguish the true quality of the reviewers in the equilibrium. Before discussing the properties of the matching rules that ensure the convergence, we first describe the best response dynamics.

At each time slot t , the best response dynamics consist of the following three updates:

$$e_i^t = \arg \max_{e_i \in [0, e_i^{\max}]} (1 - \delta_i) \cdot u_i(m, \theta_i^t, d(\theta^t), e_i, e_{-i}^t) + \delta \cdot f_i(\alpha_i, \beta_i^t, \theta_i^t, d(\theta^t), e_i); \quad (10)$$

$$\theta_i^{t+1} = \begin{cases} (1 - \mu) \cdot \theta_i^t + \mu \cdot q_i(e_i^t) & \text{if } i \text{ reviewed} \\ \theta_i^t & \text{otherwise} \end{cases}; \quad (11)$$

$$\beta_i^{t+1} = u_i(m, \theta_i^t, d(\theta^t), e^t) - \alpha_i \cdot \bar{b}_i(\theta_i^t, d(\theta^t), e_i^t). \quad (12)$$

The update of effort levels in (10) and the update of ratings in (11) are the same as (6) and (1), respectively. They are rewritten here for the convenience of reference. When determining the effort level in (10), although the current payoff $u_i(m, \theta_i^t, d(\theta^t), e_i, e_{-i}^t)$ depends on the others' effort levels e_{-i}^t , the current payoff can be separated into the benefit which depends only on the others' effort e_{-i}^t , and the cost which depends only on agent i 's own effort e_i . Hence, when solving (10), agent i can treat the benefit as a constant, and consider only the cost, which depends on its own effort level and is known to itself.

The update of the parameter β_i in (12) ensures that the conjectured future payoff equals to the current payoff, namely $f_i(\alpha_i, \beta_i^{t+1}, \theta_i^t, d(\theta^t), e_i^t) = u_i(m, \theta_i^t, d(\theta^t), e_i^t)$. When the system converges to a CE $\{\theta_i^*, e_i^*, \beta_i^*\}_{i \in \mathcal{N}}$, we will have $f_i(\alpha_i, \beta_i^*, \theta_i^*, d(\theta^*), e_i^*) = u_i(m, \theta_i^*, d(\theta^*), e_i^*)$, which fulfills the third requirement of "consistent conjectures" in the definition of CE.

Next, we will provide the design guidelines on the matching rules, such that the above dynamics (10)–(12) always converge to a CE from any initial ratings. In fact, the design guideline is simple and intuitive: the matching rule should ensure that each agent's expected benefit is concave and increasing in its own rating.

Definition 2 (Desirable Matching Rules). A matching rule m is *desirable*, if under any rating profile θ ,

- each agent i 's (conjectured) expected benefit from the reviewing of its product, namely

$$\sum_{k_j \neq k_i} [m_{k_i k_j}(d(\theta)_{k_i}, d(\theta)_{k_j}) \cdot b_i(d(\theta)_{k_j})],$$

is concave and increasing in its own rating $d(\theta)_{k_i}$;

- each agent i 's expected number of products to review is positive and fixed, namely

$$\sum_{k_j \neq k_i} m_{k_j k_i}(d(\theta)_{k_j}, d(\theta)_{k_i}) = M > 0.$$

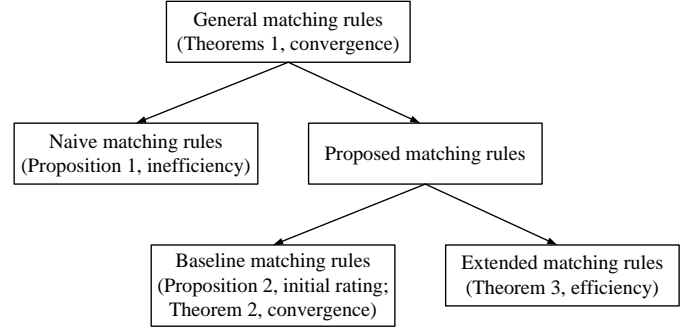


Fig. 1. Illustration of theoretical results.

The requirements of concavity and monotonicity are very reasonable. The expected benefit should be increasing in one's rating, such that one has incentives to exert high effort levels to increase its rating. In addition, if the expected benefit is concave in one's rating, since the marginal benefit is decreasing, one will not dramatically increase its effort level, which facilitates the convergence. The requirement of a fixed number of products to review ensures the fairness among the reviewers across time.

Despite the simplicity of the requirements for desirable matching rules, we are able to prove its convergence under proper rating update rules.

Theorem 1 (Convergence). Under any desirable matching rule, starting from any initial θ^0 , there exists $\bar{\mu} > 0$ such that under any small step size $\mu \in (0, \bar{\mu}]$ in the rating update rule, the system will converge to a CE through updates (10)–(12).

Proof: See Appendix B. \square

We illustrate the theoretical results in this paper in Fig. 1.

Remark 1 (Uniqueness). Theorem 1 proves that any desirable matching rule ensures the convergence to a CE. However, it is difficult to prove the convergence to a particular CE. The reasons are that there is not a single unique CE and the asymptotically reached CE (through our matching and update rules) depends on the history of the system and thus the probabilistic matching assignments and the initial ratings. In our technical analysis we prove the convergence by showing that under the updates (10)–(12), the difference (in terms of ℓ_1 norm) between two consecutive rating profiles strictly decreases over time. However, since the best responses are different under different rating profiles, the mappings from the current rating profile to the next one are different over time. Hence, the contraction mapping theorem does not apply here. In fact, in Section 5, we will show that under several desirable matching rules, there are *indeed multiple CEs*, and the system converges to different CEs under different initial ratings.

Remark 2 (Largest Step Size). For the convergence, we require the step size μ in the rating update rule to be small enough. However, we would like the step size to be as large as possible, subject to convergence, for two reasons. First, a larger step size results in a fast convergence of the ratings to the true review quality. Second, perhaps less obviously, a larger step size provides higher incentives

for agents to exert high effort levels. This is because with a larger step size, the influence of the current review quality is higher on the next rating. In Appendix B, we derive explicit upper bounds for the eligible step sizes.

5 DESIGN OF MATCHING RULES

The matching rule is the critical component of our design. In this section, we will first prove that existing matching rules are inefficient. Then we propose a baseline matching rule, and analyze the properties of this baseline matching rule in detail. Finally, we propose and study two extensions.

5.1 Inefficiency of Existing Matching Rules

We show the inefficiency of existing matching rules. In existing matching rules, the matching probabilities associated with agent i , $m_{k_i k_j}(d(\theta)_{k_i}, d(\theta)_{k_j})$, do not depend on i 's ranking k_i or its rating $d(\theta)_{k_i}$. This is true for existing matching rules in mechanism design [6]–[8], because there is no notion of effort and hence no rating. It is also true for existing matching rules in social norms [10]–[13], which are uniformly random.

Proposition 1. Under any matching rule that is independent of the rating of the agent whose product will be reviewed, namely $m_{k_i k_j}(d(\theta)_{k_i}, d(\theta)_{k_j}) = m_{k'_i k'_j}(d(\theta)_{k'_i}, d(\theta)_{k'_j})$, $\forall k_i, k'_i, k_j, k'_j, \theta$, there is a unique CE, in which $e_i^* = 0$ and $\theta_i^* = q_i(e_i^*) = 0$ for all i .

Proof: See Appendix II. \square

The above proposition shows that the matching rule that does not depend on agents' ratings is the worst-case matching rule that results in "free-riding" by everyone. This underlines the importance of designing efficient matching rules that take into account the ratings of both the reviewers and the agents whose products are reviewed.

5.2 Design of The Baseline Matching Rule

The baseline matching rule works as follows:

- 1) For the agents with the same rating, match their products among themselves using any one-to-one mapping that does not match one's product to itself.
- 2) For any agent i with a distinct rating (i.e., no other agent has the same rating),
 - a) If it has the highest rating (i.e., $k_i = 1$), match its product to a reviewer with the second highest rating with probability 1.
 - b) If it has the lowest rating (i.e., $k_i = K$), match its product to a reviewer with the second lowest rating with probability $\frac{d(\theta)_K}{d(\theta)_{K-1}}$. Hence, its product gets no review with probability $1 - \frac{d(\theta)_K}{d(\theta)_{K-1}}$.
 - c) If $1 < k_i < K$, match its product to its two "neighbors" with the following probabilities (which sum up to 1):

$$m_{k_i, k_i-1}(d(\theta)_{k_i}, d(\theta)_{k_i-1}) = \frac{d(\theta)_{k_i} - d(\theta)_{k_i+1}}{d(\theta)_{k_i-1} - d(\theta)_{k_i+1}},$$

and

$$m_{k_i, k_i+1}(d(\theta)_{k_i}, d(\theta)_{k_i+1}) = \frac{d(\theta)_{k_i-1} - d(\theta)_{k_i}}{d(\theta)_{k_i-1} - d(\theta)_{k_i+1}}.$$




Fig. 2. Illustration of the baseline matching rule with 7 agents. The highest-rating agent's product is matched to the reviewer with the second highest rating. The lowest-rating agent's product is matched to the reviewer with the second lowest rating with probability 0.5. The two agents with the same rating 0.4 are matched to each other. The rest are matched to their two nearest neighbors with probabilities inversely proportional to the distances in ratings.

The above matching rule is illustrated in Fig. 2. Agents with the same rating are matched to each other. For an agent with a distinct rating, it matches its product with its two nearest "neighbors" with probabilities that depend on how close its rating is to its neighbors' ratings.

We propose this matching rule, because it has the following desirable properties:

- No agent will have to review more than 3 products. This is because any agent will at most review a product from an agent with the same rating (if there is any), and two products from its neighbors (if they have distinct ratings).
- As we will prove later, this matching rule is a desirable matching rule as defined in Definition 2.

5.2.1 The Choice of Initial Ratings

In the considered system, it is important to choose the initial ratings correctly, because under different initial ratings, the system may converge to different CEs. Since the designer has no knowledge about the agents at the beginning, it is reasonable to assign the same initial rating to all the agents for fairness. In this case, the following proposition tells us that we should not make the initial rating too low.

Proposition 2. There always exists a rating θ , such that any initial rating profile with the same rating $\theta^0 \leq \theta$ for all agents is the equilibrium rating profile, and that each agent i chooses an equilibrium effort level e_i^* such that $q_i(e_i^*) = \theta^0$.

Proposition 2 implies that we should choose a high enough initial rating. In particular, when the initial rating is too low, it is optimal to choose an effort level e_i^* that satisfies $q_i(e_i^*) = \theta^0$. The key reason is that no agent has an incentive to reach a higher rating than the initial one,

because in this case it will get a *distinct* highest rating, and get the same benefit but a higher cost compared to choosing an effort level such that its rating remains the same as the initial rating. In other words, the initial rating determines the highest review quality produced by each agent.

5.2.2 Convergence

It is useful to classify agents into types based on their cost, review quality, and benefit functions, etc. We define agents of a certain type as follows.

Definition 3 (Types). The agents of the same *type* have the same normalized marginal benefit to cost ratio, defined as $\frac{\delta_i \alpha_i q'_i(\cdot)}{(1-\delta_i) c'_i(\cdot)}$, the same review quality function $q_i(\cdot)$, and the same marginal benefit function $b'_i(\cdot)$.

Definition 4 (Ordering of Capability). An agent i is *more capable* than an agent j , if

$$\begin{aligned} \frac{\delta_i \alpha_i q'_i(e)}{(1-\delta_i) c'_i(e)} &\geq \frac{\delta_j \alpha_j q'_j(e)}{(1-\delta_j) c'_j(e)}, \forall e, \\ q_i(e) &> q_j(e), \forall e, \\ b'_i(\theta) &\geq b'_j(\theta), \forall \theta. \end{aligned}$$

Definition 3 defines “types” of agents, in the sense that agents of the same type will always choose the same effort level and hence get the same rating. Definition 4 gives an ordering of agents in terms of their “capability”. We will prove that a more capable agent indeed gets a higher rating.

In the rest of this section, we make the following assumption about the population size.

Assumption 2 (Large Population). There is more than one agent of each type.

Assumption 2 is reasonable in practice, since the number of agents in peer review systems is indeed large. Given the same initial rating, the agents of the same type will choose the same best response effort level, and hence have the same rating. Assumption 2 ensures that for each agent, there is always another agent with the same rating. According to Property 1) in the baseline matching rule, each agent will always have exactly one product to review all the time.

Theorem 2. Suppose that the large population assumption (Assumption 2) holds. Then we have

- the baseline matching rule is a desirable matching rule;
- starting from any initial rating profile, there exists $\bar{\mu} > 0$ such that under any small step size $\mu \in (0, \bar{\mu}]$ in the rating update rule, the system will converge to a CE through the best response dynamics (10)–(12);
- if the agents have the same initial rating, at any point in the best response dynamics (10)–(12), more capable agents will always have no lower ratings than less capable agents.

Proof: See Appendix II. \square

Theorem 2 ensures the convergence of the best response dynamics to a CE. In fact, we can say something stronger about the best response dynamics. That is, a more capable agent never has lower ratings than a less capable agent *at any point* in the best response dynamics. This means that the rating mechanism can distinguish the agents of different




Fig. 3. An illustration of the first asymmetric extension of the baseline matching rule. We only show the matching of the agent with rating 0.8, who is matched to its two nearest neighbors with different probabilities than in the baseline rule.

types, and rank them in the correct order. Note that more capable agents produce reviews of high enough quality (that result in higher ratings) in their self-interest, as a result of maximizing their own payoffs; they are not obliged to do so by the designer.

5.3 Two Classes of Extended Matching Rules

Previously, we have focused on the baseline matching rule. The baseline matching rule is able to incentivize the agents to exert high effort levels by increasing the benefit obtained by an agent when its rating increases. Now we extend the baseline rule in two different ways, both of which result in a class of matching rules that allow us to tune the reward and/or punishment provided by the matching rules.

In the first extension, we assign asymmetric probabilities for matching an agent with a distinct rating to its higher and lower neighbors. In particular, the *asymmetric matching rule* is parametrized by γ such that any agent i with a distinct rating and with $k_i \in [2, K-1]$ is matched to its neighbors with the following probabilities:

$$m_{k_i, k_i-1}(d(\theta)_{k_i}, d(\theta)_{k_i-1}) = \left[\frac{d(\theta)_{k_i} - d(\theta)_{k_i+1}}{d(\theta)_{k_i-1} - d(\theta)_{k_i+1}} + \gamma \cdot \theta_i \right]_0^1,$$

and

$$m_{k_i, k_i+1}(d(\theta)_{k_i}, d(\theta)_{k_i+1}) = \left[\frac{d(\theta)_{k_i-1} - d(\theta)_{k_i}}{d(\theta)_{k_i-1} - d(\theta)_{k_i+1}} - \gamma \cdot \theta_i \right]_0^1,$$

where $[\cdot]_0^1 \triangleq \min\{\max\{\cdot, 0\}, 1\}$.

We illustrate this asymmetric extension in Fig. 3.

We can see that when $\gamma > 0$ ($\gamma < 0$), the resulting matching rule rewards (punishes) the agent by increasing its probability of being matched to the higher-rating (lower-rating) neighbor. When $\gamma = 0$, the asymmetric matching rule reduces to the baseline matching rule.

In the second extension, we allow an agent to be matched to a reviewer with even higher or even lower ratings than its two nearest neighbors. In particular, the matching rule is parametrized by $\gamma_r \in [0, 1]$ and $\gamma_p \in [0, 1]$. Then any agent i with a distinct rating and with $k_i \in [3, K-2]$ is matched to its neighbors and neighbors of neighbors with the following probabilities:

$$\begin{aligned} m_{k_i, k_i-1}(d(\theta)_{k_i}, d(\theta)_{k_i-1}) &= \frac{d(\theta)_{k_i} - d(\theta)_{k_i+1}}{d(\theta)_{k_i-1} - d(\theta)_{k_i+1}} \cdot (1 - \gamma_r), \\ m_{k_i, k_i-2}(d(\theta)_{k_i}, d(\theta)_{k_i-2}) &= \frac{d(\theta)_{k_i} - d(\theta)_{k_i+1}}{d(\theta)_{k_i-1} - d(\theta)_{k_i+1}} \cdot \gamma_r, \end{aligned}$$




Fig. 4. An illustration of the second long-range extension of the baseline matching rule. We show the matching of only one agent with rating 0.8, who is matched to its four neighbors, instead of two nearest neighbors as in the baseline rule.

and

$$m_{k_i, k_i+1}(d(\theta)_{k_i}, d(\theta)_{k_i+1}) = \frac{d(\theta)_{k_i-1} - d(\theta)_{k_i}}{d(\theta)_{k_i-1} - d(\theta)_{k_i+1}} \cdot (1 - \gamma_p),$$

$$m_{k_i, k_i+2}(d(\theta)_{k_i}, d(\theta)_{k_i+2}) = \frac{d(\theta)_{k_i-1} - d(\theta)_{k_i}}{d(\theta)_{k_i-1} - d(\theta)_{k_i+1}} \cdot \gamma_p.$$

We refer to this extension as *long-range matching rule*; see Fig. 4 for an illustration.

We can see that the parameters γ_r and γ_p reflect to what extent the agents are rewarded and punished, respectively. When $\gamma_r = \gamma_p = 0$, the matching rule reduces to the baseline rule. When $\gamma_r = 1$ ($\gamma_p = 1$), the agent is rewarded (punished) by being matched to a reviewer with the next higher (lower) rating.

We summarize the key differences among the baseline matching rule and its extensions in Fig. 5.

It is interesting to ask under each class of extended matching rules, which matching rule is optimal in terms of the equilibrium review quality? We first define the notion that one matching rule is “better” than the other.

Definition 5. We say that a matching rule m' is “better” than another matching rule m , if for any equilibrium rating profile θ^* under m , we can find an equilibrium rating profile $\theta^{*'}$ under m' that satisfies $\theta^{*'} > \theta^*$.

The following theorem tells us how to design an extended matching rule that is better than the baseline rule.

Theorem 3. Suppose that the large population assumption (Assumption 2) holds. Then we have:

- In the first asymmetric extension, there exists a $\gamma > 0$ (i.e., extra reward) under which the asymmetric matching rule is strictly better than the baseline matching rule.
- In the second long-range extension, there exists $\gamma_r = 0$ and $\gamma_p > 0$ (i.e., extra punishment) under which the long-range matching rule is strictly better than the baseline matching rule.

Theorem 3 tells us if we reward or punish by assigning higher or lower probabilities of being matched to the higher-rating neighbor, it is beneficial to reward. On the contrary, if we reward or punish by creating the possibility of being assigned to the next higher- or lower-rating neighbors, it is




Fig. 5. Comparison of the baseline matching rule and its two extensions. We illustrate how an agent with a certain rating (the red dot) is matched. In the baseline rule, it is matched to its nearest neighbors with equal probability. In the first asymmetric extension, although the distances between the ratings remain the same, the matching probabilities are slightly changed. In the second long-range extension, it can be matched to neighbors with even higher or lower ratings.

beneficial to punish. Note that we can get the benefit only when we set the correct parameters in the two extended matching rules. The technical reason is that we want to increase the marginal expected benefit when an agent’s rating is changed, in order to give more incentive for them to exert high effort levels. The main message delivered by our result is that we should carefully design the matching rule based on the way we reward and punish.

6 ILLUSTRATIVE RESULTS

We consider a system with 10 types of agents. There are 100 agents of each type. All the agents have the same patience $\delta_i = 0.8$, the same cost function $c_i(e_i) = e_i^2$, and the same benefit function $b_i(\theta) = -\theta^2 + 2\theta$. Different types of agents have different quality functions $q_i(e_i) = p_i \cdot e_i$ with $p_i = 0.2, 0.4, 0.6, \dots, 2.0$. They also have different $\alpha_i = 0.2, 0.4, 0.6, \dots, 2.0$ in the conjecture functions.

6.1 Impact of Step Sizes in Rating Update

We show the best response dynamics under step sizes $\mu = 0.1$ and $\mu = 0.3$ in Fig. 6 and Fig. 7, respectively. Note that we only show the ratings of agents of types 1, 3, 5, 7, 9. We first observe that under one-shot or exogenous matching rules, the reviewers exert 0 efforts and get 0 ratings all the time. Hence, the proposed endogenous matching greatly improves the performance of the system.

Second, we can see that under a larger step size, the agents’ equilibrium ratings are higher, indicating higher




Fig. 6. Convergence under step size $\mu = 0.1$.




Fig. 7. Convergence under step size $\mu = 0.3$.

equilibrium effort levels and higher equilibrium review quality. This is consistent with our intuition in Remark 2. Note that, when $\mu = 0.5$, the best response dynamics do not converge (type-1 agents' ratings are oscillating). This stresses the importance of choosing the step size: We want to choose a step size as high as possible for better review quality, subject to the constraint that the best response dynamics converge.

6.2 Different Matching Rules

We compare the sum review quality and the social welfare (i.e., the total benefit minus cost) at the equilibrium under different matching rules.

In Table 1, we evaluate the first asymmetric extension of matching rules under different parameters γ . We can see that in our setting, the optimal γ should be 0.1, which results in the highest sum review quality. This is consistent with our theoretical results: we can find a rewarding matching rule that outperforms the baseline rule. It is worth mentioning that the matching rule that maximizes the sum review quality may not be the one that maximizes the social welfare. This is reasonable because higher review quality also results in higher cost. In fact, in terms of social welfare, the optimal γ is -0.05, which results in lower review quality and thus lower cost. How the review quality and the social welfare are aligned depends on the benefit and cost functions.

In Table 2, we evaluate the second long-range extension of matching rules under different parameters γ_r and γ_p . We can see that the optimal sum review quality is achieved

TABLE 1
Equilibrium review quality and social welfare under the first asymmetric extension of matching rules.

γ	-0.2	-0.1	-0.05	0	0.05	0.1	0.2
Sum review quality	0.64	0.91	0.96	1.29	1.28	1.36	1.28
Social welfare	1.37	1.58	1.59	1.44	1.45	1.46	1.55

TABLE 2
Equilibrium review quality and social welfare under the second long-range extension of matching rules.

(γ_r, γ_p)	(0, 0)	(0, .5)	(0, 1)	(.5, 0)	(.5, .5)	(.5, 1)
Sum review quality	1.29	1.31	1.40	1.11	1.28	1.33
Social welfare	1.44	1.41	1.35	1.27	1.57	1.43

when $\gamma_r = 0$ and $\gamma_p = 1$, which is a matching rule that punishes to the most severe extent. The threat of being matched to an even lower-rating reviewer provides more incentive for agents to exert high effort. Again, such a matching rule does not result in the optimal social welfare. The optimal social welfare is achieved when $\gamma_r = 0.5$ and $\gamma_p = 0.5$, where the agents are also rewarded.

7 CONCLUSION

In this work, we proposed the first rating and repeated endogenous matching mechanisms to address the adverse selection and moral hazard problems simultaneously in peer review. Our proposed rating and matching mechanisms are easy to implement, require no knowledge of agents' private information, and ensure the convergence to an equilibrium in which the agents get their review quality revealed by their ratings and are incentivized to produce high-quality reviews. We thoroughly studied the design of matching rules, in terms of the initial ratings, the requirements for convergence, and the equilibrium ratings and review quality. We also studied extensions to different classes of matching rules, and proved the optimality of different rewarding/punishing mechanisms under different matching rules.

In future work, we will investigate the effect of inaccurate and possibly biased reports about the review quality, as well as more detailed modeling frameworks including multiple agents and reviewers associated with each product. An intriguing problem is the quest for the *best* matching rule maximizing the review quality.

APPENDIX A

PROPERTIES OF BEST RESPONSE

For ease of reference, we recall agent i 's best response e_i^t in time slot t here:

$$e_i^t = \arg \max_{e_i \in [0, e_i^{\max}]} (1 - \delta_i) \cdot u_i(m, \theta_i^t, d(\theta^t), e_i, e_{-i}^t) + \delta \cdot f_i(\alpha_i, \beta_i^t, \theta_i^t, d(\theta^t), e_i). \quad (13)$$

We denote the best response function by $B_i : \theta^t \mapsto e_i^t$. We need the following lemma summarizing the properties of the best response for future reference.

Lemma 1. Consider the current best response as defined in (13). The following statements hold:

- 1) The best response is unique for every $t \geq 0$;
- 2) The best response function is a continuous and almost everywhere differentiable function (uniformly in t) of θ^t ;
- 3) For any $t \geq 0$, the best response function satisfies $|B_i(\theta^{t+1}) - B_i(\theta^t)| \leq L_i(\mu) \cdot \|\theta^{t+1} - \theta^t\|_1$, where $L_i(\mu)$ is a continuous function of the step size μ that converges to zero as $\mu \downarrow 0$.

Proof: The current payoff $u_i(m, \theta_i^t, d(\theta^t), e_i, e_{-i}^t)$ in the current best response (13) has two terms:

$$\sum_{k_j \neq k_i} m_{k_i k_j}(\theta_i^t, d(\theta^t)_{k_j}) \cdot b_i(q_j(e_j^t)) - M \cdot c_i(e_i),$$

where the first term, namely the expected current benefit, does not depend on agent i 's own effort. Similarly, the conjecture function $f_i(\alpha_i, \beta_i^t, \theta_i^t, d(\theta^t), e_i) = \alpha_i \cdot \bar{b}_i(\theta_i^t, d(\theta^t), e_i^t) + \beta_i^t$, where the second term β_i^t does not depend on agent i 's own effort. Hence, we can simplify the best response as

$$e_i^t = \arg \max_{e_i \in [0, e_i^{\max}]} \begin{aligned} & -(1 - \delta_i) M c_i(e_i) \\ & + \delta_i \alpha_i \bar{b}_i(\theta_i^t, d(\theta^t), e_i^t). \end{aligned} \quad (14)$$

Since the cost $c_i(e_i)$ is strictly convex in e_i , since the grading quality $q_i(e_i)$ is concave in e_i , and since the conjectured benefit is concave and increasing in $(1 - \mu)\theta_i^t + \mu q_i(e_i)$, the objective function in (14) is strictly concave in e_i . Thus, the current best response is unique for each $t \geq 0$.

With regards to statement 2): due to strict concavity, the objective function in (14) is almost everywhere differentiable with respect to e_i and its derivative is equal almost everywhere to a decreasing function of e_i [18, Exercise 1.6.42]. Due to the Monotone Differentiation Theorem [18, Theorem 1.6.25], the derivative itself is almost everywhere differentiable with respect to e_i as well. We concluded that the objective function in (14) is almost everywhere twice differentiable with respect to e_i . As a consequence, the unique best response e_i^t is almost everywhere differentiable with respect to θ^t due to the Implicit Function Theorem [17, Ch. 5, Theorem 2.1]. In addition, since the objective function is continuous (any concave function is continuous [17, Ch. 6, Theorem 2.14]), and the feasible set does not depend on θ^t , the unique best response e_i is continuous in θ^t according to Berge's Theorem of Maximum [17, Ch. 7, Theorem 2.1]. In summary, the best response e_i is continuous and almost everywhere differentiable of θ^t . Notice that all of these properties hold uniformly in t .

With regards to statement 3): Instead of studying B_i directly, we study a closely related function $\hat{B}_i : \theta \mapsto \hat{e}_i^t$ defined as

$$\hat{e}_i^t = \arg \max_{e_i \in \mathbb{R}} -(1 - \delta_i) M c_i(e_i) + \delta_i \alpha_i \bar{b}_i(\theta_i^t, d(\theta^t), e_i^t).$$

Note that the above optimization problem is almost the same as (14), except that the feasible set is $e_i \in \mathbb{R}$ instead of $e_i \in [0, e_i^{\max}]$. In other words, the new optimization problem is unconstrained. Following the same logic, we know that \hat{B}_i is continuous and almost everywhere differentiable. Moreover, since the optimization problem is unconstrained, we do not need to worry about the boundary of the feasible set.

Hence, \hat{e}_i^t is the unique solution to the first-order condition (whenever the derivative exists):

$$\begin{aligned} 0 &= -(1 - \delta_i) M c_i'(\hat{e}_i^t) \\ &+ \delta_i \alpha_i \mu q_i'(\hat{e}_i^t) \cdot \sum_{k_j \neq k_i^+} \left[\frac{\partial m_{k_i^+ k_j}(\theta_i^{t+1}, d(\theta_i^{t+1}, \theta_i^t)_{k_j})}{\partial \theta_i^{t+1}} \right. \\ &\quad \left. \cdot b_i(d(\theta_i^{t+1}, \theta_i^t)_{k_j}) \right], \end{aligned} \quad (15)$$

where $\theta_i^{t+1} = (1 - \mu)\theta_i^t + \mu q_i(\hat{e}_i^t)$. Note that when $k_j \neq k_i^+$, $b_i(d(\theta_i^{t+1}, \theta_i^t)_{k_j}) = b_i(\theta_j^t)$ does not depend on θ_i^{t+1} . The best response function B_i is related to \hat{B}_i in the following way:

$$B_i(\theta^t) = \left[\hat{B}_i(\theta^t) \right]_0^{e_i^{\max}},$$

where $[\cdot]_0^{e_i^{\max}} = \min \{ \max \{ \cdot, 0 \}, e_i^{\max} \}$ is the projection of a real number on the interval $[0, e_i^{\max}]$. This is because at any θ^t , due to strict concavity of the objective function, the first-order derivative is decreasing and hence must be positive for all $e_i < \hat{B}_i(\theta^t)$. In other words, the objective function is increasing in e_i for all $e_i < \hat{B}_i(\theta^t)$. If $\hat{B}_i(\theta^t) > e_i^{\max}$, then the objective function is increasing in e_i for all $e_i \leq e_i^{\max}$. In this case, the maximum of the objective function in (14) is taken at $e_i^t = e_i^{\max}$. Similarly, $B_i(\theta^t) = 0$ if $\hat{B}_i(\theta^t) \leq 0$.

Now we can study $\hat{B}_i(\theta^t)$. In particular, we want to know how \hat{e}_i^t changes when θ^t changes. To this end, we apply the Implicit Function Theorem to (15) by taking the derivative of the right-hand side of (15) with respect to θ_j^t and keeping in mind that $\hat{e}_i^t = \hat{B}_i(\theta^t)$ is a function of θ_j^t . Then we obtain $\frac{\partial \hat{B}_i(\theta^t)}{\partial \theta_j^t}$ for $j \neq i$ in (16) and $\frac{\partial \hat{B}_i(\theta^t)}{\partial \theta_i^t}$ in (17) shown at the top of the next page, where $d_i^+ \triangleq d(\theta_i^{t+1}, \theta_{-i}^t)$ and $d_{i, k_j}^+ \triangleq d(\theta_i^{t+1}, \theta_{-i}^t)_{k_j} = \theta_j^t$ for $k_j \neq k_i$.

Based on the above derivations, we prove the following important property about the best response function B_i .

First, we have:

$$\begin{aligned} |B_i(\theta^{t+1}) - B_i(\theta^t)| &= \left| \left[\hat{B}_i(\theta^{t+1}) \right]_0^{e_i^{\max}} - \left[\hat{B}_i(\theta^t) \right]_0^{e_i^{\max}} \right| \\ &\leq \left| \hat{B}_i(\theta^{t+1}) - \hat{B}_i(\theta^t) \right|, \end{aligned}$$

where the inequality holds because the difference between the projections of two numbers on the interval is no larger than the difference between the two numbers.

Next, we write $\hat{B}_i(\theta^{t+1}) - \hat{B}_i(\theta^t)$ as a "telescoping sum" of N terms as follows:

$$\begin{aligned} & \hat{B}_i(\theta^{t+1}) - \hat{B}_i(\theta^t) \\ &= \sum_{j=1}^N \left[\hat{B}_i(\theta_1^{t+1}, \dots, \theta_{j-1}^{t+1}, \theta_j^{t+1}, \theta_{j+1}^t, \dots, \theta_N^t) \right. \\ &\quad \left. - \hat{B}_i(\theta_1^{t+1}, \dots, \theta_{j-1}^{t+1}, \theta_j^t, \theta_{j+1}^t, \dots, \theta_N^t) \right], \end{aligned}$$

where the j th term is the best response under the rating profile where the first j agents have ratings in period $t + 1$, minus the best response under the rating profile where the

$$\frac{\partial \hat{B}_i(\boldsymbol{\theta}^t)}{\partial \theta_j^t} = \frac{\mu \cdot \left\{ \delta_i \alpha_i q_i'(\hat{e}_i^t) \left[\frac{\partial^2 m_{k_i^+ k_j}(\theta_i^{t+1}, d_{i,k_j}^+)}{\partial \theta_i^{t+1} \partial \theta_j^t} b_i(\theta_j^t) + \frac{\partial m_{k_i^+ k_j}(\theta_i^{t+1}, d_{i,k_j}^+)}{\partial \theta_i^{t+1}} b_i'(\theta_j^t) \right] \right\}}{(1 - \delta_i) M c_i''(\hat{e}_i) - \mu \delta_i \alpha_i \left\{ q_i''(\hat{e}_i^t) \sum_{k \neq k_i^+} \left[\frac{\partial m_{k_i^+ k}(\theta_i^{t+1}, d_{i,k}^+)}{\partial \theta_i^{t+1}} b_i(d_{i,k}^+) \right] + \mu [q_i'(\hat{e}_i^t)]^2 \frac{\partial^2 m_{k_i^+ k_j}(\theta_i^{t+1}, d_{i,k_j}^+)}{\partial [\theta_i^{t+1}]^2} b_i(\theta_j^t) \right\}} \quad (16)$$

$$\frac{\partial \hat{B}_i(\boldsymbol{\theta}^t)}{\partial \theta_i^t} = \frac{\mu \cdot \left[(1 - \mu) \delta_i \alpha_i q_i'(\hat{e}_i^t) \sum_{k \neq k_i^+} \frac{\partial^2 m_{k_i^+ k}(\theta_i^{t+1}, d_{i,k}^+)}{\partial [\theta_i^{t+1}]^2} b_i(d_{i,k}^+) \right]}{(1 - \delta_i) M c_i''(\hat{e}_i) - \mu \delta_i \alpha_i \sum_{k \neq k_i^+} \left[q_i''(\hat{e}_i^t) \frac{\partial m_{k_i^+ k}(\theta_i^{t+1}, d_{i,k}^+)}{\partial \theta_i^{t+1}} b_i(d_{i,k}^+) + \mu [q_i'(\hat{e}_i^t)]^2 \frac{\partial^2 m_{k_i^+ k_j}(\theta_i^{t+1}, d_{i,k_j}^+)}{\partial [\theta_i^{t+1}]^2} b_i(d_{i,k}^+) \right]} \quad (17)$$

first $j - 1$ agents have ratings in period $t + 1$. As a result, we have:

$$\begin{aligned} & \left| \hat{B}_i(\boldsymbol{\theta}^{t+1}) - \hat{B}_i(\boldsymbol{\theta}^t) \right| \\ &= \sum_{j=1}^N \left| \hat{B}_i(\theta_1^{t+1}, \dots, \theta_{j-1}^{t+1}, \theta_j^{t+1}, \theta_{j+1}^t, \dots, \theta_N^t) \right. \\ & \quad \left. - \hat{B}_i(\theta_1^{t+1}, \dots, \theta_{j-1}^{t+1}, \theta_j^t, \theta_{j+1}^t, \dots, \theta_N^t) \right| \\ &= \sum_{j=1}^N \left| \int_{\theta_j^t}^{\theta_j^{t+1}} \frac{\partial \hat{B}_i(\theta_1^{t+1}, \dots, \theta_{j-1}^{t+1}, \theta_j, \theta_{j+1}^t, \dots, \theta_N^t)}{\partial \theta_j} d\theta_j \right| \\ &\leq \sum_{j=1}^N \left| \theta_j^{t+1} - \theta_j^t \right| \cdot \\ & \quad \text{ess sup}_{\theta_j \in [\theta_j^t, \theta_j^{t+1}]} \left| \frac{\partial \hat{B}_i(\theta_1^{t+1}, \dots, \theta_{j-1}^{t+1}, \theta_j, \theta_{j+1}^t, \dots, \theta_N^t)}{\partial \theta_j} \right|, \end{aligned}$$

where we take the essential supremum, since the derivative may not exist for all $\theta_j \in [\theta_j^t, \theta_j^{t+1}]$, but exists almost everywhere except on a subset of measure zero.

We define

$$L_{ij}(\mu) \triangleq \text{ess sup}_{\boldsymbol{\theta}} \frac{\partial \hat{B}_i(\boldsymbol{\theta})}{\partial \theta_i}(\mu), \text{ subject to } \frac{\partial \hat{B}_i(\boldsymbol{\theta})}{\partial \theta_i} \text{ exists,}$$

where we write $\frac{\partial \hat{B}_i(\boldsymbol{\theta})}{\partial \theta_i}(\mu)$ to emphasize the fact that $\frac{\partial \hat{B}_i(\boldsymbol{\theta})}{\partial \theta_i}$ is a function of the step size μ . We further define

$$L_i(\mu) \triangleq \max_{j \in \mathcal{N}} L_{ij}(\mu).$$

We can see from (16) and (17) that $\frac{\partial \hat{B}_i(\boldsymbol{\theta})}{\partial \theta_i}(\mu)$ is continuous in μ and is 0 when $\mu = 0$. Hence, due to Berge's Theorem of Maximum [17, Ch. 7, Theorem 2.1], $L_{ij}(\mu)$ is continuous in μ , and $L_{ij}(0) = 0$. Similarly, $L_i(\mu)$ is continuous in μ , and $L_i(0) = 0$. Therefore, we have

$$\begin{aligned} |B_i(\boldsymbol{\theta}^{t+1}) - B_i(\boldsymbol{\theta}^t)| &\leq \sum_{j=1}^N \left| \theta_j^{t+1} - \theta_j^t \right| \cdot L_{ij}(\mu) \\ &\leq L_i(\mu) \cdot \|\boldsymbol{\theta}^{t+1} - \boldsymbol{\theta}^t\|_1. \end{aligned}$$

□

APPENDIX B PROOF OF THEOREM 1

Before we prove the convergence of the updates (10)–(12), we show that if they converge to a triple $\{\theta_i^*, e_i^*, \beta_i^*\}_{i \in \mathcal{N}}$, this triple $\{\theta_i^*, e_i^*, \beta_i^*\}_{i \in \mathcal{N}}$ is a CE. Note that when the updates converge, each agent i has a fixed rating $\theta_i^{t+1} =$

$\theta_i^t = \theta_i^*$, a fixed effort level $e_i^t = e_i^*$, and a fixed payoff $u_i(m, \theta_i^t, d(\boldsymbol{\theta}^t), e^t) = u_i(m, \theta_i^*, d(\boldsymbol{\theta}^*), e^*)$. First, the best response (10) becomes

$$\begin{aligned} e_i^* &= \arg \max_{e_i \in [0, e_i^{\max}]} (1 - \delta_i) \cdot u_i(m, \theta_i^*, d(\boldsymbol{\theta}^*), e_i, e_{-i}^*) \\ & \quad + \delta_i \cdot f_i(\alpha_i, \beta_i^*, \theta_i^*, d^*, e_i), \end{aligned}$$

which is exactly the first requirement of ‘‘incentive compatibility’’ in the definition of CE. Second, the rating update (11) ensures $\theta_i^* = q_i(e_i^*)$, which fulfills the second requirement of ‘‘stable and correct rating’’ in the definition of CE. Finally, the update of β in (12) becomes

$$\beta_i^* = u_i(m, \theta_i^*, d(\boldsymbol{\theta}^*), e^*) - \alpha_i \cdot \bar{b}(\theta_i, d(\boldsymbol{\theta}^*), e_i^*).$$

Rearranging terms and using the definition of the conjecture function (4), we have

$$f_i(\alpha_i, \beta_i^*, \theta_i^*, d(\boldsymbol{\theta}^*), e_i^*) = u_i(m, \theta_i^*, d(\boldsymbol{\theta}^*), e^*).$$

In other words, the conjecture is indeed equal to the true payoff, which fulfills the third requirement of ‘‘correct conjectures’’ in the definition of CE.

Next, we prove that the updates (10)–(12) do converge under a small enough step size μ in the rating update rule. Consider a sequence of rating profiles $\{\boldsymbol{\theta}^t\}_{t=0}^\infty$ generated by the updates (10)–(12). Our goal is to prove that there exists a $\rho \in (0, 1)$ such that

$$\|\boldsymbol{\theta}^{t+2} - \boldsymbol{\theta}^{t+1}\|_1 \leq \rho \cdot \|\boldsymbol{\theta}^{t+1} - \boldsymbol{\theta}^t\|_1$$

for all $t \geq 0$. If the above is true, the sequence $\{\boldsymbol{\theta}^t\}_{t=0}^\infty$ will be a Cauchy sequence in \mathbb{R}_+^N , and hence will converge.

We prove the contraction property of the differences between consecutive rating profiles. Note that the rating update rule (11) is asynchronous. Hence, agent i 's rating will not be updated if it did not grade an assignment. Recall Lemma 1. For each agent i , there are two cases:

- If there is rating update at time $t + 1$, we have

$$|\boldsymbol{\theta}^{t+2} - \boldsymbol{\theta}^{t+1}| \leq L_i(\mu) \cdot \|\boldsymbol{\theta}^{t+1} - \boldsymbol{\theta}^t\|_1,$$

where $L_i(\mu)$ is as in Lemma 1.

- If there is no rating update at time $t + 1$, we have $\theta_i^{t+2} = \theta_i^{t+1}$, and hence

$$|\boldsymbol{\theta}^{t+2} - \boldsymbol{\theta}^{t+1}| = 0 \leq L_i(\mu) \cdot \|\boldsymbol{\theta}^{t+1} - \boldsymbol{\theta}^t\|_1.$$

Hence, the above inequality holds despite the asynchronous rating update.

Moreover, since q_i is concave and thus has a decreasing derivative, we have

$$\begin{aligned} |q_i(B_i(\theta^{t+1})) - q_i(B_i(\theta^t))| &\leq q'_i(0) |B_i(\theta^{t+1}) - B_i(\theta^t)| \\ &\leq q'_i(0) L_i(\mu) \cdot \|\theta^{t+1} - \theta^t\|_1. \end{aligned}$$

As a result, we have

$$\begin{aligned} \|\theta^{t+2} - \theta^{t+1}\|_1 &= \sum_{i=1}^N |\theta_i^{t+2} - \theta_i^{t+1}| \\ &\leq \sum_{i=1}^N [(1-\mu) \cdot |\theta_i^t - \theta_i^{t-1}| + \mu |q_i(B_i(\theta^{t+1})) - q_i(B_i(\theta^t))|] \\ &\leq (1-\mu) \|\theta^t - \theta^{t-1}\|_1 + \mu \cdot \left(\sum_{i=1}^N q'_i(0) L_i(\mu) \right) \cdot \|\theta^t - \theta^{t-1}\|_1 \\ &= \left[(1-\mu) + \mu \cdot \left(\sum_{i=1}^N q'_i(0) L_i(\mu) \right) \right] \cdot \|\theta^t - \theta^{t-1}\|_1. \end{aligned}$$

Due to property 3 of Lemma 1, we can find a small enough step size $\mu > 0$ such that $\sum_{i=1}^N q'_i(0) L_i(\mu) < 1$, and hence $\rho \triangleq (1-\mu) + \mu \cdot \left(\sum_{i=1}^N q'_i(0) L_i(\mu) \right) < 1$.

APPENDIX C PROOF OF RESULTS IN SECTION 5

Proof of Proposition 1: We analyze the best response defined in (14). According to (14), the best response maximizes an objective function that consists of two terms. Under a matching rule that is independent of the author's rating, the second term of the objective function in (14), namely

$$\delta_i \alpha_i \sum_{k_j \neq k_i^+} m_{k_i^+ k_j} \left((1-\mu) \theta_i^t + \mu q_i(e_i, \theta^t) \right) \cdot b_i(\theta_j^t),$$

is independent of i 's rating, and hence is independent of its effort level. Hence, i 's best response is an effort level that maximizes the first term of the objective function, namely

$$-(1-\delta_i) M c_i(e_i).$$

Since each agent i 's cost is strictly increasing in its effort level, its best response is then $e_i^* = 0$ under any rating profile θ^t . Since student i chooses $e_i^* = 0$ at any time t , its rating will converge to $\theta_i^* = q_i(0) = 0$. \square

Proof of Proposition 2: Given the same current initial rating θ^0 of everyone, based on Property 1) of the matching rule, agent i is matched to a reviewer with the same rating θ^0 with probability 1. Hence, agent i incurs a current cost of $-c_i(e_i)$ by choosing effort e_i .

Next, we look at agent i 's conjectured future payoff, which depends on the new rating $(1-\mu)\theta^0 + \mu q_i(e_i)$ when agent i exerts effort e_i . If $(1-\mu)\theta^0 + \mu q_i(e_i) \geq \theta^0$, namely $q_i(e_i) \geq \theta^0$, based on Properties 2)-a) and 2)-c) of the matching rule, agent i is matched to a reviewer with rating θ^0 with probability 1. If $(1-\mu)\theta^0 + \mu q_i(e_i) < \theta^0$, namely $q_i(e_i) < \theta^0$, based on Property 2)-b) of the matching rule, i is matched to a reviewer with rating θ^0 with probability $\frac{(1-\mu)\theta^0 + \mu q_i(e_i)}{\theta^0}$. In summary, i 's belief function at time 0, $f_i(\alpha_i, \beta_i, \theta^0, d(\theta^0), e_i)$, can be calculated as

$$\begin{cases} \alpha_i \cdot b_i(\theta^0) + \beta_i, & \text{when } q_i(e_i) \geq \theta^0 \\ \alpha_i \cdot \frac{(1-\mu)\theta^0 + \mu q_i(e_i)}{\theta^0} \cdot b_i(\theta^0) + \beta_i, & \text{when } q_i(e_i) < \theta^0 \end{cases}$$

At time 0, agent i chooses its effort level e_i^0 by solving the optimization problem (14), which is equivalent to solving

$$e_i^0 = \arg \max_{e_i} \{ -(1-\delta_i) c_i(e_i) + \delta_i f_i(\alpha_i, \beta_i, \theta^0, d(\theta^0), e_i) \}.$$

Since $f_i(\alpha_i, \beta_i, \theta^0, d(\theta^0), e_i)$ is a constant and does not increase with the effort level e_i when $q_i(e_i) \geq \theta^0$, we restrict ourselves (without loss of generality) to $q_i(e_i) \leq \theta^0$. To find the solution e_i^0 , we look at the derivative of the objective with respect to e_i for any $\theta^0 > 0$ and $e_i \in [0, q_i^{-1}(\theta^0)]$:

$$F_i(e_i, \theta^0) \triangleq -(1-\delta_i) c'_i(e_i) + \delta_i \alpha_i \mu q'_i(e_i) \frac{b_i(\theta^0)}{\theta^0}.$$

The left derivative of the objective with respect to e_i when $e_i \downarrow q_i^{-1}(\theta^0)$ is

$$\begin{aligned} F_i(q_i^{-1}(\theta^0), \theta^0) &\triangleq \lim_{e_i \uparrow q_i^{-1}(\theta^0)} F_i(e_i, \theta^0) \\ &= -(1-\delta_i) c'_i(q_i^{-1}(\theta^0)) + \delta_i \alpha_i \mu q'_i(q_i^{-1}(\theta^0)) \frac{b_i(\theta^0)}{\theta^0}. \end{aligned}$$

Since $q_i(e_i)$ is strictly increasing in e_i , $q_i^{-1}(\theta^0)$ is strictly increasing in θ^0 . Due to strict convexity of c_i , $c'_i(q_i^{-1}(\theta^0))$ is strictly increasing in θ^0 . Since q_i is concave and strictly increasing in e_i , $q'_i(q_i^{-1}(\theta^0))$ is positive and decreasing in θ^0 . Since b_i is concave and strictly increasing in θ^0 , $\frac{b_i(\theta^0)}{\theta^0}$ is positive and decreasing in θ^0 . As a result, $F_i(q_i^{-1}(\theta^0), \theta^0)$ is strictly decreasing in θ^0 .

Moreover, due to continuous differentiability assumptions of c_i , q_i , and b_i , $F_i(q_i^{-1}(\theta^0), \theta^0)$ is continuous in θ^0 .

In addition, due to our assumptions, $c'_i(0) = 0$, $q'_i(0) > 0$, and $\lim_{\theta^0 \downarrow 0} \frac{b_i(\theta^0)}{\theta^0} = b'_i(0) > 0$. Hence, we have $F_i(q_i^{-1}(0), 0) > 0$.

Since $F_i(q_i^{-1}(\theta^0), \theta^0)$ is strictly decreasing and continuous in θ^0 , and since $F_i(q_i^{-1}(0), 0) > 0$, there exists a small enough $\underline{\theta}_i$ such that $F_i(q_i^{-1}(\theta^0), \theta^0) \geq 0$ for any $\theta^0 \leq \underline{\theta}_i$. We define $\underline{\theta} \triangleq \max_i \underline{\theta}_i$. Then for any $\theta^0 \leq \underline{\theta}$, we have $F_i(q_i^{-1}(\theta^0), \theta^0) \geq 0$ for all i .

For any $\theta^0 \leq \underline{\theta}$, due to strict concavity of the objective function, the derivative of the objective with respect to e_i , $F_i(e_i, \theta^0)$, is decreasing in e_i and thus satisfies $F_i(e_i, \theta^0) \geq 0$ for all $e_i \leq q_i^{-1}(\theta^0)$. Therefore, each agent i chooses effort $e_i^0 = q_i^{-1}(\theta^0)$ at time 0, and gets the same rating $\theta_i^1 = \theta^0$ at time 1, under which it chooses the same effort $e_i^1 = q_i^{-1}(\theta^0)$. As a result, the system will stay the same. Hence, the equilibrium rating is θ^0 , and the equilibrium effort level is $e_i^* = q_i^{-1}(\theta^0)$. \square

Proof of Theorem 2:

Claims 1-2: First, we prove the first claim that the baseline matching rule is a desirable rule. Under the large population assumption (i.e., Assumption 2), there are multiple agents of the same type. These agents of the same type will choose the same effort level and hence have the same rating. According to Property 1) in the baseline matching rule, they will always have exactly one product (from another student of the same rating) to review, namely $M = 1$.

It remains to show that the conjectured expected benefit is increasing and concave in the effort level. Under the

baseline matching rule, the conjectured expected benefit can be explicitly computed as follows:

$$\begin{aligned}
& \sum_{k_j \neq k_i} [m_{k_i k_j}(\theta_i, d(\boldsymbol{\theta})_{k_j}) \cdot b_i(d(\boldsymbol{\theta})_{k_j})] \quad (18) \\
&= \frac{\theta_i - d(\boldsymbol{\theta})_{k_i+1}}{d(\boldsymbol{\theta})_{k_i-1} - d(\boldsymbol{\theta})_{k_i+1}} \cdot b_i(d(\boldsymbol{\theta})_{k_i-1}) \\
&+ \frac{d(\boldsymbol{\theta})_{k_i-1} - \theta_i}{d(\boldsymbol{\theta})_{k_i-1} - d(\boldsymbol{\theta})_{k_i+1}} \cdot b_i(d(\boldsymbol{\theta})_{k_i+1}) \\
&= \frac{b_i(d(\boldsymbol{\theta})_{k_i-1}) - b_i(d(\boldsymbol{\theta})_{k_i+1})}{d(\boldsymbol{\theta})_{k_i-1} - d(\boldsymbol{\theta})_{k_i+1}} \cdot \theta_i \\
&+ \frac{d(\boldsymbol{\theta})_{k_i-1} \cdot b_i(d(\boldsymbol{\theta})_{k_i+1}) - d(\boldsymbol{\theta})_{k_i+1} \cdot b_i(d(\boldsymbol{\theta})_{k_i-1})}{d(\boldsymbol{\theta})_{k_i-1} - d(\boldsymbol{\theta})_{k_i+1}}.
\end{aligned}$$

We can see that the expected benefit is a piecewise linear function of θ_i . When i is ranked at k_i , the slope of the function is $\frac{b_i(d(\boldsymbol{\theta})_{k_i-1}) - b_i(d(\boldsymbol{\theta})_{k_i+1})}{d(\boldsymbol{\theta})_{k_i-1} - d(\boldsymbol{\theta})_{k_i+1}}$. Since b_i is an increasing function of θ , we have $\frac{b_i(d(\boldsymbol{\theta})_{k_i-1}) - b_i(d(\boldsymbol{\theta})_{k_i+1})}{d(\boldsymbol{\theta})_{k_i-1} - d(\boldsymbol{\theta})_{k_i+1}} \geq 0$ for all k_i . Hence, the expected benefit is increasing in θ_i .

Since b_i is a concave and increasing function of θ , we have $\frac{b_i(d(\boldsymbol{\theta})_{k_i-1}) - b_i(d(\boldsymbol{\theta})_{k_i+1})}{d(\boldsymbol{\theta})_{k_i-1} - d(\boldsymbol{\theta})_{k_i+1}} \leq \frac{b_i(d(\boldsymbol{\theta})_{k_i}) - b_i(d(\boldsymbol{\theta})_{k_i+2})}{d(\boldsymbol{\theta})_{k_i} - d(\boldsymbol{\theta})_{k_i+2}}$. Hence, the slope decreases when i is ranked higher. Therefore, the expected benefit is concave in θ_i . Since $(1 - \mu)\theta_i + \mu q_i(e_i)$ is increasing and concave in e_i , the expected benefit

$$\sum_{k_j \neq k_i} [m_{k_i k_j}((1 - \mu)\theta_i + \mu q_i(e_i), d(\boldsymbol{\theta})_{k_j}) \cdot b_i(d(\boldsymbol{\theta})_{k_j})]$$

is increasing and concave in the effort level e_i .

In summary, the baseline matching rule is a desirable rule. The second claim on the convergence then follows due to Theorem 1.

Claim 3: Now we prove the third claim that a more capable agent always gets no lower ratings than a less capable agent. More specifically, we will prove that as long as a more capable agent has a no lower current rating, he/she will exert effort high enough such that his/her next rating is no lower than that of a less capable agent. Then under the same initial rating for all the agents, the third claim follows. Our proof strategy is to first derive the sufficient and necessary conditions for the best response effort levels, and then show (by contradiction) that if a less capable agent exerts a high effort level that results in a next rating higher than that of a more capable agent, this high effort level violates the conditions and cannot be the best response. The main technical difficulty arises since the best responses may be at the smooth or non-smooth points of the objective function, which results in different sufficient and necessary conditions and needs separate treatments.

Sufficient and necessary conditions for best responses: First, we look at i 's best response at time t :

$$\begin{aligned}
e_i^t &= \arg \max_{e_i \in [0, e_i^{\max}]} -(1 - \delta_i)c_i(e_i) + \\
&\delta_i \alpha_i \sum_{k_j \neq k_i} m_{k_i k_j}((1 - \mu)\theta_i^t + \mu q_i(e_i), d(\boldsymbol{\theta}^t)_{k_j}) \cdot b_i(d(\boldsymbol{\theta}^t)_{k_j}).
\end{aligned}$$

Under the baseline rule, if the next rating $\theta_i^{t+1} = (1 - \mu)\theta_i^t + \mu q_i(e_i)$ lies between $d(\boldsymbol{\theta}^t)_k$ and $d(\boldsymbol{\theta}^t)_{k+1}$ for some

$k \in \{1, \dots, N - 1\}$, the expected benefit is calculated as:

$$\begin{aligned}
& \sum_{k_j \neq k_i} m_{k_i k_j}((1 - \mu)\theta_i^t + \mu q_i(e_i), d(\boldsymbol{\theta}^t)_{k_j}) \cdot b_i(d(\boldsymbol{\theta}^t)_{k_j}) \\
&= \frac{b_i(d(\boldsymbol{\theta}^t)_k) - b_i(d(\boldsymbol{\theta}^t)_{k+1})}{d(\boldsymbol{\theta}^t)_k - d(\boldsymbol{\theta}^t)_{k+1}} \cdot [(1 - \mu)\theta_i^t + \mu q_i(e_i)] \\
&+ \frac{d(\boldsymbol{\theta}^t)_k \cdot b_i(d(\boldsymbol{\theta}^t)_{k+1}) - d(\boldsymbol{\theta}^t)_{k+1} \cdot b_i(d(\boldsymbol{\theta}^t)_k)}{d(\boldsymbol{\theta}^t)_k - d(\boldsymbol{\theta}^t)_{k+1}}.
\end{aligned}$$

Removing all the terms that are unrelated to e_i , the best response can be simplified into:

$$\begin{aligned}
e_i^t &= \arg \max_{e_i \in [0, e_i^{\max}]} -(1 - \delta_i)c_i(e_i) \quad (19) \\
&+ \delta_i \alpha_i \frac{b_i(d(\boldsymbol{\theta}^t)_k) - b_i(d(\boldsymbol{\theta}^t)_{k+1})}{d(\boldsymbol{\theta}^t)_k - d(\boldsymbol{\theta}^t)_{k+1}} \cdot \mu q_i(e_i).
\end{aligned}$$

Recall that the expected benefit is piecewise linear in θ_i^t . Hence, the objective function in (19) is a piecewise smooth function of e_i . There are N nonsmooth points where the effort e_i satisfies $\theta_i^{t+1} = (1 - \mu)\theta_i^t + \mu q_i(e_i) = d(\boldsymbol{\theta}^t)_k$ for some $k \in \{1, \dots, N\}$.

Since the objective function is strictly concave, the best response e_i^t should satisfy the following conditions. When e_i^t results in a nonsmooth point, namely $\theta_i^{t+1} = (1 - \mu)\theta_i^t + \mu q_i(e_i^t) = d(\boldsymbol{\theta}^t)_k$, the left derivative at e_i^t must be nonnegative and the right derivative at e_i^t must be nonpositive, i.e.:

$$(1 - \delta_i)c_i'(e_i^t) \leq \delta_i \alpha_i \mu q_i'(e_i^t) \cdot \frac{b_i(d(\boldsymbol{\theta}^t)_k) - b_i(d(\boldsymbol{\theta}^t)_{k+1})}{d(\boldsymbol{\theta}^t)_k - d(\boldsymbol{\theta}^t)_{k+1}},$$

and

$$(1 - \delta_i)c_i'(e_i^t) \geq +\delta_i \alpha_i \mu q_i'(e_i^t) \cdot \frac{b_i(d(\boldsymbol{\theta}^t)_{k-1}) - b_i(d(\boldsymbol{\theta}^t)_k)}{d(\boldsymbol{\theta}^t)_{k-1} - d(\boldsymbol{\theta}^t)_k}.$$

When e_i^t results in a smooth point, namely $(1 - \mu)\theta_i^t + \mu q_i(e_i^t) \in (d(\boldsymbol{\theta}^t)_k, d(\boldsymbol{\theta}^t)_{k+1})$, the derivative of the objective in (19) exists and should be 0, namely

$$\begin{aligned}
0 &= -(1 - \delta_i)c_i'(e_i^t) \quad (20) \\
&+ \delta_i \alpha_i \mu q_i'(e_i^t) \cdot \frac{b_i(d(\boldsymbol{\theta}^t)_k) - b_i(d(\boldsymbol{\theta}^t)_{k+1})}{d(\boldsymbol{\theta}^t)_k - d(\boldsymbol{\theta}^t)_{k+1}}.
\end{aligned}$$

Proof by contradiction: Now pick two arbitrary agents i and j with i being more capable than j . We prove that if $\theta_i^t \geq \theta_j^t$, their best responses must satisfy

$$(1 - \mu)\theta_i^t + \mu q_i(e_i^t) \geq (1 - \mu)\theta_j^t + \mu q_j(e_j^t).$$

We distinguish two cases depending on whether e_i^t results in a smooth point or not. In each case, we further separate our discussions into two subcases depending on whether e_j^t results in a smooth point or not.

- Case 1: e_i^t results in a smooth point, namely $(1 - \mu)\theta_i^t + \mu q_i(e_i^t) \in (d(\boldsymbol{\theta}^t)_k, d(\boldsymbol{\theta}^t)_{k+1})$. Then e_i^t must satisfy (20), which is equivalent to

$$\frac{(1 - \delta_i)c_i'(e_i^t)}{\delta_i \alpha_i \mu q_i'(e_i^t)} = \frac{b_i(d(\boldsymbol{\theta}^t)_k) - b_i(d(\boldsymbol{\theta}^t)_{k+1})}{d(\boldsymbol{\theta}^t)_k - d(\boldsymbol{\theta}^t)_{k+1}}.$$

Suppose that j chooses a best response e_j^t such that

$$(1 - \mu)\theta_i^t + \mu q_i(e_i^t) < (1 - \mu)\theta_j^t + \mu q_j(e_j^t).$$

Since $\theta_j^t \leq \theta_j^t$, and since $q_j(e) < q_i(e)$, we must have $e_j^t > e_i^t$, and hence $\frac{(1-\delta_j)c'_j(e_j^t)}{\delta_j\alpha_j q'_j(e_j^t)} > \frac{(1-\delta_j)c'_j(e_i^t)}{\delta_j\alpha_j q'_j(e_i^t)} \geq \frac{(1-\delta_i)c'_i(e_i^t)}{\delta_i\alpha_i q'_i(e_i^t)}$. Moreover, for any $k' \leq k$, we have

$$\frac{b_j(d(\theta^t)_{k'}) - b_j(d(\theta^t)_{k'+1})}{d(\theta^t)_{k'} - d(\theta^t)_{k'+1}} \leq \frac{b_i(d(\theta^t)_k) - b_i(d(\theta^t)_{k+1})}{d(\theta^t)_k - d(\theta^t)_{k+1}},$$

which leads to

$$\frac{(1-\delta_j)c'_j(e_j^t)}{\delta_j\alpha_j\mu q'_j(e_j^t)} > \frac{b_j(d(\theta^t)_{k'}) - b_j(d(\theta^t)_{k'+1})}{d(\theta^t)_{k'} - d(\theta^t)_{k'+1}}, \quad (21)$$

for any $k' \leq k$.

- Subcase 1: If e_j^t results in any smooth point $(1-\mu)\theta_j^t + \mu q_j(e_j^t) \in (d(\theta^t)_{k'}, d(\theta^t)_{k'+1})$ with $k' \leq k$, then (21) violates the condition (20).
- Subcase 2: If e_j^t results in any nonsmooth point $(1-\mu)\theta_j^t + \mu q_j(e_j^t) = d(\theta^t)_{k'}$ with $k' \leq k$, then (21) violates the optimality condition (20).

In summary, in both subcases, e_j^t cannot be the best response, which leads to contradiction. Hence, in Case 1 we must have

$$(1-\mu)\theta_i^t + \mu q_i(e_i^t) > (1-\mu)\theta_j^t + \mu q_j(e_j^t).$$

- Case 2: e_i^t results in a nonsmooth point, namely $(1-\mu)\theta_i^t + \mu q_i(e_i^t) = d(\theta^t)_k$. Then e_i^t must satisfy (20), which leads to

$$\frac{(1-\delta_i)c'_i(e_i^t)}{\delta_i\alpha_i\mu q'_i(e_i^t)} \geq \frac{b_i(d(\theta^t)_{k-1}) - b_i(d(\theta^t)_k)}{d(\theta^t)_{k-1} - d(\theta^t)_k}.$$

Suppose that j chooses a best response e_j^t such that

$$(1-\mu)\theta_i^t + \mu q_i(e_i^t) < (1-\mu)\theta_j^t + \mu q_j(e_j^t).$$

Then we must have $e_j^t > e_i^t$. Following the same logic that leads to (21), we have

$$\frac{(1-\delta_j)c'_j(e_j^t)}{\delta_j\alpha_j\mu q'_j(e_j^t)} > \frac{b_j(d(\theta^t)_{k'}) - b_j(d(\theta^t)_{k'+1})}{d(\theta^t)_{k'} - d(\theta^t)_{k'+1}}, \quad (22)$$

for any $k' \leq k-1$.

Similar to Case 1, since $(1-\mu)\theta_j^t + \mu q_j(e_j^t) > d(\theta^t)_k$, we have two subcases: $(1-\mu)\theta_j^t + \mu q_j(e_j^t) = d(\theta^t)_{k'}$, and $(1-\mu)\theta_j^t + \mu q_j(e_j^t) \in (d(\theta^t)_{k'}, d(\theta^t)_{k'+1})$, where $k' \leq k-1$. Since the inequality (22) violates (20) and (20), e_j^t cannot be the best response in either subcase.

Hence, in Case 2 we must have

$$(1-\mu)\theta_i^t + \mu q_i(e_i^t) \geq (1-\mu)\theta_j^t + \mu q_j(e_j^t).$$

We have shown that if a more capable agent has a current rating no lower than that of a less capable agent, then its next rating is no lower. Under the same initial rating, more capable agents always have no lower ratings. \square

Proof of Theorem 3: A rating profile θ^* is an equilibrium rating profile if each student i 's effort $e_i^* = q_i^{-1}(\theta_i^*)$ is the best response. Under the baseline matching rule, we have derived the sufficient and necessary conditions for e_i^* to be the best response in (20)–(20). Hence, a rating profile θ^* and the associated effort profile e^* with $e_i^* = q_i^{-1}(\theta_i^*)$ are

an equilibrium rating profile and the associated equilibrium effort profile if and only if: when $k_i = 1$,

$$(1-\delta_i)c'_i(e_i^*) \leq \delta_i\alpha_i\mu q'_i(e_i^*) \cdot \frac{b_i(d(\theta^*)_{k_i}) - b_i(d(\theta^*)_{k_i+1})}{d(\theta^*)_{k_i} - d(\theta^*)_{k_i+1}}; \quad (23)$$

when $k_i \in \{2, \dots, N-1\}$,

$$(1-\delta_i)c'_i(e_i^*) \leq \delta_i\alpha_i\mu q'_i(e_i^*) \cdot \frac{b_i(d(\theta^*)_{k_i}) - b_i(d(\theta^*)_{k_i+1})}{d(\theta^*)_{k_i} - d(\theta^*)_{k_i+1}}, \quad (24)$$

$$(1-\delta_i)c'_i(e_i^*) \geq \delta_i\alpha_i\mu q'_i(e_i^*) \cdot \frac{b_i(d(\theta^*)_{k_i-1}) - b_i(d(\theta^*)_{k_i})}{d(\theta^*)_{k_i-1} - d(\theta^*)_{k_i}}; \quad (25)$$

and when $k_i = N$,

$$(1-\delta_i)c'_i(e_i^*) \geq \delta_i\alpha_i\mu q'_i(e_i^*) \cdot \frac{b_i(d(\theta^*)_{k_i-1}) - b_i(d(\theta^*)_{k_i})}{d(\theta^*)_{k_i-1} - d(\theta^*)_{k_i}}. \quad (26)$$

Following the same procedure, we can show that under the asymmetric extension with $\gamma > 0$, a rating profile $\hat{\theta}^*$ and the associated effort profile \hat{e}^* with $\hat{e}_i^* = q_i^{-1}(\hat{\theta}_i^*)$ are an equilibrium rating profile and the associated equilibrium effort profile if and only if: when $k_i = 1$,

$$(1-\delta_i)c'_i(\hat{e}_i^*) \leq \delta_i\alpha_i\mu q'_i(\hat{e}_i^*) \left[\frac{b_i(d(\hat{\theta}^*)_{k_i}) - b_i(d(\hat{\theta}^*)_{k_i+1})}{d(\hat{\theta}^*)_{k_i} - d(\hat{\theta}^*)_{k_i+1}} + \gamma \right] \quad (27)$$

when $k_i \in \{2, \dots, N-1\}$,

$$(1-\delta_i)c'_i(\hat{e}_i^*) \leq \delta_i\alpha_i\mu q'_i(\hat{e}_i^*) \left[\frac{b_i(d(\hat{\theta}^*)_{k_i}) - b_i(d(\hat{\theta}^*)_{k_i+1})}{d(\hat{\theta}^*)_{k_i} - d(\hat{\theta}^*)_{k_i+1}} + \gamma \right] \quad (28)$$

$$(1-\delta_i)c'_i(\hat{e}_i^*) \geq \delta_i\alpha_i\mu q'_i(\hat{e}_i^*) \left[\frac{b_i(d(\hat{\theta}^*)_{k_i-1}) - b_i(d(\hat{\theta}^*)_{k_i})}{d(\hat{\theta}^*)_{k_i-1} - d(\hat{\theta}^*)_{k_i}} + \gamma \right] \quad (29)$$

and when $k_i = N$,

$$(1-\delta_i)c'_i(\hat{e}_i^*) \geq \delta_i\alpha_i\mu q'_i(\hat{e}_i^*) \left[\frac{b_i(d(\hat{\theta}^*)_{k_i-1}) - b_i(d(\hat{\theta}^*)_{k_i})}{d(\hat{\theta}^*)_{k_i-1} - d(\hat{\theta}^*)_{k_i}} + \gamma \right] \quad (30)$$

Observe that the defining inequalities for the baseline matching rule and those for the first extension are identical up to the additive term γ . For any equilibrium rating profile θ^* under the baseline rule m , we can define $\hat{\theta}^* = \theta^* + \varepsilon \cdot \mathbf{1}_N$, where $\varepsilon > 0$ is a small constant. In the following we show that a strictly positive and small enough $\varepsilon > 0$ leads to a strictly better equilibrium rating profile $\hat{\theta}^* > \theta^*$.

Since $\hat{\theta}_i^* = q_i(\hat{e}_i^*)$ and $\theta_i^* = q_i(e_i^*)$, and since q_i is strictly increasing in e_i , we have $\hat{e}_i^* = e_i^* + \varepsilon_i$, where $\varepsilon_i > 0$. Since c_i is convex and q_i is concave, we have $c'_i(\hat{e}_i^*) > c'_i(e_i^*)$ and $q'_i(\hat{e}_i^*) < q'_i(e_i^*)$. Since b_i is concave and increasing, we have

$$\frac{b_i(d(\hat{\theta}^*)_{k_i-1}) - b_i(d(\hat{\theta}^*)_{k_i})}{d(\hat{\theta}^*)_{k_i-1} - d(\hat{\theta}^*)_{k_i}} < \frac{b_i(d(\theta^*)_{k_i-1}) - b_i(d(\theta^*)_{k_i})}{d(\theta^*)_{k_i-1} - d(\theta^*)_{k_i}}$$

for any $k_i = 2, \dots, N$. Therefore, if the inequalities with “ \geq ”, i.e., (25) and (26), hold for θ^* and e^* , then the inequalities with “ \geq ”, i.e., (29) and (30), hold for $\hat{\theta}^*$ and \hat{e}^* as long as $\gamma > 0$ is small enough. For any such γ , if the inequalities with “ \leq ”, i.e., (23) and (24), hold for θ^* and e^* , then the inequalities with “ \leq ”, i.e., (27) and (28) hold for $\hat{\theta}^*$ and \hat{e}^* as long as $\varepsilon_i(\gamma) > 0$ is small enough. In conclusion, we can find an equilibrium rating profile $\hat{\theta}^*$ under the extended rule that satisfies $\hat{\theta}^* > \theta^*$.

Under the second long-range extension with $\gamma_r = 0$ and $\gamma_p > 0$, a rating profile $\hat{\theta}^*$ is an equilibrium rating profile iff:

$$0 \leq -(1 - \delta_i)c'_i(\hat{e}_i^*) + \delta_i\alpha_i\mu q'_i(\hat{e}_i^*) \cdot \left[\frac{b_i(d(\hat{\theta}^*)_{k_i}) - b_i(d(\hat{\theta}^*)_{k_i+1})}{d(\hat{\theta}^*)_{k_i} - d(\hat{\theta}^*)_{k_i+1}} + \gamma_p \frac{b_i(d(\hat{\theta}^*)_{k_i+1}) - b_i(d(\hat{\theta}^*)_{k_i+2})}{d(\hat{\theta}^*)_{k_i} - d(\hat{\theta}^*)_{k_i+1}} \right],$$

when $k_i = 1$,

$$0 \leq -(1 - \delta_i)c'_i(\hat{e}_i^*) + \delta_i\alpha_i\mu q'_i(\hat{e}_i^*) \cdot \left[\frac{b_i(d(\hat{\theta}^*)_{k_i}) - b_i(d(\hat{\theta}^*)_{k_i+1})}{d(\hat{\theta}^*)_{k_i} - d(\hat{\theta}^*)_{k_i+1}} + \gamma_p \frac{b_i(d(\hat{\theta}^*)_{k_i+1}) - b_i(d(\hat{\theta}^*)_{k_i+2})}{d(\hat{\theta}^*)_{k_i} - d(\hat{\theta}^*)_{k_i+1}} \right],$$

$$0 \geq -(1 - \delta_i)c'_i(\hat{e}_i^*) + \delta_i\alpha_i\mu q'_i(\hat{e}_i^*) \cdot \frac{b_i(d(\hat{\theta}^*)_{k_i-1}) - b_i(d(\hat{\theta}^*)_{k_i})}{d(\hat{\theta}^*)_{k_i-1} - d(\hat{\theta}^*)_{k_i}},$$

when $k_i \in \{2, \dots, N-1\}$, and

$$0 \geq -(1 - \delta_i)c'_i(\hat{e}_i^*) + \delta_i\alpha_i\mu q'_i(\hat{e}_i^*) \cdot \frac{b_i(d(\hat{\theta}^*)_{k_i-1}) - b_i(d(\hat{\theta}^*)_{k_i})}{d(\hat{\theta}^*)_{k_i-1} - d(\hat{\theta}^*)_{k_i}},$$

when $k_i = N$.

Following similar reasoning as in the first extension, we can find an equilibrium rating profile $\hat{\theta}^*$ under the second extended rule that satisfies $\hat{\theta}^* > \theta^*$. \square

REFERENCES

- [1] Y. Xiao, F. Dörfler, and M. van der Schaar, "Rating and Matching in Peer Review Systems," accepted by *Allerton 2014*.
- [2] J. Daniel, "Making sense of MOOCs: Musings in a maze of myth, paradox and possibility," *Journal of Interactive Media in Education*, 2012.
- [3] C. Piech, J. Huang, Z. Chen, C. Do, A. Ng, and D. Koller, "Tuned models of peer assessment in MOOCs," *Proc. International Conference on Educational Data Mining (EDM 2013)*, 2013.
- [4] J. Diez, O. Luaces, A. Alonso-Betanzos, A. Troncoso, and A. Bahamonde, "Peer assessment in MOOCs using preference learning via matrix factorization," *NIPS Workshop on Data Driven Education*, 2013.
- [5] N. B. Shah, J. K. Bradley, A. Parekh, M. Wainwright, and K. Ramchandran, "A case for ordinal peer evaluation in MOOCs," *NIPS Workshop on Data Driven Education*, 2013.
- [6] D. Gale and L. S. Shapley, "College admissions and the stability of marriage," *American Mathematical Monthly*, vol. 69, pp. 9-14, 1962.
- [7] A. Abdulkadiroğlu and T. Sönmez, "School choice: A mechanism design approach," *American Economic Review*, vol. 93, no. 3, pp. 729-747, 2003.
- [8] A. E. Roth, T. Sönmez, and M. Utku Üver, "Pairwise kidney exchange," *Journal of Economic Theory*, vol. 125, no. 2, pp. 151-188, 2005.
- [9] N. Arnosti, R. Johari, and Y. Kanoria, "Managing congestion in decentralized matching markets," 2014. Available at SSRN: <http://dx.doi.org/10.2139/ssrn.2427960>
- [10] M. Kandori, "Social norms and community enforcement," *Review of Economic Studies*, vol. 59, no. 1, pp. 63 - 80, 1992.
- [11] Y. Zhang and M. van der Schaar, "Peer-to-peer multimedia sharing based on social norms," *Elsevier Journal Signal Processing: Image Communication Special Issue on "Advances in video streaming for P2P networks"*, vol. 27, no. 5, pp. 383-400, May 2012.
- [12] Y. Zhang and M. van der Schaar, "Rating protocols for online communities," *ACM Transactions on Economics and Computation*, 2013.
- [13] Y. Xiao and M. van der Schaar, "Socially-optimal design of service exchange platforms with imperfect monitoring," *Submitted*, 2013. Available: <http://arxiv.org/abs/1310.2323>
- [14] D. Kurokawa, O. Lev, J. Morgenstern, and A. D. Procaccia, "Impartial Peer Review," *Working paper*, 2014. Available: <http://www.cs.cmu.edu/~arielpro/papers/impartial.pdf>
- [15] A. Spalvieri, S. Mandelli, M. Magarini, and G. Bianchi, "Weighting peer reviewers," *Proc. International Conference on Privacy, Security, and Trust (PST)*, pp. 414-419, 2014.
- [16] F. H. Hahn, "Exercises in conjectural equilibria," *The Scandinavian Journal of Economics*, Vol. 79, No. 2, pp. 210-226, 1977.

- [17] A. de la Fuente, *Mathematical methods and models for Economists*. Cambridge University Press, 2009.
- [18] T. Tao, *An introduction to measure theory*. American Mathematical Society.