Optimal Matchmaking Strategy in Two-sided Marketplaces

Peng Shi
USC Marshall School of Business, pengshi@usc.edu.

Online platforms that match customers with suitable service providers utilize a wide variety of matchmaking strategies: some create a searchable directory of one side of the market (i.e., Airbnb, Google Local Finder); some allow both sides of the market to search and initiate contact (i.e., Care.com, Upwork); others implement centralized matching (i.e., Amazon Home Services, TaskRabbit). This paper compares these strategies in terms of their efficiency of matchmaking, as proxied by the amount of communication needed to facilitate a good market outcome. The paper finds that the relative performance of the above matchmaking strategies is driven by whether the preferences of agents on each side of the market are easy to describe. Here, "easy to describe" means that the preferences can be inferred with sufficient accuracy based on responses to standardized questionnaires. For markets with suitable characteristics, each of the above matchmaking strategies can provide near-optimal performance guarantees according to an analysis based on information theory. The analysis provides prescriptive insights for online platforms.

Key words: market design, online platforms, two-sided matching, communication complexity.

---

1. Introduction

How can an online platform best help customers find suitable service providers at good prices? This is a central question faced by the designers of online platforms such as Amazon Home Services, Care.com, Google Local Finder, HomeAdvisor, TaskRabbit, Thumbtack, Upwork and Yelp. These platforms can be referred to as two-sided marketplaces, and their value proposition is that they reduce the time and effort customers and service providers expend to find suitable partners. Inasmuch as these platforms are able to maximize the quality of the matches formed while minimizing the hassle of the search process, they generate value to society and gain market share in their respective industries.

Among existing two-sided marketplaces, the design of their search and matching features can largely be classified into the following four strategies:

1. **Customers search** for providers: these platforms create a searchable directory of providers and customers initiate contact with providers they are interested in. Each provider is asked to create an informative profile that customers can browse, and the platform offers customers various options for filtering and sorting providers.

2. **Providers search** for customers: these platforms ask customers to complete certain questionnaires to post a job request, and providers initiate contact with customers to bid for jobs. The platform offers providers various options for filtering and sorting job requests, and gives them the opportunity to be notified whenever jobs satisfying certain criteria are posted.
3. **Both sides search** for partners: these platforms combine features of the above two designs and allow both sides of the market to search for and initiate contact with potential partners.
4. **Centralized matching:** these platforms ask both customers and providers to complete certain questionnaires, and use this information to recommend a limited number of suitable providers to each customer.

The academic literature provides little guidance on what matchmaking strategy to use for a given market. Meanwhile, practitioners are experimenting with various strategies. As illustrated in Table 1, all four matchmaking strategies are observed among platforms that help customers find local professionals, and platforms sometimes migrate from one strategy to another in an effort to facilitate better matches for their intended user base.

Table 1 Disparate matchmaking strategies utilized by platforms that are used to find local professionals. As noted in the table, TaskRabbit and Thumbtack migrated from a design based on providers bidding on job requests to one based on centralized matching in 2014 and 2018 respectively (Newton 2014, Shieber 2018). Yelp modified its design in 2020 from only allowing customers to reach out to providers to also allowing providers to reach out to customers via its "Nearby Jobs" product (Wu 2020).

Matchmaking strategy	Examples of platforms
1. Customers search	Google Local Finder, Houzz, Yelp (before 2020), ...
2. Providers search	Bark, Porch, TaskRabbit (before 2014), Thumbtack (before 2018), ...
3. Both sides search	Care.com, Upwork, Yelp (after 2020), ...
4. Centralized matching	Amazon Home Services, Angi, HomeAdvisor, TaskRabbit (after 2014), Thumbtack (after 2018), ...

1.1. Contributions

This paper develops a tractable framework for studying the optimal matchmaking strategy in a two-sided marketplace. The analysis is based on communication complexity theory, which is a branch of information theory that studies the minimum amount of communication needed to solve a given problem (Kushilevitz and Nisan 2006). In the context of two-sided marketplaces, the "given problem" is to facilitate a good market outcome, so that users are happy with whom they are matched with and with the transaction price. This paper shows that each of the above four matchmaking strategies is able to achieve a good market outcome, but their relative efficiencies depend on certain characteristics of the market that will be described in the next paragraph. For markets with suitable characteristics, each of the above strategies is able to approximately minimize the amount of communication needed in the worst case to achieve a good outcome. Here, the minimization is over the space of all possible platform designs, so that even if a platform were to invent a new matchmaking process that is entirely different from the four strategies modeled here, it cannot achieve a much better performance, at least according to the metrics analyzed here.

The analysis yields prescriptive guidelines for platform design, as summarized in Table 2, which gives a high level description of the characteristics of suitable markets for each of the four matchmaking strategies. The central question is whether the preferences on each side of the market are easy to describe. Here, "easy to describe" means that the preferences can be accurately approximated based on the agents' responses to a standardized questionnaire, such as those typically used by online platforms when users create a profile or post a request. As an example, providers of short-term rentals, such as Airbnb hosts, generally have easy-to-describe preferences, as they can readily communicate the availabilities of their units via a calendar and are generally okay with housing any customer who has a record of being clean, respectful and paying on time. On the other hand, customers may have picky idiosyncratic preferences on location and amenities that are hard to fully describe upfront. According to the first row of Table 2, platforms for short-term rentals should create a directory of listings and have customers search and initiate contact, which is what Airbnb already does. The near-optimality result in this paper suggests that it would not be worthwhile for Airbnb to overhaul its matchmaking strategy.

Table 2 Characteristics of suitable markets for each of the four commonly observed matchmaking strategies. The characteristics are restated in terms of mathematically precise parameters in Section 3.

Matchmaking strategy	Characteristics of suitable markets
1. Customers search	*Provider* preferences are easy to describe.
2. Providers search	*Customer* preferences are easy to describe.
3. Both sides search	*Neither* side's preferences are easy to describe.
4. Centralized matching	*Both* sides' preferences are easy to describe.

Analogously, the second row of Table 2 states that having providers search for customers is suitable when customer preferences are easy to describe. One example is the market for data entry, in which a customer's job specifications can be succinctly described upfront, whereas providers may have idiosyncratic preferences about what kind of jobs they are interested in and when they have the capacity to take on a new job. This justifies the matchmaking strategy adopted by Amazon Mechanical Turk, which is based on providers searching through a directory of job requests.

The third row of Table 2 suggests that allowing both sides of the market to search for partners is important in markets such as those for childcare providers, in which both customers and providers have picky idiosyncratic preferences that are hard to systematically describe upfront. This supports the high level design choice of platforms such as Care.com, which is a platform for finding nannies that allows both sides of the market to search for partners and initiate contact. In 2020, Yelp also migrated to a design in which a customer not only can search through provider profiles but also can post job requests, which are sent to nearby providers, who choose which customers to contact (Wu 2020). According to the analysis in this paper, allowing both sides of the market to initiate contact allows Yelp to more efficiently facilitate good matches in markets such as those for non-standard home improvement projects, in which both sides may have hard-to-describe preferences.

The final row of Table 2 suggests that centralized matching is suitable in markets for standardized services, such as carpet cleaning, furniture assembly, ground transportation, and TV mounting. In such markets, customers' requirements are easy to describe, and providers have similar costs for jobs with similar descriptions. Hence, it is unnecessary to ask agents to search over the entire directory of potential partners, as the platform can facilitate a near-optimal match simply by recommending a few suitable providers based on the scope, location and timing of service.

1.2. Comparison with previous approaches

Before going into the mathematical details, it is instructive to compare the paper's overall approach with previous studies, so as to motivate the modeling choices. The broader literature that this paper relates to is the one on reducing congestion in matching markets, which originates in a series of seminal studies by Alvin Roth and co-authors (Roth and Xing 1994, 1997, Avery et al. 2001, Roth 2008). Market congestion refers to having too many interactions in the market, so that the market is unable to process information quickly enough to reach a desirable outcome. Li and Netessine (2019) use the merger of two online platforms as a natural experiment to show that due to market congestion, increasing the market size may reduce the efficiency of matching. Therefore, as a two-sided marketplace grows its number of users, effectively managing market congestion is crucial for maintaining the platform's value proposition.

Several papers use field experiments to show that relatively simple changes to platform design can result in significant improvements in the efficiency of matchmaking. For example, Coles et al. (2010) and Lee and Niederle (2015) demonstrate that allowing market participants to signal their preferences to a limited number of partners gives them significantly higher chances of obtaining a positive response. Fradkin (2017) documents that on Airbnb, the simple intervention of tracking listing availability results in a reduction in host rejections by 59%, which is significant because conditional on being rejected from their first query, customers are 43% to 70% less likely to eventually book a listing. Horton (2017) documents that a match recommendation algorithm at oDesk led to a 20% higher fill rate for job postings, and Horton (2018) estimates that allowing workers to signal their capacity leads to a 6% increase in market surplus.

However, field experiments are insufficient to study the question of optimal matchmaking, as migrating to a different matchmaking strategy is expensive and may alienate the platform's user base, so it is best to perform extensive theoretical analysis before commencing such an experiment. Moreover, experiments on the search and matching features of a platform typically need to be conducted at a market level, and users need time to adapt to a new design. This limits the number of experiments that can be conducted, and further increases the appeal of theoretical analysis.

A theoretical analysis of congestion in two-sided marketplaces needs to account for two objectives: 1) achieving a good market outcome; and 2) limiting the number of interactions among agents, so as to limit congestion. There are two approaches in the literature for optimizing these objectives.

1. **Equilibrium-based approach:** Fix a budget on the number of possible interactions or a cost per interaction, and maximize the quality of the market outcome that would arise in equilibrium. Previous studies using this approach find that in certain platform designs, agent welfare can be improved by limiting the number of partners an agent can contact (Halaburda et al. 2018, Arnosti et al. 2019), having the side with fewer agents initiate contact (Kanoria and Saban 2017), or giving agents a small number of signals to indicate special interest for certain partners (Coles et al. 2013, Abdulkadiroğlu et al. 2015, Jagadeesan and Wei 2018). While these insights are valuable, a limitation of the equilibrium-based approach is that in order to pin down an equilibrium that is tractable for analysis, one must make strong assumptions on the exact process by which matchmaking is conducted. This implies that any result on optimality is narrow in scope, as the result may not hold if the space of feasible designs is different from what is modeled. For example, none of the above papers allow providers to set their own prices, which is a first-order feature of many markets of interest.
2. **Communication complexity approach:** Fix a given notion of a good market outcome, and minimize the amount of interactions needed to reach such an outcome, as proxied by the amount of information agents need to communicate with one another. An advantage of this approach is that it can be used to establish robust optimality guarantees that are based on information theory rather than on a particular equilibrium model: a system of matchmaking is optimal if achieves a good market outcome using the minimum amount of communication possible, where the minimum is taken over all possible systems of multi-agent communication. However, previous studies that use this approach, such as Segal (2007), Chou and Lu (2010), Gonczarowski et al. (2015) and Ashlagi et al. (2019), all define a "good market outcome" as a stable matching in the model of Gale and Shapley (1962), which results in two limitations:
    i) The model of Gale and Shapley (1962) does not allow for transfer payments, so it is a poor fit for studying matching markets for paid services, as service providers typically set prices endogenously based on supply and demand.

ii) Requiring the market to reach an exactly stable matching can be overly restrictive, which is why previous results based on this approach are mostly impossibility results: Segal (2007), Chou and Lu (2010), and Gonczarowski et al. (2015) show that when preferences are adversarially constructed, obtaining a stable matching with high probability requires $\Omega(n)$ bits of communication per agent, where n is the number of agents on each side. This asymptotic rate is the same as requiring agents to communicate their preferences for the entire market, which suggests that a high level of congestion is inevitable. Ashlagi et al. (2019) bypass the above impossibility results by assuming that the preferences of at least one side are independently drawn. Under such an assumption, they show that $\Omega(\sqrt{n})$ bits of communication per agent are required, and $O(\sqrt{n}\,\text{polylog(n)})$ bits per agent suffice. This is the strongest positive result in previous papers based on the communication complexity approach.[1]

This paper overcomes the above limitations by adopting the communication complexity approach but using an alternative definition of a "good market outcome," which is based on a relaxed notion of stability and allows for endogenous prices. The relaxed notion of stability corresponds to assuming that agents have a certain tolerance for sub-optimality, so that they have a bounded level of pickiness for their preferred partners. Using this definition of a "good market outcome," this paper obtains much stronger positive results: Section 4 shows that for large classes of markets, $O(\log n)$ bits of communication per agent suffice to achieve a good market outcome with high probability, and this is much better than the $O(\sqrt{n}\,\text{polylog}(n))$ guarantee in Ashlagi et al. (2019). More importantly, this paper's modeling framework is able to rationalize each of the four matchmaking strategies commonly observed in practice, which brings the theory much closer to practice than in previous studies using the communication complexity approach.

1.3. Organization of the paper

Section 2 describes the model and defines concrete versions of the four matchmaking strategies in Table 1. Section 3 analyzes the efficiency of each strategy and derives the characteristics of suitable markets in Table 2. Section 4 defines concepts from communication complexity theory and uses them to establish the near-optimality results. For clarity of exposition, the bulk of the paper focuses on a static one-to-one matching model with independently drawn preferences and endogenous prices. Section 5 shows that the insights continue to hold under richer assumptions, including correlated preferences, many-to-one matching, dynamic arrival of customers, and exogenous prices. Section 6 concludes with managerial insights for two-sided marketplaces.

[1] Ashlagi et al. (2019) also show that $O(\log^2 n)$ bits of communication per agent suffice for finding a stable matching with high probability if one makes the following highly stylized assumptions on preferences: both sides of the market can be partitioned into tiers; agents always prefer partners from higher tiers to those from lower tiers, and preferences are uniformly random within each tier.

2. Model

This paper measures the efficiency of a matchmaking strategy in a given market based on the amount of communication it requires to reach a good market outcome. A formal treatment requires defining a "matchmaking strategy," a "market," the "amount of communication" and a "good market outcome." Sections 2.1 and 2.2 define a "market." Section 2.3 defines a "good market outcome." Section 2.4 models the four matchmaking strategies described in the introduction. Various metrics for measuring the "amount of communication" are given in Sections 3 and 4.

2.1. Assignment game

Before defining a "market" in Section 2.2, we first review the assignment game model of Shapley and Shubik (1971), which we build on later. An *assignment game* is specified by a tuple (I, J, b, c) where I is a finite set of customers, J is a finite set of service providers, and matrices $b, c \in \mathbb{R}^{I\times J}$ are the *preferences* of customers and providers respectively. For each pair $(i, j) \in I \times J$, b_{ij} represents the benefit customer i derives from provider j's service, and c_{ij} represents the provider's cost for serving customer i; the total surplus of a match between them is $b_{ij} - c_{ij}$. Note that the same customer may have different benefit values for matching with different providers, and the same provider may have different cost values for matching with different customers. For simplicity, assume for now that agents are interested in matching with at most one partner. Without loss of generality, assume that $|I| = |J|$, as one can always add dummy customers or providers that represent being unmatched.[2] Define $n = |I| = |J|$ to be the number of agents on each side of the market, and label both I and J as $[n] := \{1, 2, \cdots, n\}$.

An *outcome* to the assignment game is a tuple (x, u, v), where x is a $n \times n$ binary matrix with $x_{ij} = 1$ if customer i is matched to provider j and $x_{ij} = 0$ otherwise; $u \in \mathbb{R}^n$ is a vector in which u_i denotes the surplus of customer $i \in I$; and $v \in \mathbb{R}^n$ is a vector in which v_j denotes the surplus of provider $j \in J$. An outcome is *feasible* if it satisfies the following constraints:

$$\text{(Unit Demand)} \qquad \sum_{j\in J} x_{ij} \le 1 \qquad \text{for each } i \in I. \tag{1}$$

$$\text{(Unit Capacity)} \qquad \sum_{i\in I} x_{ij} \le 1 \qquad \text{for each } j \in J. \tag{2}$$

$$\text{(No Side Payments)} \qquad u_i + v_j = b_{ij} - c_{ij} \qquad \text{whenever } x_{ij} = 1. \tag{3}$$

[2] For example, to add a dummy customer i, define $b_{ij} = c_{ij} = 0$ for all $j \in J$. Any provider matched to this customer might as well be unmatched.

Any $n \times n$ binary matrix that satisfies constraints (1) and (2) is called a *matching*. Constraint (3) represents the assumption that individual surpluses are based on redistributing the total surplus among each pair of matched agents.[3]

An equivalent representation of a feasible outcome (x, u, v) is a *transaction outcome* (x, p) where x is a matching and $p \in \mathbb{R}^n$ is a vector specifying how much each customer pays the matched provider: if customer i is unmatched, then $p_i = 0$. Otherwise, if customer i is matched to provider j, then $u_i = b_{ij} - p_i$ and $v_j = p_i - c_{ij}$. There is a bijection between feasible outcomes and transaction outcomes, and we may work with one or the other depending on which is more convenient.

2.2. Market

While Shapley and Shubik (1971) assume full information, this paper assumes that initially, only customer i knows the value of b_{ij} and only provider j knows the value of c_{ij}, so finding a good outcome requires communication. We assume that b_{ij} is drawn from a commonly known distribution B_{ij}, and c_{ij} is drawn from a commonly known distribution C_{ij}. For simplicity, assume that the draws are independent across $(i, j) \in I \times J$ and independent between benefits and costs. (Section 5.1 discusses various ways in which these independence assumptions can be relaxed.) The B_{ij}'s and C_{ij}'s are referred to as the *preference distributions* of customers and providers. These distributions are allowed to be different for each pair (i, j), so the model can capture rich heterogeneity in preferences.

A *market* M is defined by the tuple (I, J, B, C). Mathematically speaking, a market represents a distribution over assignment games on the same set of agents. In the context of online platforms, the preference distributions B_{ij}'s and C_{ij}'s should be interpreted as the market's beliefs about the preferences of agents after conditioning on their initial descriptions of who they are and what they want. Platforms typically collect such information in the form of questionnaires, user profiles, or search options. For customers, such information may include where they live, what type of service they are looking for, and when they desire the service. For providers, such information may include their service region, the types of services they provide, and their qualifications. After conditioning on such information, the residual uncertainty that is represented by the distributions B_{ij}'s and C_{ij}'s should be interpreted as the *hard-to-describe* components of preferences, which are idiosyncratic in nature and difficult to articulate upfront. The significance of this interpretation will be explained in Section 3.

[3] In the textbook treatment of the assignment game such as in Roth and Sotomayor (1990), (3) is not part of the definition of a feasible outcome, but the following weaker contraint is imposed:

$$\text{(Budget Balance)} \qquad \sum_{i \in I} u_i + \sum_{j \in J} v_j = \sum_{i \in I, j \in J} (b_{ij} - c_{ij}) x_{ij}. \tag{4}$$

However, Roth and Sotomayor (1990) derive (3) as a consequence of stability (see Section 2.3). For ease of exposition, we assume (3) upfront, and (3) implies (4). In the context of two-sided marketplaces, (3) is a natural assumption, as there are typically no side payments between agents who are not matched to one another.

2.3. An ideal matchmaking outcome

Given an assignment game (I, J, b, c) and a parameter $\epsilon > 0$, a feasible outcome (x, u, v) is said to be *ϵ-stable* if it satisfies the following two conditions:

$$\text{(No } \epsilon\text{-Blocking Pair)} \qquad u_i + v_j + 2\epsilon \geq b_{ij} - c_{ij} \qquad \text{for each } (i, j) \in I \times J. \tag{5}$$

$$\text{(Individual Rationality)} \qquad u, v \geq 0 \tag{6}$$

Condition (5) says that if two agents (i, j) deviate from the matching x and match on their own, they cannot each get more than ϵ of additional surplus. On the left hand side, $u_i + v_j$ is their current total surplus; after giving the two agents ϵ each, the sum is at least equal to the right hand side, $b_{ij} - c_{ij}$, which is their total surplus from matching on their own. Condition (6) represents the assumption that each agent always has the option of being unmatched and receiving zero surplus.

In order to give ϵ a scale, assume that benefits and costs are bounded between 0 and 1:

$$\text{(Bounded Support)} \qquad b_{ij}, c_{ij} \in [0, 1] \qquad \text{for each } (i, j) \in I \times J. \tag{7}$$

The literature on assignment games focuses on exact stability, in which $\epsilon = 0$, but this paper allows for a positive ϵ in order to model the idea that agents have a certain tolerance for sub-optimality. For example, if switching partners results in a fixed cost of ϵ for each agent, then agents in an ϵ-stable outcome have no incentives to search for better partners. The existence of ϵ-stable outcomes follows from the existence of exactly stable outcomes, which is established by Shapley and Shubik (1971) using linear programming duality theory.

In this paper, the goal of a platform is to achieve an ϵ-stable outcome with minimal communication. (Other goals such as profit maximization are not modeled.) Achieving an ϵ-stable outcome is a desirable goal for the platform, as it guarantees near-optimal social welfare, as stated in the following proposition.

Proposition 1 **(Near-optimality in social welfare)**. *In any ϵ-stable outcome (x, u, v) of an assignment game (I, J, b, c), the average surplus per agent is within ϵ of the maximum possible: if x' is any other matching, then*

$$\frac{1}{2n}\left(\sum_{i\in[n]} u_i + \sum_{j\in[n]} v_j\right) = \frac{1}{2n}\left(\sum_{i,j\in[n]} (b_{ij} - c_{ij})x_{ij}\right) \geq \frac{1}{2n}\left(\sum_{i,j\in[n]} (b_{ij} - c_{ij})x'_{ij}\right) - \epsilon. \tag{8}$$

Proof. Consider each pair (i, j) of matched agents in x'. Sum up the "no ϵ-blocking pair" condition in (5) for all these pairs and apply the condition in (6) that surpluses are non-negative:

$$\sum_{i\in[n]} u_i + \sum_{j\in[n]} v_j \geq \sum_{(i,j):x'_{ij}=1} (u_i + v_j) \geq \sum_{(i,j):x'_{ij}=1} (b_i - c_j - 2\epsilon)\,. \tag{9}$$

Dividing by $2n$ yields the desired result. □

Another benefit of achieving an ϵ-stable outcome is that it helps the platform maintain its market share. As explained in Section 2.2, we interpret the distributions B_{ij}'s and C_{ij}'s as modeling the platform's incomplete information on agent preferences, and a rival platform may in theory obtain better information and steal users by offering better matches. However, ϵ-stability guarantees that if each agent has a fixed cost ϵ of switching platforms, then even if a rival platform were to obtain perfect information on preferences, it cannot attract away any group of users by matching them with one another more efficiently: due to (5), any pair $(i,j) \in I \times J$ that is matched by the rival platform would get at most ϵ each in surplus over what they currently get.

A potential concern with ϵ-stability is that it might be too hard to achieve, as ϵ-stability is defined based on the actual preferences b_{ij}'s and c_{ij}'s, whereas the platform initially only knows the distributions B_{ij}'s and C_{ij}'s. This concern is addressed in the next section, which shows that a platform can facilitate an ϵ-stable outcome using a variety of matchmaking strategies.

2.4. Matchmaking protocols

In order to analyze the performance of a given matchmaking strategy, we need to make assumptions on how agents behave. In the main body of the paper, we make the simplifying assumption that agents are non-strategic, which means that they truthfully reveal whatever information that is elicited from them. Incentives to deviate from truthful behavior are discussed in Appendix A.

The entire process of match formation, which accounts for both platform design and agent behavior, is referred to as a *matchmaking protocol*, where the word "protocol" is inherited from communication complexity theory and denotes a precise sequence of steps whereby communication takes place. Examples of matchmaking protocols are given below. Each protocol is a precise model of the matchmaking process under one of the four matchmaking strategies described in Section 1. As will be shown, each of the four matchmaking protocols is able to achieve an ϵ-stable outcome, and subsequent analysis will measure the efficiency of each protocol by quantifying the amount of communication it needs to achieve this goal.

Protocol 1 models how agents interact when customers are given a searchable directory of providers along with information on the surpluses required by each provider. In this protocol, each customer i contacts providers in decreasing order of estimated surplus $b_{ij} - \hat{v}_j - \hat{c}_{ij}$, where $\hat{v}_j$ represents the minimum surplus that provider j requires to obtain his service, and $\hat{c}_{ij}$ is the customer's estimate of the provider's cost for serving her. (For concreteness, the paper always uses female pronouns to refer to customers and male pronouns to refer to providers.) As customers contact providers, they obtain more accurate cost information, and update $\hat{c}_{ij}$ accordingly. Providers also update the surplus $\hat{v}_j$ they require based on supply and demand. The protocol assumes that $\hat{v}_j$ is initialized to zero, and $\hat{c}_{ij}$ is initialized to be close to the lower support of the cost distribution

C_{ij}. This can be interpreted as providers being initially pessimistic of their competitiveness in the market and customers being optimistic that they can find providers at relatively low costs.

Protocol 1 (Customers search) *Given* $\epsilon > 0$.

1. **Initialization:** *Initialize the matching* x *to be empty and the price vector* p *to be zero. Initialize the set of active searching customers as* $A = I$. *For each provider* $j \in J$, *initialize his required surplus as* $\hat{v}_j = 0$. *For each customer* $i \in I$, *initialize her cost estimate as* $\hat{c}_{ij} = \underline{c}_{ij} + \epsilon$ *for each* $j \in J$, *where* $\underline{c}_{ij}$ *is the lower support of the cost distribution* C_{ij}. *Define her estimated surplus from being served by each provider* j *as*

$$U_{ij}(\hat{v}, \hat{c}) := b_{ij} - \hat{v}_j - \hat{c}_{ij}. \tag{10}$$

2. **Customers search for providers:**
   (a) **Customer selection:** *If* $|A| = 0$, *then terminate the protocol with the transaction outcome* (x, p). *Otherwise, select an arbitrary customer* $i \in A$ *and proceed to step (b).*
   (b) **Provider interaction:** *For the selected customer* i, *if* $U_{ij}(\hat{v}, \hat{c}) < 0$ *for all* $j \in J$, *then the customer communicates that she is not interested in any provider, and the protocol removes* i *from* A *and goes back to step (a). Otherwise, customer* i *contacts a provider*

$$j^* \in \arg\max\{U_{ij}(\hat{v}, \hat{c})\}, \tag{11}$$

   and asks whether he is willing to serve her at price $\hat{v}_{j^*} + \hat{c}_{ij^*}$.
   - *If* $c_{ij^*} \leq \hat{c}_{ij^*}$, *the provider answers "yes" and becomes matched to* i*: update* $x_{ij^*} = 1$, $p_i = \hat{v}_{j^*} + \hat{c}_{ij^*}$, *and remove* i *from* A. *If* j^* *was previously matched to another customer* i', *then* i' *becomes unmatched: update* $x_{i'j^*} = 0$, $p_{i'} = 0$, *and add* i' *to* A. *Provider* j^* *increments* $\hat{v}_{j^*}$ *by* ϵ *and communicates this to the market. Go back to step (a).*
   - *If* $c_{ij^*} > \hat{c}_{ij^*}$, *the provider answers "no" and communicates his actual cost* c_{ij^*} *to the customer. The customer updates* $\hat{c}_{ij^*} = c_{ij^*}$. *Repeat step (b).*

From a technical perspective, Protocol 1 is a generalization of the Auction mechanism of Bertsekas (1988) and Demange et al. (1986) to a setting with incomplete information: if the cost estimates $\hat{c}_{ij}$'s were initialized to the actual costs c_{ij}'s, then Protocol 2 becomes equivalent to the mechanism analyzed in Bertsekas (1988) and Demange et al. (1986), with provider surpluses being bid up in a multi-item ascending auction. The initialization of the $\hat{v}_j$'s and $\hat{c}_{ij}$'s in Step 1 are justified by the following proposition.

PROPOSITION 2 **(Effectiveness of Protocols 1)**. *Given any* $\epsilon > 0$, *Protocol 1 always terminates with an* ϵ*-stable outcome.*

Proof of Proposition 2. To see that Protocol 1 always terminates in finitely many iterations, observe that in Step 2(b), whenever a provider j answers "yes" to a customer, $\hat{v}_j$ increments by ϵ. However, $\hat{v}_j$ will never increment more than $\lceil 1/\epsilon \rceil$ times since $\hat{v}_j$ will not increment when it exceeds 1, as otherwise $U_{ij}(\hat{v}, \hat{c}) < 0$ for all $i \in I$ by the bounded support assumption in (7) and j would not be contacted by any customer. Moreover, a provider j will answer "no" to a customer i at most once, as by the second time i contacts j, $\hat{c}_{ij}$ would be equal to c_{ij}.

It remains to show that the final outcome is ϵ-stable. Observe first that the initialization of $\hat{c}_{ij}$ implies that throughout the protocol,

$$c_{ij} \geq \hat{c}_{ij} - \epsilon \qquad \text{for all } (i,j) \in I \times J. \tag{12}$$

Let (x, u, v) be the outcome corresponding to the final transaction outcome (x, p). Observe that the definition of the transaction price in Step 2(b) is such that whenever a pair is matched, both agents obtain a non-negative surplus, so $u \geq 0$ and $v \geq 0$. Moreover,

$$u_i \geq b_{ij} - \hat{v}_j - \hat{c}_{ij} \qquad \text{for all } (i,j) \in I \times J, \tag{13}$$

$$v_j \in [\hat{v}_j - \epsilon, \hat{v}_j] \qquad \text{for all } j \in J. \tag{14}$$

To see (13), observe that if i is unmatched, then the left hand side is zero and the right hand side is negative; if i is matched to j^*, then (13) follows from (11). Moreover, (14) follows from the definition of the transaction price in Step 2(b) and from (12), as well as from the fact that $\hat{v}_j$ would have incremented by ϵ after the last time that provider j was tentatively matched. Combining (12), (13) and (14), we have $u_i \geq b_{ij} - v_j - c_{ij} - 2\epsilon$ for all $(i,j) \in I \times J$, so (x, u, v) is ϵ-stable. □

Protocol 2 (Providers search) *This protocol is the mirror image of Protocol 1 with the roles of customers and providers interchanged. A full statement is given in Appendix B.*

By symmetry, Protocol 2 also guarantees an ϵ-stable outcome upon termination.

Protocol 3 is a generalization of Protocol 1 in which the cost estimate $\hat{c}_{ij}$ is not necessarily initialized to be close to the lower support $\underline{c}_{ij}$, but may be based on another quantile z_{ij} of the commonly known cost distribution C_{ij}. The exact choice of z_{ij} is a parameter of the protocol that will be optimized later. When $z_{ij} = \underline{c}_{ij}$, Protocol 3 is identical to Protocol 1; when z_{ij} takes on higher values, Protocol 3 involves an additional step in which providers initiate contact with customers for whom $c_{ij} < z_{ij}$. One can interpret this as providers browsing through a searchable directory of customers and offering special discounts to their preferred customers, which in the model corresponds to those whom they can serve at especially low costs. As will be shown in Section 3.2, the advantage of this design is that if certain providers are very picky about their preferred customers, it would be more efficient for these providers to initiate contact first.

Protocol 3 (Both sides search) *Given $\epsilon > 0$, as well as an $I \times J$ matrix z that depends only on the commonly known market information (I, J, B, C).*

1. **Initialization:** *Exactly as in Step 1 of Protocol 1, except that for each customer $i \in I$, her cost estimate is initialized as $\hat{c}_{ij} = z_{ij} + \epsilon$ for each $j \in J$.*
2. **Providers search for customers:** *Each provider $j \in J$ initiates contact with each customer $i \in I$ such that $c_{ij} < z_{ij}$, and communicates c_{ij} to customer i, who updates $\hat{c}_{ij} = c_{ij}$.*
3. **Customers search for providers:** *Exactly as in Step 2 of Protocol 1.*

Proposition 3 (**Effectiveness of Protocol 3**). *Given any $\epsilon > 0$ and any choice of z, Protocol 3 always terminates with an ϵ-stable outcome.*

Proof of Proposition 3. Observe that the number of customer-provider pairs that interact in Step 2 is bounded above by $|I| \times |J|$. Moreover, at the end of Step 2, we have $c_{ij} \geq \hat{c}_{ij} - \epsilon$ for all $(i, j) \in I \times J$, as in condition (12) in the proof of Proposition 2. The same argument as in the proof of Proposition 2 can now be applied to show that Step 3 always terminates in finitely many iterations with an ϵ-stable outcome. □

Protocol 4 represents a version of centralized matching, in which the platform uses agents' responses to questionnaires to recommend a limited number of suitable providers to each customer along with the transaction price. To account for idiosyncratic preferences unknown to the platform, agents have the option of rejecting non-ideal recommendations. The final matching is determined so as to maximize the number of matches. This can either be interpreted as throughput optimization on the part of the platform, or as an approximation to a decentralized process of match formation among the recommended pairs.

Protocol 4 (Centralized matching) *Given $\epsilon > 0$ and the following parameters, each of which can only depend on the commonly known market information (I, J, B, C): a price vector $p \in \mathbb{R}^n$, and a set $S_i \subseteq J$ of match recommendations for each customer $i \in I$. Define $\bar{b}_{ij}$ to be the upper support of the benefit distribution B_{ij} and $\underline{c}_{ij}$ to be the lower support of the cost distribution C_{ij}.*

1. **Preference elicitation:** *For each pair (i, j) where $j \in S_i$, ask customer i whether $b_{ij} \geq \bar{b}_{ij} - \epsilon$ and ask provider j whether $c_{ij} \leq \underline{c}_{ij} + \epsilon$. If both agents answer "yes," then the pair is said to be an accepted recommended pair.*
2. **Match determination:** *Compute a maximum cardinality matching x among the accepted recommended pairs. Terminate the protocol with the transaction outcome (x, p).*

Compared to the other protocols, Protocol 4 has the following qualitative advantages:

1. All communication can be done simultaneously, so agents do not have to wait for certain potential partners to respond before knowing what to do next.

2. Agents can limit their attention to the set of recommended partners, and do not have to know their exact preferences over the entire directory of possible partners.

Unlike Protocols 1-3, Protocol 4 does not necessarily yield ϵ-stable outcomes for all markets. Section 3.3 derives regularity conditions under which it achieves an ϵ-stable outcome with high probability, and shows how the price vector and match recommendations can be selected.

3. Suitable market characteristics

This section derives the market characteristics under which each of the four matchmaking protocols achieves an ϵ-stable outcome using low amounts of communication. For now, the "amount of communication" is measured by the number of interactions, where an *interaction* is defined in Protocols 1, 2 and 3 as each instance in which one agent initiates contact with another, and is defined in Protocol 4 as each match recommendation. (Section 4 analyzes a complementary metric, which is the number of bits of information communicated.) The main takeaway is that the driver of matchmaking efficiency is whether preferences are easy to describe, as modeled by the following parameters of the preference distributions.

Definition 1 (Preference density). For any market (I, J, B, C), define the preference density d^I of customers as the largest value that satisfies the following condition:

$$\mathbb{P}(b_{ij} \geq \bar{b}_{ij} - \epsilon) \geq d^I \qquad \text{for all } (i,j) \in I \times J. \tag{15}$$

Similarly define the preference density d^J of providers as the largest value such that

$$\mathbb{P}(c_{ij} \leq \underline{c}_{ij} + \epsilon) \geq d^J \qquad \text{for all } (i,j) \in I \times J. \tag{16}$$

Here, $\bar{b}_{ij}$ is the upper support of the distribution B_{ij} and $\underline{c}_{ij}$ is the lower support of C_{ij}. Note that the $\bar{b}_{ij}$'s and $\underline{c}_{ij}$'s are allowed to be fully heterogeneous across (i,j) pairs.

We interpret a high value of d^I as customers having easy-to-describe preferences. This is because the upper support $\bar{b}_{ij}$ is based on agents' initial descriptions of who they are and what they want. (See Section 2.2 for a more detailed explanation of the relationship between the distribution B_{ij} and agents' descriptions.) If agents' descriptions already capture the most relevant information about customer benefits, then $\bar{b}_{ij}$ is a good estimate of the true benefit b_{ij}. Here, being a "good estimate" does not mean that $\bar{b}_{ij}$ is always close to b_{ij}, but only that $\bar{b}_{ij}$ is close to b_{ij} with some probability d^I, which represents the degree that customer benefts are well-approximated by their upper supports. The reason that we use the upper support $\bar{b}_{ij}$ of benefit and the lower support $\underline{c}_{ij}$ of cost is that for the purpose of identifying an ϵ-stable outcome, agent pairs (i,j) with high benefit and low cost are the most relevant, as matched pairs tend to have a high value of $b_{ij} - c_{ij}$. From a technical perspective, having an upper bound of benefits and a lower bound of costs allows us to conclude that certain (i, j) pairs are not ϵ-blocking, without needing to know the exact values of b_{ij} and c_{ij}. This will be important in the analysis later.

Similarly, a high value of d^J corresponds to providers having easy-to-describe preferences. Table 3 restates Table 2 from Section 1.1 by expressing the characteristics of suitable markets in terms of d^I and d^J. The following subsections derive formal results that support this table.

Table 3 Market characteristics under which each matchmaking protocol presented in Section 2.4 achieves an ϵ-stable outcome with high probability using low amounts of communication.

Matchmaking protocol	Characteristics of suitable markets	Interpretation	Analysis
1. Customers search	d^J is high.	Provider preferences are easy to describe.	Sec. 3.1.
2. Providers search	d^I is high.	Customer preferences are easy to describe.	By symmetry.
3. Both sides search	Neither d^I or d^J is high.	Neither side's preferences are easy to describe.	Sec. 3.2.
4. Centralized matching	Both d^I and d^J are high.	Both sides' preferences are easy to describe.	Sec. 3.3.

3.1. Protocol 1 (customers search) is suitable when d^J is high

Recall from Proposition 2 that Protocol 1 always yields an ϵ-stable outcome upon termination, so the main question is how quickly it terminates. The following result implies that the total number of interactions in this protocol is inversely related to d^J. In other words, when d^J is high (i.e., provider preferences are easy to describe), then "customers search" is relatively efficient; when d^J is very low, then "customers search" may not be suitable.

PROPOSITION 4 **(Performance of Protocol 1)**. *In Protocol 1, with probability at least* $1 - e^{-n/2}$*, the number of interactions per customer (i.e., the total number of interactions divided by* n*) does not exceed*

$$2\left\lceil \frac{1}{\epsilon} \right\rceil \frac{1}{d^J}. \tag{17}$$

Proof of Proposition 4. An interaction in Protocol 1 corresponds to one iteration of Step 2(b). Observe that if $\epsilon \geq 1$, then by the initialization of $\hat{c}_{ij}$ in Protocol 1, no provider will answer "no" to a customer in Step 2(b), and the number of interactions per customer is at most 1. For the remainder of the proof, assume that $\epsilon < 1$.

As discussed in the proof of Proposition 2, the number of increments of $\hat{v}_j$ cannot exceed $\lceil 1/\epsilon \rceil$, as $\hat{v}_j$ is initialized to 0, does not increment when it exceeds 1, and each increment increases it by ϵ. Moreover, whenever a customer i contacts a provider j in Step 2(b), this interaction leads to an increment in $\hat{v}_j$ with probability at least d^J: if i has not contacted j before, then this holds by (16) and by the assumption that c_{ij} is independent from everything else; if i has contacted j before, then i already has already updated $\hat{c}_{ij}$ to be equal to c_{ij}, and $\hat{v}_j$ increments with probability 1. Therefore, the total number of interactions is stochastically dominated by the sum of $n\lceil 1/\epsilon\rceil$ i.i.d. geometric random variables, each with mean $1/d^J$. The chance that this exceeds $r := 2n\lceil 1/\epsilon\rceil/d^J$ is equal to the chance that a $Binomial(r, d^J)$ random variable is less than $rd^J/2$, which by a standard Chernoff bound is bounded above by

$$\exp\left(-\frac{rd^J}{8}\right) = \exp\left(-\frac{\lceil 1/\epsilon\rceil n}{4}\right) \leq \exp\left(-\frac{n}{2}\right). \tag{18}$$

The last inequality follows from $\epsilon < 1$. □

In interpreting the bound in (17), the most important term is the $1/d^J$, as this is an upper bound for the expected number of providers a customer needs to contact in order to receive a favorable response. In the proof, the factor of 2 is used to obtain the high probability guarantee, and can be reduced to any constant larger than 1 if one allows for a weaker probabilistic guarantee. The $\lceil 1/\epsilon\rceil$ arises from the bidding process for provider surpluses, and is an artifact of the assumption that the $\hat{v}_j$'s are initialized to zero. In reality, one would expect providers to have a better sense of how much surplus they can expect when they join the market, and if the $\hat{v}_j$'s were initialized at their final values, then the $\lceil 1/\epsilon\rceil$ term would disappear.

Since Protocol 2 (provider search) is the mirror analog of Protocol 1, the above analysis implies that Protocol 2 is suitable when d^I is high (i.e., customer preferences are easy to describe).

3.2. Protocol 3 (both sides search) is suitable when neither d^I or d^J is high

The following proposition illustrates that when d^I are d^J are very low, Protocols 1 and 2 may both perform badly when the market size is large, whereas Protocol 3 may perform very well.

PROPOSITION 5 **(Illustration of the benefit of allowing both sides to search)**. *Consider the following market: for all* (i, j),

$$b_{ij} = \begin{cases} 1 & \text{with probability } d^I := 1/(2n), \\ 0.6 & \text{otherwise;} \end{cases} \tag{19}$$

$$c_{ij} = \begin{cases} 0 & \text{with probability } d^J := 1/(2n), \\ 0.4 & \text{otherwise.} \end{cases} \tag{20}$$

In this market, for any $\epsilon < 0.4$, *the expected number of interactions per customer is at least* $0.47n$ *under Protocol 1 or 2, but is equal to* 1.5 *under Protocol 3 if the parameter* $z_{ij} \equiv 0.4$.

Proof of Proposition 5. Consider applying Protocol 1 in the above market. We say that a customer has no preferred provider if $b_{ij} = 0.6$ for all $j \in J$. Such a customer would contact every provider in order to find one who can serve with zero cost. The expected number of interactions before she either finds a zero-cost provider or has obtained cost information from every provider is $\mathbb{E}[\min(X,n)]$, where X is a geometrically distributed random variable with mean $2n$. By a straightforward calculation, $\mathbb{E}[\min(X,n)] = 2n(1-q)$, where $q := [1-1/(2n)]^n$. The expected number of customers with no preferred providers is qn. Hence, in Protocol 1, the expected number of interactions per customer is at least $2(1-q)qn \geq 2(1-e^{-0.5})e^{-0.5}n > 0.47n$. The first inequality follows from the fact that $f(y) = (1-y)y$ is minimized if y is furthest from 0.5, and $0.5 \leq q \leq e^{-0.5}$.

By symmetry of the given benefit and cost distributions, the same bound applies for Protocol 2.

In Protocol 3, an interaction corresponds to each instance in which a provider contacts a customer in Step 2, or a customer contacts a provider in Step 3. If $z_{ij} = 0.4$ for all (i,j), then the expected number of interactions in Step 2 is $0.5n$, as only zero-cost providers reach out. At the end of Step 2, every customer can be certain of the cost of every provider: customer i can safely assume that provider j's cost of serving her is 0.4, unless she was contacted by the provider, in which case the cost is 0. Since provider costs are no longer uncertain, every customer would contact exactly one provider in provider 3. The expected number of interactions per customer is hence $0.5n/n + 1 = 1.5$, as desired. □

For best performance, Protocol 3 requires a good choice of the parameters z_{ij}'s, which can be interpreted as customers having access to a good initial estimate of provider costs based on the commonly known information. The following proposition analyzes the performance of Protocol 3 under an arbitrary choice of z.

Proposition 6 **(Performance of Protocol 3)**. *Given any market and any $\epsilon > 0$, the following expression upper-bounds the expected number of interactions in Protocol 3 under a given choice of z.*

$$N(z) := \left(\sum_{i \in I, j \in J} \mathbb{P}(c_{ij} < z_{ij}) \right) + \left\lceil \frac{1}{\epsilon} \right\rceil \left(\sum_{j \in J} \frac{1}{\min_{i \in I}\{\mathbb{P}(c_{ij} \leq z_{ij} + \epsilon | c_{ij} \geq z_{ij})\}} \right). \tag{21}$$

Proof of Proposition 6. The first term in (21) is the expected number of interactions in Step 2. The second term is an upper-bound of the expected number of interactions in Step 3: As in the proof of Proposition 4, for a given provider j, if $\mathbb{P}(c_{ij} \leq z_{ij} + \epsilon | c_{ij} \geq z_{ij}) \geq d_j$, then the provider is contacted by at most $1/d_j$ customers in expectation before incrementing his required surplus $\hat{v}_j$ once, which can increment at most $\lceil 1/\epsilon \rceil$ times in total. □

The above implies that a good initial cost estimate z_{ij} should have the following two properties:

1. It is "**optimistic**," in the sense that $\mathbb{P}(c_{ij} < z_{ij})$ is small. This ensures that a provider does not need to reach out to too many customers in Step 2. In Proposition 5, $\mathbb{P}(c_{ij} < z_{ij}) = 1/(2n)$.
2. It is "**attainable**," in the sense that $\mathbb{P}(c_{ij} \leq z_{ij} + \epsilon | c_{ij} \geq z_{ij})$ is large. This ensures that a customer does not need to reach out to too many providers in Step 3 before finding one who agrees to provide service at the desired price. In Proposition 5, $\mathbb{P}(c_{ij} \leq z_{ij} + \epsilon | c_{ij} \geq z_{ij}) = 1$.

A good choice of z for a given market can be found by minimizing $N(z)$. For the market in Proposition 5, minimizing $N(z)$ yields the optimal choice of $z_{ij} \equiv 0.4$. The following choice of z approximately minimizes $N(z)$ for worst-case preference distributions:

$$z^*_{ij} := \inf\{y \in [\underline{c}_{ij}, 1] : \mathbb{P}(c_{ij} \in [y, y+\epsilon]) \geq 1/\sqrt{n}\}. \tag{22}$$

Whenever $\epsilon \geq 1/\sqrt{n}$, the infimum is guaranteed to exist by the Pigeonhole Principle because $\mathbb{P}(c_{ij} \in [0,1]) = 1$. The following proposition shows that the above choice of z induces a performance guarantee that scales according to $\sqrt{n}$, which will be shown in Section 4 to be asymptotically the best possible if both d^I and d^J are allowed to be very low.

Proposition 7 **(Distribution-independent guarantee of Protocol 3)**. *Given any market and any $\epsilon \geq 1/\sqrt{n}$, consider Protocol 3 (both sides search) with parameter z^* as in (22). With probability at least $1 - 2e^{-n/2}$, the number of interactions per customer does not exceed*

$$4\left\lceil \frac{1}{\epsilon} \right\rceil \sqrt{n}. \tag{23}$$

Proof of Proposition 7. If $\epsilon \geq 1$, then $z^*_{ij} \equiv \underline{c}_{ij}$ and the number of interactions per customer in Protocol 3 is at most 1. In the remainder of the proof, assume $\epsilon < 1$, so $\lceil 1/\epsilon \rceil \geq 2$.

By the construction of z^*,

$$\mathbb{P}(c_{ij} < z^*_{ij}) < z^*_{ij} \frac{1}{\epsilon\sqrt{n}} \leq \frac{1}{\epsilon\sqrt{n}}. \tag{24}$$

Define $r_1 := ((1/\epsilon) + 0.5)n^{1.5}$. The chance that the total number of interactions in Step 2 exceeds r_1 is upper-bounded by the chance that a $Binomial(n^2, (1/\epsilon)n^{-0.5})$ random variable exceeds r_1, which by Hoeffding's inequality is no more than $\exp(-2(0.5n^{1.5})^2/n^2) = \exp(-n/2)$.

Moreover, the construction of z^* implies that $\mathbb{P}(c_{ij} \leq z^*_{ij} + \epsilon | c_{ij} \geq z^*_{ij}) \leq \mathbb{P}(c_{ij} \in [z^*_{ij}, z^*_{ij} + \epsilon]) \geq 1/\sqrt{n}$. Therefore, the total number of interactions in Step 3 of Protocol 3 is stochastically dominated by the sum of $n\lceil 1/\epsilon \rceil$ i.i.d. geometric random variables, each with mean $\sqrt{n}$. The chance that this exceeds $r_2 := 2\lceil 1/\epsilon \rceil n^{1.5}$ is upper-bounded by the chance that a $Binomial(r_2, 1/\sqrt{n})$ random variable being less than $r_2/(2\sqrt{n})$, which by a standard Chernoff bound is upper-bounded by $\exp(-r_2/(8\sqrt{n})) \leq \exp(-n/2)$.

Combining the above, the probability that the number of interactions per customer exceeds $4\lceil 1/\epsilon \rceil \sqrt{n} > (r_1 + r_2)/n$ is upper-bounded by $2\exp(-n/2)$, as desired. □

3.3. Protocol 4 (centralized matching) is suitable when both d^I and d^J are high

Intuitively speaking, when the preferences of both sides are easy to describe, then the platform can suggest a good matching using only agents' initial descriptions, and it is unnecessary for either side to search through the set of all potential partners. If there are idiosyncratic preferences that are not captured by the initial descriptions, then agents need to be allowed to choose among several recommended options, but the number of options needed is small. These ideas are formalized by the analysis in this section.

For clarity of exposition, consider first the simpler case in which both sides of the market are ex-ante homogeneous: i.e., the distributions B_{ij} and C_{ij} do not depend on (i,j). Let $\overline{b}$ denote the upper support of the common benefit distribution and $\underline{c}$ the lower support of the common cost distribution. Given $\epsilon > 0$, assume that $\overline{b} - \underline{c} \geq 2\epsilon$, as otherwise the empty matching is ϵ-stable. The following proposition shows that in this case, whenever the product $d^I d^J$ is sufficiently large, centralized matching with randomly chosen recommendation sets yields an ϵ-stable matching with high probability, and the expected number of recommendations per agent is small.

PROPOSITION 8. *Choose any $y \in [\underline{c}+\epsilon, \overline{b}-\epsilon]$. Consider Protocol 4 with price $p_i \equiv y$, and each recommendation set S_i is randomly generated so that each provider $j \in J$ is contained in S_i with probability $a := \min(1, (3\ln n)/(nd^I d^J))$, independently across (i,j) pairs. If $d^I d^J \geq \min(1, (3\ln n)/n)$, then this protocol yields an ϵ-stable outcome with probability at least $1 - 3/n^2$.*

Note that the expected number of recommendations per customer is

$$\mathbb{E}[|S_i|] = an \leq \frac{3\ln n}{d^I d^J}, \tag{25}$$

which is small if $d^I d^J$ is relatively large.

Proof of Proposition 8. By construction of Protocol 4, a recommended pair (i,j) is accepted if both $b_{ij} \geq \overline{b} - \epsilon$ and $\underline{c}_{ij} \leq \underline{c} + \epsilon$, which guarantees a total surplus of $b_{ij} + c_{ij} \geq \overline{b} - \underline{c} - 2\epsilon$. Each recommended pair is accepted with probability at least $d^I d^J$. The final outcome is ϵ-stable whenever there exists a perfect matching within the subgraph of accepted recommended pairs. This is because when such a perfect matching exists, $u_i \geq \overline{b} - \epsilon - y \geq 0$ for all $i \in I$ and $v_j \geq y - \underline{c} - \epsilon \geq 0$ for all $j \in J$. Moreover, $u_i + v_j + 2\epsilon \geq \overline{b} - \underline{c} \geq b_{ij} - c_{ij}$ for all (i,j). Now, if $a = 1$ and $d^I d^J = 1$, then every pair (i,j) is an accepted recommended pair and there is trivially a perfect matching. On the other hand, if $a < 1$ or $d^I d^J < 1$, then $ad^I d^J = (3\ln n)/n$, and the probability that each (i,j) is an accepted recommended pair is at least $r := (3\ln n)/n$. The following lemma implies that a perfect matching exists with probability at least $1 - 2.1ne^{-rn} = 1 - 2.1/n^2 > 1 - 3/n^2$. The proof of the lemma is in Appendix D.1 and is related to the analysis of Erdös and Rényi (1964) regarding the probability that a perfect matching exists in a random bipartite graph. □

LEMMA 1. *Given any $n \geq 1$ and $r \geq (3\ln n)/n$, consider a random bipartite graph with n nodes on each side, such that there exists an edge between each pair of nodes (i,j) with probability at least $\min(1, r)$, independently across (i,j) pairs. The final graph contains a perfect matching with probability at least $1 - 2.1ne^{-rn}$.*

Returning to the general case in which agents are not ex-ante homogeneous, so that B_{ij} and C_{ij} may depend on (i,j). The following definition allows us to parameterize the degree of homogeneity on one side of the market using a certain parameter l. The definition can be applied to either side of the market, but for concreteness we focus on the case of homogeneity among the providers.

Definition 2 (Relatively homogeneous). Given a market (I,J,B,C), the providers are said to be *relatively homogeneous* with minimum segment size l if the set J can be partitioned into segments of size at least l: i.e., $J = \dot{\bigcup}_{k\in K} J_k$, where K is the set of segments, J_k is the set of providers in segment k, and $|J_k| \geq l$ for all $k \in K$. The segmentation is such that the supports $\overline{b}_{ij}$ and $\underline{c}_{ij}$ depend only on the customer i and the segment k that contains the provider j, but not on the exact identity of the provider $j \in J_k$.

Note that every market satisfies the above definition if the minimum segment size l is 1. So the above definition is only restrictive when l is large. We will show that whenever the product $d^I d^J$ is weakly larger than the threshold $\min(1,(3\ln n)/l)$, then we can construct a price vector and recommendation sets such that Protocol 4 yields an ϵ-stable matching with high probability, and the expected number of recommendations per agent is at most $(3\ln n)/(d^I d^J)$, which is inversely related to $d^I d^J$. This implies that in any market, if both d^I and d^J are sufficiently high, then centralized matching is suitable. When $l \leq 3\ln n$, $d^I d^J \geq \min(1,(3\ln n)/l)$ implies that $d^I = d^J = 1$. With a larger value of l, the requirement on the minimum value of $d^I d^J$ is correspondingly reduced.

The construction of the price vector and recommendation sets is based on computing a stable transaction outcome $(\hat{x}, p)$ under the following optimistic estimates of benefits and costs:

$$\hat{b}_{ij} := \overline{b}_{ij} - \epsilon, \quad \hat{c}_{ij} := \underline{c}_{ij} + \epsilon \qquad \text{for each } (i,j) \in I \times J. \tag{26}$$

More concretely, let $\hat{x}$ be an optimal basic solution to the following linear program (LP):

$$\text{Maximize:} \quad \sum_{i\in I, j\in J} (\hat{b}_{ij} - \hat{c}_{ij})\hat{x}_{ij} \tag{27}$$

$$\text{s.t.:} \quad \sum_{j\in J} \hat{x}_{ij} \leq 1 \qquad \text{for each } i \in I. \tag{28}$$

$$\sum_{i\in I} \hat{x}_{ij} \leq 1 \qquad \text{for each } j \in J. \tag{29}$$

$$\hat{x} \geq 0 \tag{30}$$

Let $\hat{u}_i$ be an optimal dual variable corresponding to (28). Define the price vector p such that

$$p_i := \begin{cases} \hat{b}_{ij} - \hat{u}_i & \text{if } i \text{ is matched to } j \text{ in } \hat{x}; \\ 0 & \text{if } i \text{ is unmatched in } \hat{x}. \end{cases} \tag{31}$$

Define the recommendation set S_i as follows: if customer i is matched in $\hat{x}$ to a provider of segment $k \in K$, then S_i is a random subset of J_k where each provider $j \in J_k$ is independently included with probability $a_k := \min(1, (3\ln n)/(|J_k| d^I d^J))$; if i is unmatched in $\hat{x}$, then $S_i := \emptyset$. Note that the expected number of match recommendations for each customer satisfies the upper bound:

$$\mathbb{E}[|S_i|] \leq \max_{k\in K}\{a_k |J_k|\} \leq \frac{3\ln n}{d^I d^J}. \tag{32}$$

PROPOSITION 9 **(Effectiveness and performance of Protocol 4)**. *Given $\epsilon > 0$ and a market in which the providers are relatively homogeneous with minimum segment size l. Consider Protocol 4 with the above choices of p and S_i's. If $d^I d^J \geq \min(1, (3\ln n)/l)$, then this protocol yields an ϵ-stable outcome with probability at least $1 - 3/n^2$.*

The proof of the above proposition is given in Appendix D.2. The proof is based on showing that whenever $d^I d^J$ is sufficiently large, with high probability the protocol matches every customer i who is matched in $\hat{x}$ to a recommended provider j for whom the optimistic estimates $\hat{b}_{ij}$ and $\hat{c}_{ij}$ are accurate, meaning that $b_{ij} \geq \hat{b}_{ij}$ and $c_{ij} \leq \hat{c}_{ij}$. When this occurs, every agent is at least as happy as under $(\hat{x}, p)$, which implies ϵ-stability by the construction of $(\hat{x}, p)$ and by (26).

The above analysis quantifies the number of match recommendations needed to achieve an ϵ-stable outcome. A follow up paper, Shi (2022), studies a complementary question: if the platform has a fixed budget on the number of recommendations it is allowed to provide to each customer, how can the platform optimize its match recommendations to maximize the total market surplus.

4. Near-optimality based on communication complexity theory

The analyses in Section 3 identify the characteristics of suitable markets for each of the four matchmaking protocols defined in Section 2.4. This section addresses a complementary question: for markets with the given characteristics, what is the *best possible* way of matchmaking? Here, "best possible" means that we are no longer restricting ourselves to the four protocols, but are interested in establishing performance bounds that hold under *any* system of matchmaking. As discussed in Section 1.2, the benefit of the communication complexity approach is its ability to establish such bounds using information theory, without needing to make strong assumptions on the process of matchmaking. We will show that for suitable markets, the four matchmaking protocols are in fact near-optimal among *all* possible systems of matchmaking.

In order to compare the performance of disparate platform designs, we cannot use a design-specific metric such as the number of interactions, as an interaction under "customers search" has a different meaning from an interaction under "centralized matching." A trivial way of minimizing interactions is to have each agent interact once with the platform, where each interaction entails communicating all of the agent's preferences for all potential partners. Such a system requires only one interaction per agent by construction, but the amount of information that needs to be communicated in each interaction is enormous. A universal metric for quantifying the amount of information that needs to be communicated is the number of binary bits needed to encode that information. This is the metric used in this section to quantify the amount of communication in different matchmaking strategies.

More precisely, this section addresses the following question: for a given class of markets, in the worst case, how many bits of communication per agent is needed to obtain an ϵ-stable outcome with high probability? A lower bound to the minimum number of bits can be derived using information theory by analyzing a worst-case market within the given class; an upper bound can be established by constructing a protocol that guarantees good performance for every market in the given class. If we identify a protocol whose performance guarantee is close to the lower bound, then that protocol is said to provide near-optimal performance guarantees.

The findings are summarized in Table 4. The first column describes the space of matchmaking protocols being considered, which are explained in Section 4.1. The second column describes the class of markets being considered, and each class of market corresponds to whether d^I or d^J is required to be high, which is consistent with the classification of markets in Section 3. The next two columns state lower and upper bounds on the minimum number of bits of communication per agent needed in the worst case to obtain an ϵ-stable outcome with high probability. The last column of Table 4 describes a protocol whose performance guarantee matches the upper bound.

Table 4 The number of bits of communication per agent needed in the worst case by any matchmaking protocol that obtains an ϵ-stable outcome with high probability in every market within a given class. All bounds assume a constant $\epsilon > 0$, which is incorporated into the constants in the big-O notation.

Space of matchmaking protocols (See Sec. 4.1)	Class of markets	Lower bound (See Sec. 4.3)	Upper bound (See Sec. 4.2)	Protocol achieving the upper bound
Unrestricted	$d^J \geq \Omega(1)$	$\Omega(\log n)$	$O(\log n)$	Protocol 1 (customers search)
	$d^J \geq 1/\sqrt{n}$	$\Omega(1/d^J)$	$O((\log n)/d^J)$	
	$d^I \geq \Omega(1)$	$\Omega(\log n)$	$O(\log n)$	Protocol 2 (providers search)
	$d^I \geq 1/\sqrt{n}$	$\Omega(1/d^I)$	$O((\log n)/d^I)$	
	All markets	$\Omega(\sqrt{n})$	$O(\sqrt{n}\log n)$	Protocol 3 (both sides search)
One-round protocols	$d^I d^J \geq \Omega(1)$, and providers are relatively homogeneous.	$\Omega(\log n)$	$O(\log n)$	Protocol 4 (centralized matching)

The main point of the table is that in each row, the upper and lower bounds differ by either a constant or a logarithmic factor, which shows that Protocols 1-4 provide near-optimal performance guarantees for their respective class of suitable markets. The findings corroborate Table 2 from the introduction. In the first two rows of Table 4, Protocol 1 provides near-optimal performance guarantees, and the corresponding classes of markets are those in which d^J is sufficiently high, meaning that provider proferences are easy to describe. The next two rows give the analogous results for Protocol 2 when the roles of customers and providers are interchanged. The fifth row shows that when we allow for arbitrary markets, in which case both d^I and d^J might be very low, there are markets in which order $\sqrt{n}$ bits of communication per agent is necessary for obtaining an ϵ-stable outcome with high probability, which means that the $O(\sqrt{n}\log n)$ guarantee provided by Protocol 3 is near-optimal. The last row shows that when both d^I and d^J are high, centralized matching yields near-optimal performance.

The results in Table 4 are explained in the following subsections. Section 4.1 formally defines a communication protocol, and uses it to define the space of all matchmaking protocols. Section 4.2 derives the upper bounds by counting the number of bits of communication used by the four protocols. Section 4.3 explains the derivation of the lower bounds, which establish the near-optimality of the four protocols for their respective class of suitable markets.

4.1. Space of matchmaking protocols

The following concepts allow us to rigorously define the space of all possible matchmaking protocols and compare the amount of communication across disparate platform designs.

Definition 3 (Communication protocol and communication cost). A *communication protocol* with m agents is defined as follows. Each agent $i \in [m] := \{1, 2, \cdots, m\}$ is endowed with her privately known input y_i, where the vector of everyone's inputs $y = (y_1, y_2, \cdots, y_m)$ comes from a certain set Y of possible input vectors. Without knowing y, the protocol begins by designating a certain agent to send the first message. Whenever agent $i \in [m]$ is designated to send the next message, the message sent is a deterministic function of the agent's private information and the history of messages, which is defined as the sequence of all messages sent since the beginning of the protocol by any agent. Each message is a sequence of binary bits, and the length of a message is defined as the number of bits in the sequence. (For example, a message may be "110010," which involves 6 binary bits, so the length is 6.) After an agent sends a message, it is appended to the history of messages. At that point, as a deterministic function of the history of messages, the protocol either terminates with an output or designates another agent to send the next message. Given the input vector y, the *length of the protocol* is defined to be the total number of bits in all the messages that are sent before the protocol terminates. The *per-agent communication cost* of the protocol is defined to be the maximum length of the protocol divided by the number of agents m. (Here, the maximum is taken over all possible input vectors $y \in Y$.)

In the communication complexity literature, the above is referred to as the *blackboard model of communication*, as one can think of the agents as writing their messages to a common blackboard. This model is very general, as almost all conceivable systems of multi-agent communication can be simulated within this framework. For example, an alternative model of communication is one in which each agent can only communicate to one receiver, whose identity needs to be communicated to the protocol before the message is sent. Let us refer to this as the *pairwise model of communication*. This is a strictly more restrictive model because given any communication protocol formulated in the pairwise model, one can replicate it in the blackboard model as follows: instead of having agent i sending a message s to agent j, have agent i send the message (j,s), where the first $\lceil \log_2 n \rceil$ bits are used to encode the identity of the receiver j and the following bits contain the message s. The amount of communication does not increase from this modification, as even in the pairwise model, agent i needs to communicate $\lceil \log_2 n \rceil$ bits to the platform to specify the identity of the receiver. The only difference is that under the blackboard model, agents other than j also receive the message, but having extra communication is not a problem as agents are non-strategic. Hence, if we prove a lower bound that any communication protocol in the blackboard model requires B bits of communication to achieve a certain outcome, the same lower bound would hold in the pairwise model. However, the reverse does not hold, as the pairwise model cannot simulate a protocol in the blackboard model without greatly increasing the number of messages. For the purpose of proving lower bounds, it is desirable to establish the bounds under the most general model of communication possible, as that makes the result the strongest. This is why we define a communication protocol using the blackboard model in Definition 3.

Having defined a communication protocol, we are ready to define a matchmaking protocol.

Definition 4 (Matchmaking protocol). Given a market (I,J,B,C), a *matchmaking protocol* is a communication protocol in which the set of agents is $I \cup J$; the input vector is $(y_l : k \in I \cup J)$, where $y_i = (b_{ij} : j \in J)$ for every customer $i \in I$ and $y_j = (c_{ij} : i \in I)$ for every provider $j \in J$; the support Y of the input vector y is induced by the support of the preference distributions B and C; the output of the protocol is a transaction outcome (x,p) where $x \in \{0,1\}^{I \times J}$ and $p \in \mathbb{R}^I$.

One technical issue in defining a matchmaking protocol based on Definition 4 is that since it builds on the definition of a communication protocol, all messages need to be encoded as binary bits, as required in Definition 3. However, this is not a restrictive requirements, as it is straightforward to modify the statements of the four protocols in Section 2.4 to satisfy this. The only non-trivial change is that in Protocols 1 and 3, instead of having provider j communicate the exact value of c_{ij} to customer i, the provider would communicate the integer $\Delta_{ij} = \lceil (c_{ij} - \hat{c}_{ij})/\epsilon \rceil$, which can be encoded using $\lceil \log_2(1/\epsilon) \rceil$ bits, and the customer would update her cost estimate to $\hat{c}_{ij} = \hat{c}_{ij} + \epsilon \Delta_{ij}$. This is similar to rounding costs up to the nearest integer multiple of ϵ. Appendix C gives formal restatements of the four protocols that fit strictly within Definitions 3 and 4, as all messages are encoded as bits. The appendix also shows that all of the propositions derived in Sections 2 and 3 continue to hold under these restated definitions.

The last row of Table 4 limits the space of matchmaking protocols to be those that allow only one round of communication. This is formalized below:

Definition 5 (One-round matchmaking protocol). A *one-round communication protocol* with m agent is a communication protocol in which each agent $i \in [m]$ sends a message μ_i based on her private information y_i alone, without seeing the messages sent by anyone else. After every agent has sent a message, the protocol terminates with an output, which is a deterministic function of the vector of messages $\mu = (\mu_1, \mu_2, \cdots, \mu_m)$. A *one-round matchmaking protocol* is defined exactly as in Definition 4, except that the phrase "communication protocol" is replaced by the phrase "one-round communication protocol."

Appendix C shows that Protocol 4 is a one-round matchmaking protocol in the sense of Definition 5. As discussed in Section 2.4, one qualitative advantage of centralized matching is that it allows all communication to be simultaneous. This advantage is not captured by the notion of communication cost in Definition 3, but is captured by Definition 5.

4.2. Derivation of the upper bounds in Table 4

The upper bounds in the first five rows of Table 4 are implied by the following proposition.

Proposition 10. *Given any fixed $\epsilon > 0$, there exists a constant $\kappa > 0$ such that for any market with preference density parameters d^I and d^J, there exists a matchmaking protocol that yields an ϵ-stable outcome*

a) with probability at least $1 - e^{-n/2}$ using per-agent communication cost $\kappa(\log n)/d^J$;

b) with probability at least $1 - e^{-n/2}$ using per-agent communication cost $\kappa(\log n)/d^I$;

c) with probability at least $1 - 2e^{-n/2}$ using per-agent communication cost $2\kappa\sqrt{n}\log n$.

The proof is given in Appendix D.4, and the three parts are based on counting the number of bits of communication used in each interaction of Protocols 1, 2, and 3 respectively, and terminating the protocol when the amount of communication reaches an upper-threshold. Since Section 3 shows that the number of interactions in these protocols is small with high probability, the threshold is not reached with high probability.

The upper bound in the last row of Table 4 is implied by the following result, which is a straightforward corollary of Proposition 9. A formal proof is given in Appendix D.3, and shows that although Proposition 9 is stated in terms of randomized recommendation sets, it also implies the existence of good deterministic recommendation sets with similar guarantees.

PROPOSITION 11. *Given $\epsilon > 0$ and any market that is relatively homogeneous with minimum segment size l, with the preference density parameters satisfying $d^I d^J \geq \min(1, (3\ln n)/l)$. There exists a one-round matchmaking protocol for this market that yields an ϵ-stable outcome with probability at least $1 - 6/n^2$ using per-agent communication cost $(6\ln n)/(d^I d^J)$.*

4.3. Derivation of the lower bounds in Table 4

The lower bounds in Table 4 are implied by the following proposition.

PROPOSITION 12 **(Lower bounds in communication cost)**. *Define $M(n, d^I, d^J)$ as the market with n agents on each side and binary preferences as given below: for each $(i,j) \in I \times J$,*

$$b_{ij} = \begin{cases} .99 & \text{with probability } d^I, \\ 0 & \text{otherwise;} \end{cases} \tag{33}$$

$$c_{ij} = \begin{cases} .01 & \text{with probability } d^J, \\ 1 & \text{otherwise.} \end{cases} \tag{34}$$

There exists universal constants $\kappa_1, \kappa_2, \kappa_3 > 0$ such that for any $\epsilon < .49$ and $n \geq 64$,

a) any matchmaking protocol that yields an ϵ-stable outcome with at least 99% probability in the market $M(n, 1/n, 1)$ requires a per-agent communication cost of at least $\kappa_1 \log n$;

b) for any $d^J \in [1/\sqrt{n}, 1/2]$, any matchmaking protocol that yields an ϵ-stable outcome with at least 99% probability in the market $M(n, 1/(nd^J), d^J)$ requires a per-agent communication cost of at least κ_2/d^J.

c) any one-round matchmaking protocol that yields an ϵ-stable outcome with at least 99% probability in the market $M(n, .5, .5)$ requires a per-agent communication cost of at least $\kappa_3 \log n$.

To see how Proposition 12 is related to Table 4 given at the beginning of Section 4, observe that the lower bound of $\Omega(\log n)$ in the first row of Table 4 is implied by part a) of Proposition 12, as $M(n, 1/n, 1)$ is a market with $d^J = 1 = \Omega(1)$. The lower bound of $\Omega(1/d^J)$ in the second row follows from part b). The lower bounds in the third and forth rows follow by symmetry, after interchanging d^I and d^J. The lower bound of $\Omega(\sqrt{n})$ in the fifth row of Table 4 follows from part b) of Proposition 12 with d^J set to $1/\sqrt{n}$. The lower bound of $\Omega(\log n)$ in the last row follows from part c), as $M(n, .5, .5)$ is ex-ante homogeneous on both sides with $d^I d^J = 1/4 = \Omega(1)$.

The proof of Proposition 12 is given in Appendices D.5 and D.6. While the proofs are written in terms of information-theoretic constructs such as entropy, mutual information and KL-divergence, the high level ideas are intuitive and are summarized below. The construction of the market $M(n, d^I, d^J)$ implies that each pair (i,j) is *profitable*, meaning that $(b_{ij}, c_{ij}) = (.99, .01)$, with probability $d^I d^J$. Unprofitable pairs have negative surpluses. For any $\epsilon < .49 = (.99 - .01)/2$, the agents in a profitable pair (i,j) cannot both be unmatched, as it is better for them to be matched to one another. For each agent, the expected number of profitable partners is equal to $d^I d^J n$.

In parts a) and b), $d^I d^J = 1/n$, which implies that each agent has one profitable partner in expectation. If a profitable pair (i, j) is such that the two agents do not have any other profitable partners, then refer to it as a *perfect* pair. One can show that the expected number of perfect pairs is at least $e^{-2}n$, so each agent is in a perfect pair with constant probability. Any ϵ-stable outcome must match every perfect pair. In part a), this requires each customer i in a perfect pair (i, j) to identify her unique profitable partner j, which requires $\lceil \log_2 n \rceil$ bits of communication. This explains the $\Omega(\log n)$ bound in per-agent communication cost.

In part b), even if an agent is in a perfect pair, finding a profitable partner requires significant interaction between agents. If customers reach out to providers, each would need to reach out to $1/d^J$ providers in expectation before finding one who can serve at a low cost. On the other hand, if providers reach out to customers, each would need to reach out to $1/d^I$ customers in expectation. Since $1/d^I \geq 1/d^J$ by construction, it is more efficient for customers to initiate contact. This explains the $\Omega(1/d^J)$ bound in part b). While the above reasoning is intuitive, a more elaborate argument is needed to prove that no matchmaking protocol can beat this bound, and the argument in the proof is closely related to the information-theoretic analysis of the Disjointness problem from the communication complexity literature (Bar-Yossef et al. 2004, Braverman 2015).

In part c), $d^I d^J = 1/4$, so each agent has many profitable partners. Suppose that each customer indicates high preferences for a certain subset of k providers, such a message would require at least $\Omega(k)$ bits. If $k \leq \alpha \log_2 n$ for some $\alpha < 1$, then in expectation, there are $n2^{-k} \geq n^{1-\alpha}$ customers for whom none of the providers they like reciprocate interest. Refer to such customers as being *unreciprocated*. Having many unreciprocated customers makes it impossible for the protocol to obtain an ϵ-stable outcome with high probability. To see this, note that if the protocol matches an unreciprocated customer, it must be to a provider for whom she has not indicated interest. The surplus of this match may be negative, which violates individual rationality. On the other hand, if all unreciprocated customers are unmatched, then in expectation there would be $n^{1-\alpha}$ unmatched customers. However, this cannot be, as in this market, the maximum number of unmatched customers in any ϵ-stable outcome is at most $O(\log n)$ with high probability.[4] Hence, $k \geq \Omega(\log n)$, which explains the lower bound in part c).

5. Extensions

The results derived in this paper continue to hold in more general settings.

[4] In the market $M(n, .5, .5)$, the chance that there are m unmatched agents in an ϵ-stable matching is upper-bounded by $\binom{n}{m}^2 (3/4)^{m^2}$, as this is the probability that there are m customers and providers with no profitable matches between them. This bound implies that $m \leq O(\log n)$ with high probability.

5.1. Correlated preferences

The assumption that the b_{ij}'s and c_{ij}'s are independently drawn for each $(i, j) \in I \times J$ is unnecessary for much of the analysis. The proofs of Propositions 2 and 3 do not use the independence assumption at all, so Protocols 1, 2 and 3 are guaranteed to terminate with an ϵ-stable outcome even if preferences are arbitrarily correlated. The performance guarantees for Protocols 1 and 3 in Propositions 4, 6, 7 and 10 only require each c_{ij} to be independently drawn and independent of the b_{ij}'s, but the matrix b of customer benefits can be arbitrarily correlated within itself.

If correlation in preferences exhibit certain structure, the platform can utilize the structure to estimate the correlated components, so that the residual preferences are effectively independent. For example, Appendix E shows how Protocols 1 and 3 can be modified to handle the following type of correlation in provider costs: for each $(i, j) \in I \times J$, $c_{ij} = \alpha_i + \beta_j + \gamma_{ij}$, where α_i is an unobserved fixed effect for each customer i, β_j is an unobserved fixed effect for each provider j, and γ_{ij} is a residual term that is assumed to be independently drawn from known distributions.

5.2. Many-to-one matching

The analysis in this paper can easily be generalized to allow one side of the market to have multiple match partners. For example, suppose that each provider j has the capacity to serve q_j customers. This can be transformed into a one-to-one matching model by replacing each provider j with q_j identical copies of himself, each with unit capacity. Each ϵ-stable outcome in this one-to-one model maps back to an ϵ-stable outcome in the original many-to-one model. (See Appendix F for a formal definition of ϵ-stable outcomes in the many-to-one model.) Hence, it suffices to consider the one-to-one model in which certain providers have identical cost vectors. In this case, Protocols 1 and 3 can be modified by having customers who receive updated cost information from a provider to also update her cost estimates for the identical copies of that provider. In Protocol 2, no modification is needed, as the current analysis already allows for arbitrary preference correlations among agents on the searching side, as pointed out in Section 5.1. In Protocol 4, the recommendation sets need to be selected such that two copies of the same provider are recommended to disjoint sets of customers. This implies that a provider with a larger capacity should be recommended to more customers, so as to maximize the chance that his capacity can be fully utilized.

5.3. Dynamic arrival of customers

While the paper abstracts away from issues of timing by studying a static model, Appendix G shows that the main insights continue to hold in a model where customers arrive to the market according to a stationary stochastic process while providers remain perpetually in the market. Each provider is associated with a certain capacity, which represents the maximum rate at which he can serve customers in the long run.

The appendix shows that suitably modified versions of the four matchmaking protocols provide near-optimal performance guarantees for their respective class of suitable markets, exactly as in Table 4. The protocols match agents whenever m new customers arrive, where m is referred to as the *batch size.* A new insight is that the four protocols have different requirements on batch size: the performance guarantees for the dynamic versions of Protocols 1 and 4 allow $m = 1$, which means that customers can be matched immediately upon arrival. On the other hand, optimal performance of Protocols 2 and 3 requires m to be sufficiently large. This corresponds to the intuitive fact that in order for providers to efficiently search across customers, there needs to be a certain market thickness on the customer side, as it is inefficient to require providers to make a decision on whether to reach out to each customer as soon as she arrives.

5.4. Alternative model without transfer payments

While the results in this paper are presented a model with transfer payments, Appendix H shows that the insights also hold under an analogous model without transfer payments. Such a model is a better fit when monetary transfers are unimportant or when payments are exogenously determined, such as in dating markets or in the academic job market. The appendix presents modified versions of the four protocols under such a model and rederives all of the paper's results. The main difference is that instead of helping agents to estimate the transaction price, the protocols help agents to estimate which partners they can realistically be matched with, which plays the role of "prices" in such markets and is also determined endogenously based on supply and demand.

6. Managerial insights for two-sided marketplaces

The analysis in this paper suggests the following guidelines for the practical implementation of two-sided marketplaces. First, it is important for platforms to understand the nature of preferences for the intended user base: i.e., whether they are easy to describe using a short questionnaire. If preferences on both sides are easy to describe, then centralized matching is a good strategy. If preferences on one side are hard to describe, then the platform should not attempt to do centralized matching, but instead help agents with hard-to-describe preferences search among potential partners and initiate contact. If both sides have hard-to-describe preferences, then the platform should allow both sides to search and initiate contact. One way to estimate whether agents' preferences are hard to describe is to observe how many partners they compare before deciding to match, or how often they reciprocate interest when contacted. (In the analysis, the primary role of the parameters d^I and d^J is to capture the likelihood of reciprocating interest.) It is possible that among agents on the same side of the market, some have hard-to-describe preferences while others have easy-to-describe preferences. In this case, the platform may find it beneficial to implement different matchmaking strategies at different sections of its website, and direct users to the most appropriate section based on the complexity of their needs.

Secondly, the platform should do its best to make preferences easier to describe. This could be done, for example, by employing a more specific, descriptive or quantitative questionnaire, or by asking users to choose among a set of carefully curated keywords so as to segment them into homogeneous sub-markets. By reducing the uncertainty in user preferences, the platform can minimize the amount of back and forth communication needed during the matchmaking process.

Last but not least, the platform should ensure that good price estimates are available. Many two-sided marketplaces already invest efforts in helping users to estimate prices. For example, the home services platforms Angi, HomeAdvisor and Thumbtack all have dedicated portions on their websites giving reasonable ranges of costs for various projects depending on the customer's location. Labor platforms such as Taskrabbit and Upwork ask providers to post their hourly rate, and the platforms publish a reasonable range of hourly rate for each type of work. Amazon Home Services goes one step further and posts standardized prices for each service item that is supported by the platform. The paper suggests that these platforms should ensure that their price estimates are "optimistic but attainable," which is a common feature of the near-optimal protocols in this paper. Instead of providing a wide range of prices or estimating the average or median price, the platforms should focus on estimating what is a "good reasonable price," which should be better than the average quoted price but is nevertheless possible to obtain after comparing several quotes. Given such price estimates, customers would naturally contact providers based on an optimistic expectation of prices, as prescribed in Protocols 1 and 3, and providers would not need to communicate separately with many customers about special deals, since many of these deals would already be incorporated in the platform's price estimates. As an alternative to providing price estimates directly, a platform can also direct users to third-party websites that specialize in estimating what is a good reasonable price for various services. Such websites exist in many industries and the paper's analysis suggests that they also can play an important role in improving the efficiency of two-sided marketplaces in their respective industries.

References

Abdulkadiroğlu A, Che YK, Yasuda Y (2015) Expanding "choice" in school choice. *American Economic Journal: Microeconomics* 7(1):1–42.

Arnosti N, Johari R, Kanoria Y (2019) Managing congestion in decentralized matching markets. Manufacturing and Service Operations Management, Forthcoming.

Ashlagi I, Braverman M, Kanoria Y, Shi P (2019) Clearing matching markets efficiently: Informative signals and match recommendations. *Management Science* 66(5).

Avery C, Jolls C, Posner RA, Roth AE (2001) The market for federal judicial law clerks. *The University of Chicago Law Review* 68(3):793–902.

Bar-Yossef Z, Jayram TS, Kumar R, Sivakumar D (2004) An information statistics approach to data stream and communication complexity. *Journal of Computer and System Sciences* 68(4):702–732.

Bertsekas DP (1988) The auction algorithm: A distributed relaxation method for the assignment problem. *Annals of Operations Research* 14(1):105–123, ISSN 1572-9338.

Braverman M (2015) Interactive information complexity. *SIAM Journal on Computing* 44(6):1698–1739.

Chou JH, Lu CJ (2010) Communication requirements for stable marriages. *Proceedings of the 7th International Conference on Algorithms and Complexity*, 371–382 (Springer).

Coles P, Cawley J, Levine PB, Niederle M, Roth AE, Siegfried JJ (2010) The job market for new economists: A market design perspective. *The Journal of Economic Perspectives* 24(4):187–206.

Coles P, Kushnir A, Niederle M (2013) Preference signaling in matching markets. *American Economic Journal: Microeconomics* 5(2):99–134.

Cover TM, Thomas JA (2012) *Elements of Information Theory* (John Wiley & Sons).

Demange G, Gale D, Sotomayor M (1986) Multi-item auctions. *Journal of Political Economy* 94(4):863–872.

Erdös P, Rényi A (1964) On random matrices. *Publ. Math. Inst. Hungar. Acad. Sci.* (8).

Fradkin A (2017) Search, matching, and the role of digital marketplace design in enabling trade: Evidence from Airbnb. Working paper, Boston University.

Gale D, Shapley LL (1962) College admissions and the stability of marriage. *American Mathematical Monthly* 69:9–15.

Gonczarowski YA, Nisan N, Ostrovsky R, Rosenbaum W (2015) A stable marriage requires communication. *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, 1003–1017 (SIAM).

Halaburda H, Jan Piskorski M, Yildirim P (2018) Competing by restricting choice: The case of matching platforms. *Management Science* 64(8):3574–3594.

Horton JJ (2017) The effects of algorithmic labor market recommendations: Evidence from a field experiment. *Journal of Labor Economics* 35(2):345–385.

Horton JJ (2018) Buyer uncertainty about seller capacity: Causes, consequences, and a partial solution. *Management Science* 65(8).

Jagadeesan M, Wei A (2018) Varying the number of signals in matching markets. Christodoulou G, Harks T, eds., *Web and Internet Economics*, 232–245 (Cham: Springer International Publishing).

Kanoria Y, Saban D (2017) Facilitating the search for partners on matching platforms: Restricting agent actions. Working paper, Columbia University.

Kushilevitz E, Nisan N (2006) *Communication Complexity* (Cambridge University Press).

Lee S, Niederle M (2015) Propose with a rose? Signaling in internet dating markets. *Experimental Economics* 18(4):731–755.

Li J, Netessine S (2019) Higher market thickness reduces matching rate in online platforms: Evidence from a quasi-experiment. *Management Science* 66(1).

Newton C (2014) Taskrabbit is blowing up its business model and becoming the Uber for everything. The Verge (June 17, 2014). Available at `https://www.theverge.com/2014/6/17/5816254/taskrabbit-blows-up-its-auction-house-to-offer-services-on-demand`.

Roth AE (2008) What have we learned from market design? *The Economic Journal* 118(527):285–310.

Roth AE, Sotomayor O (1990) *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis* (Cambridge University Press).

Roth AE, Xing X (1994) Jumping the gun: Imperfections and institutions related to the timing of market transactions. *The American Economic Review* 84(4):992–1044.

Roth AE, Xing X (1997) Turnaround time and bottlenecks in market clearing: Decentralized matching in the market for clinical psychologists. *Journal of political Economy* 105(2):284–329.

Segal I (2007) The communication requirements of social choice rules and supporting budget sets. *Journal of Economic Theory* 136(1):341–378.

Shapley LS, Shubik M (1971) The assignment game I: The core. *International Journal of game theory* 1(1):111–130.

Shi P (2022) Optimal match recommendations in two-sided marketplaces with endogenous prices. Available at SSRN: `https://ssrn.com/abstract=4034950`.

Shieber J (2018) How did Thumbtack win the on-demand services market? by being the everything store. Tech Crunch (May 5, 2018). Available at `https://techcrunch.com/2018/05/05/how-did-thumbtack-win-the-on-demand-services-market/`.

Wu Y (2020) Yelp reinvents the hiring experience for home and local services. Yelp Official Blog (August 18,2020). Available at `https://blog.yelp.com/2020/08/yelp-reinvents-the-hiring-experience-for-home-and-local-services`.

Online Appendix

Table of Appendices

Appendix A: Strategic agents

The main body of the paper assumes that agents are non-strategic. This section argues that this is a reasonable approximation in large markets. As shown in Section 7.2 of Roth and Sotomayor (1990), one opportunity for strategic behavior in an assignment game is if agents on the same side collude. For example, all customers can collude as in a bidding ring and decide on a matching. They can reduce payments so as to pay all providers zero surplus, and re-distribute the gains among themselves via side payments. However, such a high level of coordination is unrealistic in practice. In this section, we analyze the incentive for an agent to unilaterally deviate from the protocols, assuming that everyone else partcipates truthfully.

Among the four protocols, centralized matching has the strongest incentive guarantee, as the match recommendations and transaction prices are computed using commonly known information alone, and deviating from the protocol means either to reject a near-optimal match or to accept a sub-optimal one, neither of which can result in much gain for the agent.

Proposition A.1 **(Incentive guarantees for centralized matching)**. *Under the assumptions of Proposition 9, with probability at least* $1-3/n^2$*, Protocol 4 returns an* ϵ*-stable outcome, and no agent can unilaterally deviate from the protocol and obtain a surplus that is more than* ϵ *higher than what the agent gets when following the protocol.*

Proof of Proposition A.1. Let $\hat{I}$ denote the set of customers matched in $\hat{x}$, as defined in the statement of Proposition 9. Under the assumptions of Proposition 9, with probability at least $1-3/n^2$, Protocol 4 when everyone follows the protocol matches every customer in $\hat{I}$. Every customer not in $\hat{I}$ receives zero surplus and cannot do anything about it because the protocol does not elicit any inputs from her. Every customer $i \in \hat{I}$ receives a surplus of at least $\overline{b}_{ij} - \epsilon - p_i \geq 0$, whereas the maximum possible surplus she can get from the protocol is $\overline{b}_{ij} - p_i$, since she cannot influence the transaction price p_i nor the segment of providers she is matched to. Hence, she cannot improve her surplus by more than ϵ regardless of what she does.

Consider now the providers. For providers in a segment k that is unmatched in $\hat{x}$ (i.e., a segment k such that $\sum_{i\in I, j\in J_k} \hat{x}_{ij} = 0$), they receive zero surplus and cannot do anything about it. For providers in a fully-matched segment k (i.e., a segment k such that $\sum_{i\in I, j\in J_k} \hat{x}_{ij} = |J_k|$), the same argument applies as with customers in $\hat{I}$. The only remaining case to consider is a partially matched segment (i.e., a segment k for which $0 < \sum_{i\in I, j\in J_k} \hat{x}_{ij} < |J_k|$). For such a segment k, the optimal dual variable $\hat{v}_k$ to (29) would be equal to zero, which implies that providers in such segments are paid $\underline{c}_{ij} + \epsilon$ for being matched to any customer i. Therefore, regardless of what they do, the maximum surplus they can possibly get is ϵ, whereas they are guaranteed a non-negative surplus if they follow the protocol. □

For the search-based protocols, since Protocol 1 is a special case of Protocol 3 and Protocol 2 is the mirror image of Protocol 1, it is sufficient to analyze the incentives to unilaterally deviate from Protocol 3. For a customer i, deviating from Protocol 3 means to contact a provider j who does not maximize her estimated surplus $U_{ij}(\hat{v}, \hat{c}) = b_{ij} - \hat{v}_j - \hat{c}_{ij}$. Such an action is risky since it is possible that the provider j^* who maximizes $U_{ij}(\hat{v}, \hat{c})$ might be able to serve her at cost $\hat{c}_{ij^*}$, in which case she would be forgoing a near-optimal match by not contacting j^*. For a provider, deviating from the protocol means to misrepresent his cost when communicating with customers. However, communicating a cost that is lower than actual would reduce the customer's offer price and cut into his surplus, and communicating a cost that is higher than actual would risk diverting customers to his competitors. When the market is large and there are many competitors, an individual provider has limited market power and can gain little by raising prices beyond that which balances his supply and demand. For a formal treatment of the idea that an individual participant in a large assignment game has little market power, see Hassidim and Romm (2014) and Kanoria et al. (2018).

References

Hassidim A, Romm A (2014) An approximate "law of one price" in random assignment games. Available at `https://arxiv.org/abs/1404.6103`.

Kanoria Y, Saban D, Sethuraman J (2018) Convergence of the core in assignment markets. *Operations Research* 66(3):620–636.

Appendix B: Precise statement of Protocol 2 (providers search)

Given $\epsilon > 0$.

1. **Initialization:** Initialize the matching x to be empty and the price vector p to be zero. Initialize the set of active searching providers as $A = J$. For each customer i, initialize her required surplus as $\hat{u}_i = 0$. For each provider $j \in J$, initialize his benefit estimate as $\hat{b}_{ij} = \bar{b}_{ij} - \epsilon$ for each $i \in I$, where $\bar{b}_{ij}$ is the upper support of the benefit distribution B_{ij}. Define his estimated surplus from serving each provider i as
$$V_{ij}(\hat{u}, \hat{b}) := \hat{b}_{ij} - \hat{u}_i - c_{ij}. \tag{B.1}$$
2. **Providers search for customers:**
   (a) **Provider selection:** If $|A| = 0$, then terminate the protocol with the transaction outcome (x, p). Otherwise, select an arbitrary provider $j \in A$ and proceed to step (b).
   (b) **Provider interaction:** For the selected provider j, if $V_{ij}(\hat{u}, \hat{b}) < 0$ for all $i \in I$, then the provider communicates that he is no longer interested in any customer, and the protocol removes j from A and goes back to step (a). Otherwise, provider j contacts a customer
$$i^* \in \arg\max_{i \in I} \{V_{ij}(\hat{u}, \hat{b})\}, \tag{B.2}$$
   and asks whether she is willing to be served by him at price $\hat{b}_{ij} - \hat{u}_i$.
   - If $b_{i^*j} \geq \hat{b}_{i^*j}$, the customer answers "yes" and becomes matched to j: update $x_{i^*j} = 1$, $p_{i^*} = \hat{b}_{i^*j} - \hat{u}_{i^*}$, and remove j from A. If i^* was previously matched to another provider j', then j' becomes unmatched: update $x_{i^*j'} = 0$, and add j' to A. Customer i^* increments $\hat{u}_{i^*}$ by ϵ and communicates this to the market. Go back to step (a).
   - If $b_{i^*j} < \hat{b}_{i^*j}$, the customer answers "no" and communicates her actual benefit b_{i^*j} to the provider. The provider updates $\hat{b}_{i^*j} = b_{i^*j}$. Repeat step (b).

Appendix C: Restatements of the four protocols which satisfy Definition 4

The following restated definitions of the four protocols make it explicit how all communication is encoded using binary bits, thus satisfying the the formal framework in Definition 4, which is based on Definition 3. Without loss of generality, assume that agents on each side of the market are indexed from 1 to n, so that the identity of a potential partner can be encoded using $\lceil \log_2 n \rceil$ bits.

Among Protocols 1, 2 and 3, note that Protocol 1 is a special case of Protocol 3 with $z_{ij} \equiv \underline{c}_{ij}$, and Protocol 2 is the mirror image of Protocol 1. Therefore, it suffices to state Protocol 3.

Restatement of Protocol 3 *Given $\epsilon > 0$, as well as an $I \times J$ matrix z that depends only on the commonly known market information (I, J, B, C).*

1. **Initialization:** *Initialize the matching x to be empty, the price vector p to be zero, and the set $A = I$. For each $j \in J$, initialize $\hat{v}_j = 0$. For each $(i, j) \in I \times J$, initialize $\hat{c}_{ij} = z_{ij} + \epsilon$. The variables x, p, A, $\hat{v}$ and $\hat{c}$ can be interpreted as state variables of the protocol, which are initialized based on commonly known information. Since the protocol is deterministic, they remain commonly known throughout the protocol. Define $U_{ij}(\hat{v}, \hat{c}) := b_{ij} - \hat{v}_j - \hat{c}_{ij}$.*

2. **Providers search for customers:** *Each provider* $j \in J$ *sends a message* $(i, \Delta_{ij} := \lceil (c_{ij} - \hat{c}_{ij})/\epsilon \rceil)$ *for each* $i \in I$ *such that* $c_{ij} < z_{ij}$*. The protocol updates* $\hat{c}_{ij} = \hat{c}_{ij} + \epsilon\Delta_{ij}$*. Note that each such message can be encoded using* $\lceil \log_2 n \rceil + \lceil \log_2(1/\epsilon) \rceil$ *bits. After each provider is finished contacting customers, he uses one bit to communicate that he is done.*
3. **Customers search for providers:**
   (a) **Customer selection:** *If* $|A| = 0$*, then terminate the protocol with the transaction outcome* (x, p)*. Otherwise, select the customer* $i \in A$ *with the smallest index*[5] *and proceed to step 3(b).*
   (b) **Provider interaction:** *The selected customer* i *sends the message "0" if* $U_{ij}(\hat{v}, \hat{c}) < 0$ *for all* $j \in J$*, in which case the protocol removes* i *from* A *and go back to step 3(a). Otherwise, customer* i *selects the provider* $j^* \in \arg\max_j \{U_{ij}(\hat{v}, \hat{c})\}$ *with the smallest index and sends the index of* j^**. In either case, customer* i*'s message can be encoded in* $\lceil \log_2(1+n) \rceil$ *bits. The provider* j^* *responds with one bit indicating whether* $c_{ij^*} \le \hat{c}_{ij^*}$*.*
   - *If "yes," then the protocol increments* $\hat{v}_{j^*}$ *by* ϵ*, updates* $x_{ij^*} = 1$*,* $p_i = \hat{v}_{j^*} + \hat{c}_{ij^*}$*, and removes* i *from* A*. If* j^* *was previously matched to another customer* i'*, then the protocol updates* $x_{i'j^*} = 0$*,* $p_{i'} = 0$*, and adds* i' *to* A*. Go back to step 3(a).*
   - *If "no," then provider* j^* *sends another message using* $\lceil \log_2(1/\epsilon) \rceil$ *bits encoding* $\Delta_{ij^*} := \lceil (c_{ij^*} - \hat{c}_{ij^*})/\epsilon \rceil$*. The protocol updates* $\hat{c}_{ij^*} = \hat{c}_{ij^*} + \epsilon\Delta_{ij^*}$*. Repeat step 3(b).*

Besides rewording the description so it is clear how it fits within the definition of a communication protocol in Definition 3, the only change in the above from the original definition of Protocol 3 in Section 2.4 is that provider j communicates $\Delta_{ij} = \lceil (c_{ij} - \hat{c}_{ij})/\epsilon \rceil$ instead of the exactly cost c_{ij}, and the cost estimate of customer i is updated to $\hat{c}_{ij} = \hat{c}_{ij} + \epsilon\Delta_{ij}$. By construction, this ensures that $\hat{c}_{ij} \le c_{ij} + \epsilon$ throughout Step 3 of the protocol. Moreover, a provider j who communicates his updated cost information Δ_{ij} to customer i would never respond with "no" the next time customer i indicates interest in Step 3(b). Finally, under this modification, the probability that provider j responds with "yes" in Step 3(b) remains the same as before, which is equal to $\mathbb{P}(c_{ij} \le z_{ij} + \epsilon | c_{ij} \ge z_{ij})$ if j has not yet sent updated cost information to customer i, and equal to 1 otherwise. Therefore, the statements and proofs of Propositions 2, 3, 4, 5, 6 and 7 continue to hold under the restated description of Protocols 1, 2 and 3, without requiring any change at all.

The following definition of Protocol 4 is equivalent to the one in Section 2.4, but the format by which communication takes place is spelled out more precisely. Note that the total number of bits of communication is equal to $2\sum_{i\in I} |S_i|$.

Restatement of Protocol 4 *Given* $\epsilon > 0$*, price vector* $p \in \mathbb{R}^n$*, a set* $S_i \subseteq J$ *for each customer* $i \in I$*. Define the inverse mapping* $T_j := \{i \in I : j \in S_i\}$*.*

1. **Preference elicitation:** *Each customer* $i \in I$ *sends a message with* $|S_i|$ *bits. Each bit* $k \in \{1, 2, \cdots |S_i|\}$ *indicates whether* $b_{ij} \ge \bar{b}_{ij} - \epsilon$*, where* j *is the provider in* S_i *with the* k*th smallest index. Similarly, each provider* $j \in J$ *sends a message with* $|T_j|$ *bits indicating whether* $c_{ij} \le \underline{c}_{ij} + \epsilon$ *for each recommended partner* $i \in T_j$*. For each recommended pair* (i, j)*, if both agents indicated "yes" regarding this pair, then the pair is said to be accepted.*

[5] Any other deterministic tie-breaking rules can be used.

2. **Match determination:** *Compute a maximum cardinality matching x among the accepted recommended pairs. Terminate the protocol with the transaction outcome (x, p).*

Appendix D: Omitted proofs

D.1. Proof of Lemma 1

If $n \le 4$, then $(3\ln n)/n > 1$, and a perfect matching always exists since we have a complete graph. This is also the case if $r \ge 1$. Assume for the remainder of the proof that $n \ge 5$ and $r < 1$.

Let the node sets in the bipartite graph be denoted as I and J, with $|I| = |J| = n$. By Hall's Theorem, if a perfect matching does not exist, then there exists sets $S \subseteq I$ and $T \subseteq J$ with $|S| + |T| = n + 1$ such that there are no edges between S and T. (This observation is also the starting point of the analysis in Erdös and Rényi (1964), but their analysis is optimized for the case in which $r \approx (\ln n)/n$, whereas we derive tighter bounds for the case in which $r \ge (3\ln n)/n$.) Let $Q(n, r)$ be the probability that a perfect matching does not exist, then,

$$Q(n,r) \le \sum_{k=1}^{n} \binom{n}{k} \binom{n}{n+1-k} (1-r)^{k(n+1-k)} \tag{D.1}$$

$$\le n e^{-rn} \sum_{k=1} q(n,r,k), \tag{D.2}$$

$$\text{where } q(n,r,k) := \frac{1}{n} \binom{n}{k} \binom{n}{k-1} (1-r)^{k(n+1-k)-n}. \tag{D.3}$$

Since $q(n, r, k)$ is non-increasing in r, it suffices to prove that $F(n) < 2.1$ for all $n \ge 5$, where

$$F(n) := \sum_{k=1}^{n} q\left(n, \frac{3\ln n}{n}, k\right). \tag{D.4}$$

It is straightforward to verify numerically that

$$\max_{5 \le n \le 49} \{F(n)\} = F(16) = 2.0083... < 2.1. \tag{D.5}$$

It remains to show that $F(n) < 2.1$ for all $n \ge 50$. Note that $q(n, (3\ln n)/n, k) \le f(n, k)$, where

$$f(n,k) := \frac{1}{n} \binom{n}{k} \binom{n}{k-1} \exp\left(-\frac{3\ln n}{n}[k(n+1-k)-n]\right). \tag{D.6}$$

Observe that $f(n, k) = f(n, n+1-k)$, and $f(n, 1) = 1$. Moreover, for $n \ge 50$, $3(n-2)/n \ge 2.88$, so

$$f(n,2) = \frac{n(n-1)}{2} \exp\left(-\frac{3\ln n}{n}(n-2)\right) \le \frac{1}{2} n^{-.88}. \tag{D.7}$$

Now,

$$\frac{f(n,k+1)}{f(n,k)} = \frac{(n-k)(n-k+1)}{k(k+1)} \exp\left(-\frac{3\ln n}{n}(n-2k)\right). \tag{D.8}$$

When $k \in [2, n/6]$, $(n-2k)/n \ge 2/3$, so (D.8) is bounded above by $1/(k(k+1)) \le 1/6$. When $k \in [n/6, n/4]$, $(n-2k)/n \ge 1/3$, and $(n-k)(n-k+1)/(k(k+1)) \le 25$, so (D.8) is bounded above by $25/n \le 1/2$. Thus,

$$\sum_{k=2}^{\lceil n/3 \rceil} f(n,k) \le f(n,2) \left(1 + \frac{1}{2} + \frac{1}{2^2} + \cdots\right) \le n^{-.88}. \tag{D.9}$$

Now, when $k \in [n/3, 2n/3]$, since $\binom{n}{k} \le 2^n$, and $2(n-3)/3 \ge n/2 + 1$ for $n \ge 18$, we have

$$f(n,k) \le \frac{1}{n} 2^{2n} \exp\left(-\frac{3\ln n}{n}\left[\frac{2}{9} n(n-3)\right]\right) \le \frac{1}{n^2}\left(\frac{16}{n}\right)^{n/2} \le \frac{1}{n^2}. \tag{D.10}$$

Combining the above, we have that for all $n \ge 50$,

$$F(n) \le \sum_{k=1}^{n} f(n,k) \le 2 + 2n^{-.88} + n^{-1} < 2.1, \tag{D.11}$$

as desired. □

D.2. Proof of Proposition 9

Note that the LP given in (27)-(30) has a totally unimodular constraint set, so that in an optimal basic solution, $\hat{x}_{ij} \in \{0,1\}$ for all $(i,j) \in I \times J$. The dual LP is given by

$$\text{Minimize} \quad \sum_{i \in I} \hat{u}_i + \sum_{j \in J} \hat{v}_j \tag{D.12}$$

$$\text{s.t.} \quad \hat{u}_i + \hat{v}_j \geq \hat{b}_{ij} - \hat{c}_{ij} \qquad \text{for each } (i,j) \in I \times J, \tag{D.13}$$

$$\hat{u}, \hat{v} \geq 0 \tag{D.14}$$

At an optimal solution, if $j, j' \in J_k$, then $\hat{v}_j = \hat{v}_{j'}$, as otherwise it would be possible to reduce the objective (D.12). Therefore, by complementary slackness, if any provider $j' \in J_k$ is unmatched in $\hat{x}$, then $\hat{v}_j = 0$ for all $j \in J_k$. Also by complementary slackness, if a customer $i \in I$ is unmatched in $\hat{x}$, then $\hat{u}_i = 0$. Let $\hat{I}$ be set of customers who are matched in $\hat{x}$.

We show that if the matching x returned by Protocol 4 matches every customer in $\hat{I}$, then the corresponding outcome (x,u,v) is ϵ-stable. To see this, observe that if (i,j) is matched in x, then

$$u_i = b_{ij} - p_i \geq \hat{b}_{ij} - p_i = \hat{u}_i \tag{D.15}$$

$$v_j = p_i - c_{ij} \geq p_i - \hat{c}_{ij} = \hat{v}_j \tag{D.16}$$

Moreover, if i is not matched in x, then $i \notin \hat{I}$ by assumption, which implies that $\hat{u}_i = 0 = u_i$. If j is not matched in x, then $v_j = 0$, but the assumption implies that not everyone in J_k is matched in $\hat{x}$, so $\hat{v}_j = 0 = v_j$. Summarizing, we have $u_i \geq \hat{u}_i$ for all $i \in I$, and $v_j \geq \hat{v}_j$ for all $j \in J$. Combining this with (D.13), we have that for all $(i,j) \in I \times J$,

$$u_i + v_j \geq \hat{b}_{ij} - \hat{c}_{ij} = (\bar{b}_{ij} - \epsilon) - (\underline{c}_{ij} + \epsilon) \geq b_{ij} - c_{ij} - 2\epsilon. \tag{D.17}$$

Moreover, by (D.14), we have $u \geq \hat{u} \geq 0$ and $v \geq \hat{v} \geq 0$. So (x,u,v) is ϵ-stable.

If suffices to show that the matching x returned by the protocol matches every customer in $\hat{I}$ with probability at least $1 - 3/n^2$. We make use of the following slight generalization of Lemma 1.

LEMMA D.1. *Given $n \geq m \geq 1$ and $r \geq (3 \ln n)/n$. Consider a random bipartite graph with m nodes on one side and n nodes on the other side. For each pair of nodes (i,j), there is an edge between them with probability $q_{ij} \geq \min(1,r)$, independently across (i,j) pairs. The final graph contains a matching of cardinality m with probability at least $1 - 2.1ne^{-rn}$.*

Proof of Lemma D.1. Let the node sets be I and J with $|I| = m$ and $|J| = n$. The above graph can be embedded in a larger random graph with node sets $I' \supseteq I$ and J, with $|I'| = |J| = n$, as follows: for each of the new nodes $i \in (I' \backslash I)$, include an edge to every node $j \in J$. The original random graph has a matching of cardinality m if and only if the larger graph has a perfect matching, which occurs with probability at least $1 - 2.1ne^{-rn}$ by Lemma 1. □

Returning to the proof of Proposition 9, define $\hat{I}_k$ as the set of customers matched in $\hat{x}$ to a provider of segment k. For each provider segment $k \in K$, suppose that $a_k = 1$ and $d^I d^J = 1$, then every customer in $\hat{I}_k$

is matched in x, as every pair $(i,j)\in \hat{I}_k \times J_k$ is an accepted recommended pair. On the other hand, if either $a_k < 1$ or $d^I d^J < 1$, then each pair $(i,j)\in \hat{I}_k \times J_k$ is an accepted recommended pair with probability at least

$$a_k d^I d^J \geq \frac{3\ln n}{|J_k|}. \tag{D.18}$$

By Lemma D.1, the probability that the protocol does not match every customer in $\hat{I}_k$ is at most $2.1|J_k|n^{-3}$. Taking an union bound over the provider segments, we have that the probability that the protocol does not match every customer in $\hat{I}$ is at most

$$\sum_{k\in K} 2.1|J_k|n^{-3} \leq \frac{2.1}{n^2} < \frac{3}{n^2}, \tag{D.19}$$

as desired. □

D.3. Proof of Proposition 11

Given $\epsilon > 0$ and a market (I,J,B,C) that is relatively homogeneous with minimum segment size l, define $Q(p,S)$ as the probability that Protocol 4 with price vector p and recommendation sets $S := (S_i : i \in I)$ yields an ϵ-stable matching in this market. Proposition 9 implies that there exists a distribution F such that

$$\mathbb{E}_{S\sim F}[Q(p,S)] \geq 1 - \frac{3}{n^2}, \tag{D.20}$$

$$\text{and } \mathbb{E}_{S\sim F}\left[\frac{1}{n}\sum_{i\in I}|S_i|\right] \leq \frac{3\ln n}{d^I d^J}. \tag{D.21}$$

Define

$$\Gamma = \left\{S : \frac{1}{n}\sum_{i\in I}|S_i| \leq \frac{6\ln n}{d^I d^J}\right\}. \tag{D.22}$$

By Markov's inequality, $\mathbb{P}_{S\sim F}(S\in\Gamma) \geq 1/2$. Moreover, since $Q(p,S)$ can never exceed 1, we have

$$\mathbb{E}_{S\sim F}\left[Q(p,S)|S\in\Gamma\right] \geq 1 - \frac{6}{n^2}. \tag{D.23}$$

Therefore, there exists a deterministic choice of $S^* \in \Gamma$ such that $Q(p,S^*) \geq 1 - 6/n^2$.

Consider Protocol 4 with parameters ϵ, p and S^*. As shown in Appendix C, this is a one-round matchmaking protocol in which the total number of bits of communication is equal to $\sum_{i\in I} 2|S_i^*|$. By construction, the protocol yields an ϵ-stable matching with probability at least $1-6/n^2$, and the per-agent communication cost is

$$\frac{1}{2n}\sum_{i\in I} 2|S_i^*| \leq \frac{6\ln n}{d^I d^J}. \tag{D.24}$$

□

D.4. Proof of Proposition 10.

Since all of the probabilistic guarantees require $n \geq 2$ to be meaningful, assume that $n \geq 2$. Define

$$\kappa = \max_{n\geq 2}\left\{\left\lceil\frac{1}{\epsilon}\right\rceil \frac{\lceil\log_2(n+1)\rceil + \lceil\log_2(1/\epsilon)\rceil + 1}{\log n} + \frac{1}{2\log n}\right\}. \tag{D.25}$$

Note that κ is guaranteed to be finite since for any fixed $\epsilon > 0$, the maximand is less than $2\lceil 1/\epsilon\rceil + 1$ for large n.

For part a), consider running the slightly modified version of Protocol 1 described in Section 4.1, but terminating it after the rth interaction, where $r := 2\lceil 1/\epsilon \rceil n/d^J$. By Propositions 2 and 4, the probability that this protocol yields an ϵ-stable matching is at least $1 - e^{-n/2}$. The number of bits communicated in each interaction is at most $y := \lceil \log_2(n+1) \rceil + \lceil \log_2(1/\epsilon) \rceil + 1$. This is because $\lceil \log_2(n+1) \rceil$ bits suffices for a customer to communicate whether she wants to reach out to another provider and if so, the identity of the provider; $\lceil \log_2(1/\epsilon) \rceil$ bits suffices for the provider to communicate the updated cost information $\Delta_{ij} = \lceil (c_{ij} - \hat{c}_{ij})/\epsilon \rceil$; and one bit is used to encode the provider's yes/no response. The total number of bits communicated per agent is at most

$$\frac{1}{2n}(ry + n) \leq \kappa \frac{\log n}{d^J}. \tag{D.26}$$

The "$+n$" in the above corresponds to allowing each customer to have one additional bit to communicate that she is done contacting providers. Part b) follows from part a) by symmetry.

Part c) follows from a similar argument as in a), except that we use Protocol 3 instead of Protocol 1, and re-define $r = 4\lceil 1/\epsilon \rceil n^{1.5}$. By Propositions 3 and 7, this protocol terminates with an ϵ-stable outcome with probability at least $1 - 2e^{-n/2}$. The number of bits communicated per agent is at most $(ry + n + n)/2n \leq 2\kappa\sqrt{n}\log(n)$, where y is defined as in part a), and the additional "$+n$" corresponds to each provider in Step 2 needing one bit to communicate that he is done. □

D.5. Proof of Proposition 12 a) and b)

The proof for a) and b) begins in the same way. Given a matchmaking protocol that yields an ϵ-stable outcome with at least 99% probability in a market $M(n, d^I, d^J)$, where $d^I d^J = 1/n$. Define the *transcript* Π of the protocol as the entire sequence of bits sent in all the messages until the protocol terminates. Note that Π is a random variable whose randomness is induced by the randomly drawn inputs. let L be the maximum length of the protocol for any input. By Definition 3, L/n is the per-agent communication cost. Define the following random variables:

$$G_{ij} = \mathbb{1}\left((b_{ij}, c_{ij}) = (.99, .01)\right), \tag{D.27}$$

$$P_{ij} = G_{ij}\left[\prod_{j' \neq j}(1 - G_{ij'})\right]\left[\prod_{i' \neq i}(1 - G_{i'j})\right], \tag{D.28}$$

$$x_{ij}^{\Pi} = \mathbb{1}(i \text{ is matched to } j \text{ according to } \Pi) \tag{D.29}$$

$$N_{ij}^{\Pi} = P_{ij}(1 - x_{ij}^{\Pi}) + (1 - G_{ij})x_{ij}^{\Pi} \tag{D.30}$$

Intuitively, G_{ij} is one when (i,j) is a "Good" pair, with high benefit and low cost. P_{ij} is one when (i,j) is a "Perfect" pair, as each is incident to only one good pair, which is each other. Note that

$$\mathbb{E}[P_{ij}] = \frac{1}{n}\left(1 - \frac{1}{n}\right)^{2(n-1)} = \frac{1}{n}\Big/\left(1 + \frac{1}{n-1}\right)^{2(n-1)} \geq \frac{e^{-2}}{n}. \tag{D.31}$$

Now, N_{ij}^{Π} should be zero whenever the protocol returns an ϵ-stable outcome, since $N_{ij}^{\Pi} = 1$ implies that the protocol either did not match the perfect pair (i,j), or it matched a pair (i,j) that has negative net surplus. In other words, whenever the protocol succeeds in obtaining an ϵ-stable outcome, it must be that $\sum_{ij} N_{ij}^{\Pi} = 0$.

Moreover, in any realization, it must be that $\sum_{ij} N_{ij}^{\Pi} \leq 2n$, because there can be at most n perfect pairs and n matched pairs at the same time. Since Π fails to return a stable outcome with probability at most 0.01,

$$\mathbb{E}\left[\sum_{i,j\in[n]} N_{ij}^{\Pi}\right] \leq 0.02n. \tag{D.32}$$

Define a low-error pair (i,j) to be one such that $\mathbb{E}[N_{ij}^{\Pi}] \leq 0.04/n$, and denote the set of low-error pairs by S. By (D.32), $|S| \geq n^2/2$. The lower bound on per-agent communication cost is based on the following information theoretic argument:

$$L \geq H(\Pi) \geq \mathcal{I}(b,c;\Pi) \geq \sum_{i,j\in[n]} \mathcal{I}(b_{ij},c_{ij};\Pi) \geq \sum_{(i,j)\in S} \mathcal{I}(b_{ij},c_{ij};\Pi), \tag{D.33}$$

where $H(\cdot)$ denotes the entropy and $\mathcal{I}(\cdot;\cdot)$ the mutual information. The first inequality follows from the fact that if the maximum length of the protocol is L, the number of bit sequences is at most 2^L, so the entropy is at most $\log_2(2^L) = L$. The second inequality comes from the identity $\mathcal{I}(b,c;\Pi) = H(\Pi) - H(\Pi|b,c)$, and $H(\Pi|b,c) \geq 0$. The third inequality comes from the tuple (b_{ij}, c_{ij}) being independently drawn for each $(i,j) \in I \times J$, and the fact that $\mathcal{I}(X,Y;Z) \geq \mathcal{I}(X;Z) + \mathcal{I}(Y;Z)$ if X and Y are independent. The last inequality follows from the fact that mutual information is always non-negative. For a review of these facts from information theory, see Chapter 1 of Cover and Thomas (2012).

Suppose that we obtain a lower bound Z on the mutual information $\mathcal{I}(b_{ij}, c_{ij};\Pi)$ for any low-error pair $(i,j) \in S$, then combining (D.33) with the fact that $|S| \geq n^2/2$ yields the following lower bound on the per-agent communication cost,

$$\frac{L}{n} \geq \frac{n}{2} Z. \tag{D.34}$$

The desired lower bound Z is given by the following lemma. Part a) of Proposition 12 follows from part a) of Lemma D.2. Part b) of Proposition 12 follows from part b) of the lemma. □

LEMMA D.2 **(Lower bound on mutual information)**. *Let $(i,j) \in S$ be any low-error pair. Assume $n \geq 64$.*

a) If $d^J = 1$ and $d^I = 1/n$, then

$$\mathcal{I}(b_{ij}, c_{ij};\Pi) > \frac{.05\log_2 n}{n}. \tag{D.35}$$

b) If $d^J \in [1/\sqrt{n}, 1/2]$ and $d^I = 1/(nd^J)$, then

$$\mathcal{I}(b_{ij}, c_{ij};\Pi) > \frac{.0004}{nd^J}. \tag{D.36}$$

Proof of Lemma D.2 a). We use the following facts about mutual information. See Chapter 2 of Cover and Thomas (2012) for derivations.

- Symmetry: $\mathcal{I}(X;Y) = \mathcal{I}(Y;X)$.
- Data processing inequality: $\mathcal{I}(X;Y) \geq \mathcal{I}(X; f(Y))$ for any function $f(\cdot)$.
- Relationship to KL-divergence: $\mathcal{I}(X;Y) = \mathbb{E}_X\left[D(Y|_X \| Y)\right]$, where $D(\cdot\|\cdot)$ is the KL-divergence in bits. For discrete random variables X and Y with the same support and probability mass functions $p(\cdot)$ and $q(\cdot)$ respectively, $D(X\|Y) := \mathbb{E}_X[\log_2(p(X)/q(X))]$.

- Non-negativity of KL-divergence: $D(X\|Y) \geq 0$.

Applying the above, we have

$$\mathcal{I}(b_{ij}, c_{ij}; \Pi) \geq \mathcal{I}(b_{ij}; x_{ij}^{\Pi}) \geq \mathbb{E}[x_{ij}^{\Pi}] D\left(b_{ij}|_{x_{ij}^{\Pi}=1} \| b_{ij}\right) \tag{D.37}$$

To avoid clutter, omit the superscript Π on x and N for the rest of this proof, as the dependence of these random variables on the transcript Π is assumed. Now,

$$\mathbb{P}(b_{ij} = .99|x_{ij} = 1) = \frac{\mathbb{E}[G_{ij}x_{ij}]}{\mathbb{E}[G_{ij}x_{ij}] + \mathbb{E}[(1-G_{ij})x_{ij}]} \geq \frac{\mathbb{E}[G_{ij}x_{ij}]}{\mathbb{E}[G_{ij}x_{ij}] + .04/n} \geq \frac{e^{-2}/n - .04/n}{e^{-2}/n} > .7. \tag{D.38}$$

The first inequality follows from $\mathbb{E}[(1-G_{ij})x_{ij}] \leq \mathbb{E}[N_{ij}] \leq 0.04/n$ since (i,j) is a low-error pair. The second inequality follows from $\mathbb{E}[G_{ij}x_{ij}] \geq \mathbb{E}[P_{ij}] - \mathbb{E}[N_{ij}]$ and (D.31). (D.38) implies that

$$D(b_{ij}|_{x_{ij}=1} \| b_{ij}) \geq D(Bernoulli(.7) \| Bernoulli(1/n)) \tag{D.39}$$

$$= .7\log_2(.7) + .3\log_2(.3) + .7\log_2(n) + .3\log_2\left(\frac{n}{n-1}\right) \tag{D.40}$$

$$> .55\log_2 n \tag{D.41}$$

The last inequality follows from $.7\log_2(.7) + .3\log_2(.3) + .15\log_2(n) > 0$ for all $n \geq 64$. Now, since $(i,j) \in S$,

$$\mathbb{E}[x_{ij}] \geq \mathbb{E}[P_{ij}] - \mathbb{E}[N_{ij}] \geq \frac{e^{-2} - .04}{n}. \tag{D.42}$$

Plugging (D.42) and (D.41) into (D.37), we obtain the desired result for part a). □

Proof of Lemma D.2 b). Define

$$\theta_1 := \mathbb{P}((b_{ij}, c_{ij}) = (.99, 1)) \tag{D.43}$$

$$\theta_2 := \mathbb{P}((b_{ij}, c_{ij}) = (0, .01)) \tag{D.44}$$

$$\theta_3 := \mathbb{P}((b_{ij}, c_{ij}) = (0, 1)) \tag{D.45}$$

Note that under the assumptions of Lemma D.2, $n\theta_3 \geq 10$.

Without using the independence of b_{ij} and c_{ij}, we will show that

$$\mathcal{I}(b_{ij}, c_{ij}; \Pi) > \frac{0.006}{\left(\sqrt{1/\theta_1} + \sqrt{1/\theta_2} + 2\sqrt{1/\theta_3}\right)^2}. \tag{D.46}$$

The bound implies the desired result if we plug in $\theta_1 = d^I(1-d^J)$, $\theta_2 = (1-d^I)d^J$, $\theta_3 = (1-d^I)(1-d^J)$. To see this, note that the assumptions of the lemma implies that $d^I \leq d^J \leq 1/2$ and $d^I \leq 1/\sqrt{n} \leq 1/8$, so $1/(1-d^J) \leq 2$ and $1/(1-d^I) \leq 8/7$. Thus,

$$\left(\sqrt{\frac{1}{\theta_1}} + \sqrt{\frac{1}{\theta_2}} + 2\sqrt{\frac{1}{\theta_3}}\right)^2 \leq \left(\sqrt{\frac{2}{d^I}} + \sqrt{\frac{8/7}{d^I}} + 2\sqrt{\frac{2(8/7)}{8d^I}}\right)^2 < \frac{13}{d^I} = 13nd^J. \tag{D.47}$$

Plugging (D.47) into (D.46), we obtain the desired bound (D.36).

To show (D.46), define π to be the distribution of Π, and define the following conditional distributions.

- π_{11}: the distribution of Π conditional on $b_{ij} = .99$ and $c_{ij} = .01$;
- π_{10}: Π conditional on $b_{ij} = .99$ and $c_{ij} = 1$;
- π_{01}: Π conditional on $b_{ij} = 0$ and $c_{ij} = .01$;

- π_{00}: Π conditional on $b_{ij}=0$ and $c_{ij}=1$.

Define $\mu=\mathcal{I}(b_{ij},c_{ij};\Pi)$. Using the relationship between mutual information and KL-divergence,

$$\mu=\frac{1}{n}D(\pi_{11}\|\pi)+\theta_1 D(\pi_{10}\|\pi)+\theta_2 D(\pi_{01}\|\pi)+\theta_3 D(\pi_{00}\|\pi). \tag{D.48}$$

Now, KL-divergence is always non-negative. Moreover, it is related to the total variation distance by Pinsker's inequality $\|P-Q\|_1\le\sqrt{2(\ln 2)D(P\|Q)}$. Hence,

$$\|\pi_{10}-\pi_{00}\|_1\le\|\pi_{10}-\pi\|_1+\|\pi_{00}-\pi\|_1\le\sqrt{2\ln 2}(\sqrt{\mu/\theta_1}+\sqrt{\mu/\theta_3}). \tag{D.49}$$

Similarly,

$$\|\pi_{01}-\pi_{00}\|_1\le\|\pi_{10}-\pi\|_1+\|\pi_{00}-\pi\|_1\le\sqrt{2\ln 2}(\sqrt{\mu/\theta_2}+\sqrt{\mu/\theta_3}). \tag{D.50}$$

However, since $(i,j)\in S$ and $n\theta_3\ge 10$, we have

$$\mathbb{P}\left(x_{ij}^{\Pi}=1|(b_{ij},c_{ij})=(.99,.01)\right)\ge\frac{\mathbb{E}[P_{ij}]-0.04/n}{1/n}\ge e^{-2}-0.04>0.0953, \tag{D.51}$$

$$\mathbb{P}\left(x_{ij}^{\Pi}=1|(b_{ij},c_{ij})=(0,1)\right)\le\frac{0.04/n}{\theta_3}\le 0.004. \tag{D.52}$$

Thus, the total variation distance

$$\frac{1}{2}\|\pi_{11}-\pi_{00}\|_1>0.0953-0.004=0.0913. \tag{D.53}$$

The desired bound (D.46) follows from combining (D.49), (D.50), (D.53) with the following lemma, which is a result taken from the analysis of the disjointness problem from Mark Braverman's COS597D lecture notes on communication complexity. ☐

Lemma D.3 **(Consequence of the Rectangle Property of Communication Protocols)**.

$$\frac{1}{2}\|\pi_{11}-\pi_{00}\|_1\le\|\pi_{10}-\pi_{00}\|_1+\|\pi_{01}-\pi_{00}\|_1. \tag{D.54}$$

Proof of Lemma D.3. Let $\pi_{11}(z)$ denote the probability that given the inputs $b_{ij}=.99$ and $c_{ij}=.01$, the protocol has transcript z, where the randomness comes from the other inputs. Similarly define $\pi_{xy}(z)$ for $x\in\{0,1\}$, and $y\in\{0,1\}$. A well-known fact that holds for any communication protocol is that $\pi_{xy}(z)=P_x(z)Q_y(z)$ for some functions $P_x(\cdot)$ and $Q_y(\cdot)$ which only depend on the subscripted input bit and z. This follows from the rectangle property of deterministic protocols (Kushilevitz and Nisan 2006). A proof is found in Lemma 6.7 of Bar-Yossef et al. (2004).

Now, suppose that $\pi_{00}(z)\ge\pi_{11}(z)$. Then it must be that either $P_0(z)\ge P_1(z)$ or $Q_0(z)\ge Q_1(z)$. If $P_0(z)\ge P_1(z)$, then

$$Q_0(z)|P_0(z)-P_1(z)|+P_0(z)|Q_0(z)-Q_1(z)|\ge Q_0(z)(P_0(z)-P_1(z))+P_1(z)(Q_0(z)-Q_1(z)). \tag{D.55}$$

If $Q_0(z)\ge Q_1(z)$, then

$$Q_0(z)|P_0(z)-P_1(z)|+P_0(z)|Q_0(z)-Q_1(z)|\ge Q_1(z)(P_0(z)-P_1(z))+P_0(z)(Q_0(z)-Q_1(z)). \tag{D.56}$$

In either case, the right hand side of the above is equal to $P_0(z)Q_0(z)-P_1(z)Q_1(z)$, so we have

$$|\pi_{00}(z)-\pi_{10}(z)|+|\pi_{00}(z)-\pi_{01}(z)|\ge\pi_{00}(z)-\pi_{11}(z) \tag{D.57}$$

Summing over all transcripts z for which $\pi_{00}(z)\ge\pi_{11}(z)$, we get the desired inequality (D.54). ☐

If we plug (D.46) directly into (D.34), then we obtain the following result, which implies that efficient matchmaking is not possible if we allow the tuple (b_{ij}, c_{ij}) to be adversarially correlated, so that the partners whom an agent likes tend to be the ones who do not reciprocate interest. We do not expect such adversarial correlation to cause problems in practice, because most correlations either exhibit a predictable structure as given in the example in Section 5.1, or are related with observable characteristics, which makes it possible for the platform to capture them using a well-designed questionnaire.

PROPOSITION D.1 **(Exact lower bound under correlated preferences)**. *Consider a market with $|I| = |J| = n$. For each $(i,j) \in I \times J$, the tuple (b_{ij}, c_{ij}) is drawn independently from the following joint distribution:*

$$(b_{ij}, c_{ij}) = \begin{cases} (.99, .01) & \text{with probability } 1/m, \\ (.99, 1) & \text{with probability } \theta_1, \\ (0, .01) & \text{with probability } \theta_2, \\ (0, 1) & \text{with probability } \theta_3, \end{cases} \tag{D.58}$$

where $\theta_1 + \theta_2 + \theta_3 = 1 - 1/n$ and $n\theta_3 \geq 10$. For any $\epsilon < 0.49$, any matchmaking protocol that yields an ϵ-stable outcome for the above market with at least 99% probability requires a per-agent communication cost of at least

$$\frac{0.003n}{\left(\sqrt{1/\theta_1} + \sqrt{1/\theta_2} + 2\sqrt{1/\theta_3}\right)^2}. \tag{D.59}$$

Observe that this lower bound is $\Omega(n)$ if θ_1, θ_2 and θ_3 are lower bounded by a constant.

D.6. Proof of Proposition 12 c)

Given $\epsilon < .49$, $n \geq 64$ and a one-round protocol that yields an ϵ-stable outcome with at least 99% probability in the market $M(n, .5, .5)$. Define Π to be the transcript of the protocol, which is the sequence of all bits communicated before termination. Let L be the maximum length of the transcript. The per-agent communication cost is equal to L/n. Since the protocol is a one-round protocol, we can write $\Pi = (\beta, \gamma)$, where β_i is the message of customer $i \in I$ and depends only on $(b_{ij}; j \in J)$, and γ_j is the message of provider $j \in J$ and depends only on $(c_{ij}; i \in I)$. Without knowing the preference realizations, β can be viewed as a random variable correlated with b but independent of γ and c; γ is a random variable correlated with c but independent of β and b. Each transcript Π leads to a transaction outcome, which we denote as (x^Π, p^Π).

For each realization of β, define

$$G_i(\beta) := \{j : \mathbb{P}(b_{ij} = .99|\beta) \geq .75\}, \tag{D.60}$$

$$U(\beta, c) := \{i : c_{ij} = 1 \text{ for all } j \in G_i(\beta)\}, \tag{D.61}$$

$$m(\beta) := \sum_{i=1}^{n} |G_i(\beta)|. \tag{D.62}$$

Intuitively, $G_i(\beta)$ is the set of providers which we can infer with at least .75 confidence that customer i likes given her message β_i. Given provider preference realizations c, customer i is in the set $U(\beta, c)$ if none of the providers in $G_i(\beta)$ likes her back. We refer to such a customer as being *unreciprocated*.

Define $z := n(\log_2 n)/4$. We will show that the total cardinality of the sets $G_i(\beta)$ satisfying the following probabilistic bound:

$$\mathbb{P}_\beta\Big(m(\beta) \geq z\Big) > .744. \tag{D.63}$$

From this, we can bound the entropy of the random variable β as follows,

$$H(\beta) \geq \mathcal{I}(b;\beta) \geq \sum_{i,j\in[n]} \mathcal{I}(b_{ij};\beta) \tag{D.64}$$

$$= \sum_{i,j\in[n]} \mathbb{E}_\beta\Big[D\left(b_{ij}|_\beta \| b_{ij}\right)\Big]. \tag{D.65}$$

$$\geq \sum_{i,j\in[n]} \mathbb{P}_\beta\Big(m(\beta)\geq z\Big)\mathbb{E}_\beta\Big[D\left(b_{ij}|_\beta \| b_{ij}\right)\Big| m(\beta)\geq z\Big] \tag{D.66}$$

$$> (.744)(.1887)z > .035n\log_2(n) \tag{D.67}$$

Here, the first inequality is due to the fact that $\mathcal{I}(b;\beta)\equiv H(\beta)-H(\beta|b)$ and $H(\beta|b)\geq 0$. The second inequality is due to the independence of b_{ij} for each $(i,j)\in I\times J$, and the fact that $\mathcal{I}(X,Y;Z)\geq \mathcal{I}(X;Z)+\mathcal{I}(Y;Z)$ whenever X and Y are independent. The equality in (D.65) is based on the relationship between mutual information and the KL-divergence in bits, as in the proof of Lemma D.2. Inequality (D.66) is due to the fact that KL-divergence is always non-negative. Inequality (D.67) is based on (D.60) and the fact that for any probability $y\geq .75$, $D(Bernoulli(y)\|Bernoulli(.5))\geq D(Bernoulli(.75)\|Bernoulli(.5))>.1887$.

By symmetry, the same argument can be used to show $H(\gamma)>.035\log_2(n)$, from which it follows that the per-agent communication cost satisfies the following lower bound:

$$\frac{L}{n}\geq \frac{H(\beta)+H(\gamma)}{n} > .07\log_2(n). \tag{D.68}$$

To complete the proof, it suffices to show (D.63). Observe that for any β such that $m(\beta)<z$, we have by Jensen's inequality,

$$\mathbb{E}_c\Big[|U(\beta,c)|\Big] = \sum_{i=1}^n 2^{-|G_i(\beta)|} \geq \sum_{i=1}^n 2^{-m(\beta)/n} > n2^{-z/n} = n^{3/4}. \tag{D.69}$$

In other words, conditional on $m(\beta)<z$, the expected number of unreciprocated customers is at least $n^{3/4}$.

Define $k:=.75n^{3/4}$. Since $|U(\beta,c)|$ can be expressed as the sum of independent binary random variables, a standard Chernoff bound implies that for any β such that $m(\beta)<z$,

$$\mathbb{P}_c\Big(|U(\beta,c)|>k\Big)\geq 1-\exp\left(-\frac{(.25)^2n^{3/4}}{2}\right) > .5 \qquad \text{(for any } n\geq 64). \tag{D.70}$$

Define the following events:

- N: (x^Π,p^Π) is not ϵ-stable.
- E: $m(\beta)<z$ and $|U(\beta,c)|>k$.
- E_1: E and there exists (i,j) such that $x_{ij}^{(\beta,\gamma)}=1$ but $j\notin G_i(\beta)$.
- E_2: $E\backslash E_1$.
- E_3: there exists an ϵ-stable matching in which more than k customers are unmatched.

The event E corresponds to $m(\beta)<z$ and there being many unreciprocated customers. By (D.70),

$$P(E) > .5\mathbb{P}_\beta\Big(m(\beta)<z\Big). \tag{D.71}$$

The event E_1 corresponds to the subset of E in which the protocol matches a customer i to a provider j for whom $\mathbb{P}(b_{ij}=0|\beta)>.25$. This is a risky move because with probability at least .25, she does not like the provider. This implies that $P(N|E_1)>.25$.

The event E_2 corresponds to avoiding such risky moves and only matching customer i to a provider j if $j \in G_i(\beta)$. However, this implies that none of the unreciprocated customers $i \in U(\beta, c)$ can be matched. This implies that $(E_2 \backslash N) \subseteq E_3$. However, a straightforward union bound implies that

$$\mathbb{P}(E_3) \leq \left(\frac{en}{k}\right)^{2k} \left(\frac{3}{4}\right)^{k^2} < .022, \tag{D.72}$$

where the last inequality holds for all $n \geq 64$ and $k = .75n^{3/4}$.

Since the protocol yields an ϵ-stable outcome with probability at least 99%, we have

$$.01 \geq \mathbb{P}(N) \geq P(N \cap E_1) + P(N \cap E_2) \geq .25P(E_1) + (P(E_2) - .022) \tag{D.73}$$

So $P(E) \leq 4(.01 + .022) = .128$. By (D.71),

$$\mathbb{P}_\beta\Big(m(\beta) < z\Big) < \frac{\mathbb{P}(E)}{.5} \leq .256, \tag{D.74}$$

which implies (D.63), as desired. □

Appendix E: Modifications of Protocols 1 and 3 under correlated provider costs

As discussed in Section 5.1, the performance guarantees for Protocols 1 and 3 derived in Sections 2 and 3 already allow arbitrary correlation with the matrix b of customer benefits. This section shows how to modify the protocols so as to obtain similar guarantees when provider costs are also correlated, according to the following structure:

$$c_{ij} = \alpha_i + \beta_j + \gamma_{ij} \qquad \text{for all } (i,j) \in I \times J, \tag{E.1}$$

where α_i and β_j are unobserved fixed effects for each customer i and provider j, and γ_{ij} is an unobserved idiosyncratic term that is assumed to be independently drawn from a commonly known distribution Γ_{ij}. The benefit matrix b is allowed to be arbitrarily correlated within itself and with (α, β), but the γ_{ij}'s are assumed to be independent of (b, α, β). Proposition D.1 implies that the independence of γ_{ij}'s and b cannot be dropped, as allowing adversarial correlations between benefits and costs would make it impossible to achieve ϵ-stable outcomes with sub-linear per-agent communication cost.

The high level idea is to run the following protocol first in order to estimate the α_i's and β_j's, and use these estimates to calibrate the initial cost estimates $\hat{c}_{ij}$'s in Protocols 1 and 3. Without loss of generality, assume that $I = [n] = \{1, 2, \cdots, n\}$ and $J = [n]$.

Protocol 5 (Cost estimate calibration) *Given an integer $m > 0$ and the means of the distributions Γ_{ij}'s. Elicit the cost c_{ij} from every provider $j \in J$ for every customer $i \in [m]$, as well as from every provider $j \in [m]$ for every customer $i \in I$. Define for each $(i,j) \in I \times J$,*

$$\hat{e}_{ij} := \frac{1}{m} \sum_{k=1}^{m} \left(c_{ik} + c_{kj} - c_{kk} - \mathbb{E}[\gamma_{ik}] - \mathbb{E}[\gamma_{kj}] + \mathbb{E}[\gamma_{kk}]\right). \tag{E.2}$$

Return the matrix $\hat{e}$.

Proposition E.1. *After running Protocol 5, with probability at least $1 - 2n^2 e^{-m\epsilon^2/24}$,*

$$|\alpha_i + \beta_j - \hat{e}_{ij}| \leq \frac{\epsilon}{4} \qquad \textit{for all } (i,j) \in I \times J. \tag{E.3}$$

Proof of Proposition 5. Define $y_{ij}^k = c_{ik} + c_{kj} - c_{kk} - \mathbb{E}[\gamma_{ik}] - \mathbb{E}[\gamma_{kj}] - \mathbb{E}[\gamma_{kk}] - \alpha_i - \beta_j$. Note that $y_{ij}^k = (\gamma_{ik} - \mathbb{E}[\gamma_{ik}) + (\gamma_{kj} - \mathbb{E}[\gamma_{kj}]) - (\gamma_{kk} - \mathbb{E}[\gamma_{kk}])$. Thus, y_{ij}^k is the sum of three independent random variables, each of mean 0 and the support of each has length of at most 1. Therefore, by Hoeffding's inequality,

$$\mathbb{P}\left(|\hat{e}_{ij}| > \frac{\epsilon}{4}\right) = \mathbb{P}\left(\left|\frac{1}{3m}\sum_{k=1}^{m} y_{ij}^k\right| > \frac{\epsilon}{12}\right) \leq 2\exp\left(-6m\left(\frac{\epsilon}{12}\right)^2\right) = 2e^{-m\epsilon^2/24}. \tag{E.4}$$

A union bound over the n^2 pairs $(i,j) \in I \times J$ yields the desired result. □

The modified Protocol 1 is as follows. In Phase 1, we run Protocol 5 with

$$m = \left\lceil \frac{72\ln(n)}{\epsilon^2} \right\rceil. \tag{E.5}$$

In Phase 2, we run Protocol 1 as defined in Section 2.4 but modify the initial cost estimates to

$$\hat{c}_{ij} = \left(\hat{e} - \frac{\epsilon}{4}\right) + \underline{\gamma}_{ij} + \epsilon, \tag{E.6}$$

where $\underline{\gamma}_{ij}$ is the lower support of γ_{ij}. The following result is similar to Proposition 4.

Proposition E.2 **(Performance of Protocol 1 under correlated costs)**. *When provider costs are correlated according to* (E.1)*, the modified Protocol 1 described above returns an ϵ-stable outcome with probability at least $1-2/n$. Moreover, with probability at least $1-2/n-e^{-n/2}$, the number of provider interactions per customer in Phase 2 is no more than*

$$2\left\lceil\frac{1}{\epsilon}\right\rceil\frac{1}{\tilde{d}^J},$$

where $\tilde{d}^J$ is the largest value such that

$$\mathbb{P}\left(\gamma_{ij} \leq \underline{\gamma}_{ij} + \frac{\epsilon}{2}\right) \geq \tilde{d}^J \qquad \textit{for all } (i,j) \in I \times J. \tag{E.7}$$

Proof of Proposition E.2. By Proposition 5, with the value of m given in (E.5), the guarantee in (E.3) holds with probability at least $1-2/n$. When (E.3) holds, $\hat{c}_{ij} \leq c_{ij} + \epsilon$ for all (i,j), so the proof of Proposition 2 implies that the above protocol terminates with an ϵ-stable outcome.

Moreover, when (E.3) holds, (E.6) implies that

$$\gamma_{ij} \leq \underline{\gamma}_{ij} + \frac{\epsilon}{2} \implies c_{ij} \leq \hat{c}_{ij}. \tag{E.8}$$

Therefore, the argument in the proof of Proposition 4 implies the desired bound on the number of interactions in Phase 2. □

Similarly, the modified Protocol 3 is defined as follows. In Phase 1, we run Protocol 1 with m given by (E.5) as above. In Phase 2, we run Protocol 3 as defined in Section 2.4 with parameter

$$z_{ij} = \left(\hat{e}_{ij} - \frac{\epsilon}{4}\right) + g_{ij}, \tag{E.9}$$

where the g_{ij}'s are parameters that depend only on commonly known information. The following result is similar to Propositions 6 and 7.

Proposition E.3 **(Performance of Protocol 3 under correlated costs)**. *The modified Protocol 3 described above always terminates with an ϵ-stable outcome. Given any choice of g_{ij}'s, the expected number of interactions in Phase 2 is at most*

$$\left(\sum_{i\in I, j\in J}\mathbb{P}(\gamma_{ij}<g_{ij})\right)+\left\lceil\frac{1}{\epsilon}\right\rceil\left(\sum_{j\in J}\frac{1}{\min_{i\in I}\{\mathbb{P}(\gamma_{ij}\le g_{ij}+\frac{\epsilon}{2}|\gamma_{ij}\ge g_{ij}-\frac{\epsilon}{2})\}}\right)+\frac{2}{n}. \tag{E.10}$$

If the parameters g_{ij}'s are chosen as follows,

$$g^*_{ij}:=\inf\{y\in[\underline{\gamma}_{ij},1]:\mathbb{P}(\gamma_{ij}\in[y,y+\epsilon/2])\ge 1/\sqrt{2n}\}, \tag{E.11}$$

then with probability at least $1-2/n-2e^{-n/2}$, the number of interactions per customer in Phase 2 is no more than

$$4\left\lceil\frac{1}{\epsilon}\right\rceil\sqrt{2n}. \tag{E.12}$$

Proof of Proposition E.3. The fact that the above protocol always terminates with an ϵ-stable outcome follows from Proposition 3, which does not require any assumptions on the preferences of agents. To see the bound in (E.10), suppose first that (E.3) holds, then the first term upper bounds the expected number of interactions in Step 2 of Phase 2, as under (E.3), $c_{ij}<z_{ij}$ implies that $\gamma_{ij}<g_{ij}$; the second term upper bounds the expected number of interactions in Step 3 of Phase 2, as under (E.3),

$$\gamma_{ij}\le g_{ij}+\frac{\epsilon}{2}\implies c_{ij}\le z_{ij}+\epsilon, \tag{E.13}$$

$$\text{and}\qquad c_{ij}\ge z_{ij}\implies \gamma_{ij}\ge g_{ij}-\frac{\epsilon}{2}. \tag{E.14}$$

Therefore, for every provider j, if $\tilde{d}_j=\min_{i\in I}\{\mathbb{P}(\gamma_{ij}\le g_{ij}+\epsilon/2|\gamma_{ij}\ge g_i-\epsilon/2\}$, then the number of interactions in Step 3 of Phase 2 until $\hat{v}_j$ increments once is stochastically dominated by a geometry random variable of mean $1/d_j$. Now, if (E.3) is violated for a certain pair (i,j), then the above guarantees may not apply, but the maximum number of additional interaction is one, as the true cost c_{ij} would be known after one interaction between customer i and provider j. Therefore, since the probability that (E.3) is violated for each (i,j) is at most $2/n^3$ and there are n^2 pairs, the expected total number of additional interactions due to these violations is at most $2/n$.

Finally, to show the bound in (E.12), observe under (E.3), for each (i,j),

$$\mathbb{P}(c_{ij}<z_{ij})\le\frac{2}{\epsilon}\frac{1}{\sqrt{2n}}. \tag{E.15}$$

Moreover, in Step 3 of Phase 2, the provider j answers "yes" to customer i whenever $\gamma_{ij}\in[\underline{\gamma}_{ij},\underline{\gamma}_{ij}+\epsilon/2]$, which occurs with probability at least $1/\sqrt{2n}$ by (E.11), independent of everything else. The desired bound in (E.12) follows from the same argument as in the proof of Proposition 9, except that we re-define the intermediate quantities $r_1:=((\sqrt{2}/\epsilon)+0.5)n^{1.5}$ and $r_2:=2\sqrt{2}\lceil 1/\epsilon\rceil n^{1.5}$. □

Appendix F: Definition of ϵ-stable outcomes under many-to-one matching

Suppose that each provider j has the capacity to serve q_j customers, we define the corresponding notion of an ϵ-stable outcome below. In this section, we do not assume that $|I| = |J|$.

A matching x is *feasible* if

$$\sum_{j\in J} x_{ij} \le 1 \qquad \text{for each } i \in I, \tag{F.1}$$

$$\sum_{i\in I} x_{ij} \le q_j \qquad \text{for each } j \in J, \tag{F.2}$$

$$x_{ij} \in \{0,1\} \qquad \text{for each } (i,j) \in I \times J. \tag{F.3}$$

A tuple (x,p) is an *individually rational transaction outcome* if x is a feasible matching, and $p \in \mathbb{R}^I$ satisfies

$$c_{ij} \le p_i \le b_{ij} \qquad \text{whenever } x_{ij} = 1. \tag{F.4}$$

The tuple (x,p) is ϵ*-stable* if in addition to being individually rational,

$$u_i + v_j + 2\epsilon \ge b_{ij} - c_{ij} \qquad \text{for each } (i,j) \in I \times J, \tag{F.5}$$

$$\text{where} \quad u_i := \begin{cases} 0 & \text{if } \sum_{j\in J} x_{ij} = 0, \\ b_{ij} - p_i & \text{if } x_{ij} = 1; \end{cases} \tag{F.6}$$

$$v_j := \begin{cases} 0 & \text{if } \sum_{i\in I} x_{ij} < q_j, \\ \min_{i:x_{ij}=1}\{p_i - c_{ij}\} & \text{otherwise.} \end{cases} \tag{F.7}$$

Proposition F.1. *Given any $\epsilon \ge 0$, an ϵ-stable transaction outcome (x,p) always exists. Moreover, any ϵ-stable transaction outcome (x,p) approximately maximizes total market surplus, in the sense that*

$$\sum_{i\in I, j\in J} (b_{ij} - c_{ij}) x_{ij} \ge \left(\sum_{i\in I, j\in J} x'_{ij} \right) - 2\epsilon |I| \qquad \textit{for any feasible matching } x'. \tag{F.8}$$

Proof of Proposition F.1. The maximum total surplus W^* is the optimal objective of the LP with objective

$$\text{Maximize:} \quad \sum_{i\in I, j\in J} (b_{ij} - c_{ij}) x_{ij}, \tag{F.9}$$

and constraints (F.1), (F.2), and $x \ge 0$. The dual LP is

$$\text{Minimize:} \quad \sum_{i\in I} \hat{u}_i + \sum_{j\in J} q_j \hat{v}_j \tag{F.10}$$

$$\text{s.t.} \quad \hat{u}_i + \hat{v}_j \ge b_{ij} - c_{ij} \qquad \text{for all } (i,j) \in I \times J, \tag{F.11}$$

$$\hat{u}, \hat{v} \ge 0. \tag{F.12}$$

Given a set of optimal primal and dual solutions $(x, \hat{u}, \hat{v})$. Assume that $x_{ij} \in \{0,1\}$. This is without loss of generality since the constraint set to the primal is totally unimodular. Define

$$p_i := \begin{cases} 0 & \text{if } \sum_{j\in J} x_{ij} = 0, \\ b_{ij} - \hat{u}_i & \text{if } x_{ij} = 1. \end{cases} \tag{F.13}$$

By complementary slackness, $x_{ij} = 1$ implies that $\hat{u}_i + \hat{v}_j = b_{ij} - c_{ij}$. Hence, the transaction outcome (x, p) leads to $u = \hat{u}$ and $v = \hat{v}$ under (F.6) and (F.7), which implies that (x, p) is exactly stable by (F.11) and (F.12). This transaction outcome is also ϵ-stable for any $\epsilon \geq 0$.

Now, given any ϵ-stable transaction outcome (x, p). Note that if u and v are defined according to (F.6) and (F.7), then $x_{ij} = 1$ implies that $b_{ij} - c_{ij} \geq u_i + v_j$. Hence,

$$\sum_{i\in I, j\in J}(b_{ij} - c_{ij})x_{ij} \geq \sum_{i\in I, j\in J}(u_i + v_j)x_{ij} = \sum_{i\in I}(u_i + 2\epsilon) + \sum_{j\in J} q_j v_j - 2\epsilon|I| \geq W^* - 2\epsilon|I|, \tag{F.14}$$

where the last inequality follows from weak inequality, because setting $\hat{u} = u + 2\epsilon$ and $\hat{v} = v$ yields a feasible solution $(\hat{u}, \hat{v})$ to the dual LP. □.

Appendix G: Dynamic arrival of customers

This section shows that the main insights in the paper continue to hold under a natural alternative model with dynamic arrival of customers. Let I be a set of customer segments and $J := [n]$ be a set of providers. The identity $i \in I$ of the customer segment represents the information collected by the platform about the customer via questionnaires. Customers of each segment $i \in I$ arrives to the market according to a Poisson process of rate λ_i, and each customer seeks to be served by at most one provider, after which she departs from the market. Providers on the other hand remain perpetually in the market, and each provider $j \in J$ is associated with a capacity q_j, which represents the maximum rate at which he can serve customers in the long run. Each individual customer is associated with a type $\theta = (i, b, c)$, where $i \in I$ is her segment and $b, c \in [0,1]^J$, with b_j and c_j representing the benefit and cost associated with her potential transaction with provider j. When a customer of type $\theta = (i, b, c)$ enters the market, she knows her segment i and the benefit vector b, while provider j knows her segment i and the cost c_j. For each customer, the benefit b_j and cost c_j are independently drawn from commonly known distributions B_{ij} and C_{ij}, which can depend both on the customer segment i and the provider j. For technical convenience, assume that these preference distributions are all continuous.

Without loss of generality, normalize the total arrival rate of customers to one, so $\sum_{i\in I} \lambda_i = 1$. Label the providers as $J = [n] = \{1, 2, \cdots, n\}$. The cardinality of I can be arbitrary and does not necessarily equal n. Define Θ to be the support of customer types $\theta = (i, b, c)$, and define F to be a probability distribution on Θ where the segment i is distributed according to the probability vector λ, and conditional on i, we have $b_j \sim B_{ij}$ and $c_j \in C_{ij}$, with each term being independent of everything else.

The type of matchmaking systems considered are those that periodically run a static matchmaking protocol every time m new customers arrive, where m is a parameter called the batch size. Any customer who enters the static protocol but is not matched by it leaves the market, so that the next time the static protocol is run, it is with m new customers. To maximize the generality of the analysis, the following definition allows the static protocol that is run each time to be randomly drawn from a known distribution,[6] and it does not always have to include every provider.

[6] In the communication complexity literature, a distribution over deterministic communication protocols is referred to as a "public coin randomized protocol" (Kushilevitz and Nisan 2006).

Definition G.1 (Dynamic matchmaking protocol). A *dynamic matchmaking protocol* with batch size m is a distribution over communication protocols, each of which has the following form: the set of agents includes a set A of m newly arrived customers and a subset $J' \subseteq J$ of providers. Each customer $a \in A$ is associated with her own type (i_a, b_a, c_a) drawn from the distribution F. Here, $i_a \in I$ is her segment, b_{aj} is her benefit for being served by provider j, and c_{aj} is provider j's cost for serving her. The segment of each customer $a \in A$ is known by the protocol and all the agents. The private input of each customer $a \in A$ is $(b_{aj} : j \in J')$, while the private input of each provider $j \in J'$ is $(c_{aj} : a \in A)$. The output of the protocol is an individually rational transaction outcome (x, p) as defined in the many-to-one model in Appendix F, with agent sets A and J' and provider capacities set to infinity. (In other words, since the dynamic model only constrains provider capacities over the long run, providers can accommodate arbitrary surges in customer demand in the short run.) The *average per-customer communication cost* of the dynamic matchmaking protocol is the expected length of the communication protocol divided by n, with the expectation taken over the random draws of customer types and the random draw of the communication protocol itself.

Each dynamic matchmaking protocol induces a *generalized transaction outcome* $(x, \underline{p}, \overline{p})$ as follows. x is referred to a *generalized matching*, $\underline{p}$ as the *minimum payment function*, and $\overline{p}$ as the *maximum payment function*. For each $j \in J$, $x_j : \Theta \to [0,1]$ is a function that specifies the probability that each customer of type $\theta \in \Theta$ is matched to the provider j by the protocol; $\underline{p}_j : \Theta \to \mathbb{R}$ is a function that specifies the minimum payment when a type θ customer is matched to provider j, and $\overline{p}_j : \Theta \to \mathbb{R}$ is a function that specifies the maximum payment. The generalized matching x is said to be *feasible* if

$$\sum_{j \in J} x_j(\theta) \le 1 \qquad \text{for all } \theta \in \Theta, \tag{G.1}$$

$$\mathbb{E}_{\theta \sim F}[x_j(\theta)] \le q_j \qquad \text{for all } j \in J. \tag{G.2}$$

The generalized transaction outcome $(x, \underline{p}, \overline{p})$ is said to be *individually rational* if

$$c_j \le \underline{p}_j(i,b,c) \le \overline{p}_j(i,b,c) \le b_j \qquad \text{for all } j \in J \text{ and } (i,b,c) \in \Theta \text{ such that } x_j(i,b,c) > 0. \tag{G.3}$$

A generalized transaction outcome $(x, \underline{p}, \overline{p})$ is said to be ϵ-stable if in addition to being feasible, it satisfies

$$u(i,b,c) + v_j + 2\epsilon \ge b_j - c_j \qquad \text{for each } (i,b,c) \in \Theta \text{ and } j \in J, \tag{G.4}$$

$$\text{where} \qquad u(i,b,c) := \begin{cases} 0 & \text{if } \sum_{j \in J} x_j(i,b,c) < 1, \\ \inf\{b_j - \overline{p}_j(i,b,c) : x_j(i,b,c) > 0\} & \text{otherwise;} \end{cases} \tag{G.5}$$

$$v_j := \begin{cases} 0 & \text{if } \mathbb{E}_{\theta \sim F}[x_j(\theta)] < q_j, \\ \inf\{\underline{p}_j(i,b,c) - c_j : x_j(i,b,c) > 0\} & \text{otherwise.} \end{cases} \tag{G.6}$$

The existence of ϵ-stable generalized transaction outcomes is established in the next section. The following proposition is analogous to Proposition 1 and states that such an outcome induces near-optimal social welfare.

Proposition G.1. *If $(x, \underline{p}, \overline{p})$ is an ϵ-stable generalized transaction outcome, then the market surplus per customer is at most 2ϵ away from the maximum possible: if x' is another feasible generalized matching, then*

$$\mathbb{E}_{(i,b,c) \sim F}[(b-c) \cdot x(i,b,c)] \ge \mathbb{E}_{(i,b,c) \sim F}[(b-c) \cdot x'(i,b,c)] - 2\epsilon. \tag{G.7}$$

Proof of Proposition G.1. By the construction in (G.5) and (G.6), if $(x, \underline{p}, \overline{p})$ is ϵ-stable, then $x_j(i,b,c) > 0$ implies that $u(i,b,c) + v_j \le b_j - c_j$. Define $W = \mathbb{E}_{(i,b,c)\sim F}[u(i,b,c)] + \sum_{j\in J} q_j v_j$. We have,

$$\mathbb{E}_{(i,b,c)\sim F)}[(b-c)\cdot x(i,b,c)] \ge \mathbb{E}_{(i,b,c)\sim F}\left[\sum_{j\in J}(u(i,b,c)+v_j)x_j(i,b,c)\right] \ge W. \tag{G.8}$$

However, by (G.4), if x' is feasible, then

$$\mathbb{E}_{(i,b,c)\sim F}[(b-c)\cdot x'(i,b,c)] \le \mathbb{E}_{(i,b,c)\sim F}\left[\sum_{j\in J}(u(i,b,c)+v_j+2\epsilon)x'_j(i,b,c)\right] - 2\epsilon \le W - 2\epsilon. \tag{G.9}$$

Combining the above two lines yields the desired results. □

As in Section 3, define the preference density parameters d^I and d^J to be the largest values such that for every segment $i \in I$, if $b_j \sim B_{ij}$ and $c_j \sim C_{ij}$, then $\mathbb{P}(b_j \ge \overline{b}_{ij} - \epsilon) \ge d^I$ and $\mathbb{P}(c_j \le \underline{c}_{ij} + \epsilon) \ge d^J$, where $\overline{b}_{ij}$ is the upper support of B_{ij} and $\underline{c}_{ij}$ is the lower support of C_{ij}.

G.1. Protocol 1 (customers search) in the dynamic model

Given $\epsilon > 0$ and a provider surplus vector $\hat{v} \in [0,1]^J$ to be determined later, consider the following dynamic matchmaking protocol with batch size 1, so there is only one customer that needs to be considered at a time. Let her type be (i,b,c)

1. **Initialization:** Initialize $\hat{c}_j = \underline{c}_{ij} + \epsilon$. Define $U_j(\hat{v},\hat{c}) = b_j - \hat{v}_j - \hat{c}_j$ for every $j \in J$.
2. **Customer searches for provider:** If $U_j(\hat{v},\hat{c}) < 0$ for all $j \in J$, then the customer sends the message "0," at which point the protocol terminates without matching her. Otherwise, the customer communicates the identity of the provider $j^* \in \arg\max_j\{U_j(\hat{v},\hat{c})\}$ with the smallest index. The provider responds whether $c_j \le \hat{c}_j$.
   - If "yes," then the protocol terminates by matching the customer to j^* at payment $\hat{v}_{j^*} + \hat{c}_{j^*}$.
   - If "no," then provider j^* communicates $\Delta_j := \lceil (c_j - \hat{c}_j)/\epsilon \rceil$ and the protocol updates $\hat{c}_j = \hat{c}_j + \epsilon\Delta_j$. Repeat step 2.

Definition G.2 (Equilibrium provider surplus). Given ϵ, define an *equilibrium provider surplus vector* $\hat{v} \in [0,1]^J$ for the above protocol as one in which the generalized transaction outcome $(x, \underline{p}, \overline{p})$ induced by the protocol is such that x is a feasible generalized matching, and

$$\hat{v} > 0 \implies \mathbb{E}_{\theta\sim F}[x_j(\theta)] = q_j. \tag{G.10}$$

The following result is analogous to Propositions 2 and 4.

Proposition G.2 **(Performance guarantee of customers search)**. *Given any $\epsilon > 0$, an equilibrium provider surplus vector $\hat{v}$ of the above protocol exists. Under such a $\hat{v}$, the generalized transaction outcome induced by the above protocol is ϵ-stable, and the average per-agent communication cost is at most $O((\log n)/d^J)$.*

Note that the $\lceil 1/\epsilon \rceil$ factor from Proposition 4 disappears in the above performance guarantee, as such a factor is associated in with the search process over the vector $\hat{v}$ in the proof of Proposition 4. However, in the dynamic model, the vector $\hat{v}$ would converge in the long run to its market clearing value, and after convergence, the only uncertainty facing new customers is their cost vector c. Since the performance metric in the dynamic model is averaged over the long run, it suffices to analyze the performance after convergence, which is why the above proposition assumes an equilibrium surplus vector.

Proof of Proposition G.2. The existence of an equilibrium provider surplus vector $\hat{v}$ follows from Tarski's fixed-point theorem as follows. Given a fixed value of ϵ, for any $v \in [0,1]^n$, define $D_j(v)$ to be the mass of agents matched to provider j when the surplus vector is v. Define the operator $\Gamma : [0,1]^n \to [0,1]^n$,

$$\Gamma_j(v) := \inf\{y \in [0,1] : D_j(r(v,j,y)) \leq q_j\}, \tag{G.11}$$

where $r(v,j,y)$ is defined to be the vector v with the jth component replaced by the value y. Note that $\Gamma_j(v)$ is well defined since $D_j(r(v,j,1)) = 0$ for all $j \in J$. Moreover, $\Gamma_j(v)$ is a monotone operator, in the sense that if $v \leq v'$ component-wise, then $\Gamma(v) \leq \Gamma(v')$ component-wise. This is because increasing the surplus required by every provider $j' \neq j$ would only increase the demand for provider j. By Tarski's fixed point theorem, there exists a fixed point $\hat{v}$ such that $\Gamma(\hat{v}) = \hat{v}$. Moreover, since the preference distributions are continuous, $D_j(r(v,j,y))$ is a continuous function of y, so the fixed point satisfies (G.10). Such a fixed point $\hat{v}$ is an equilibrium surplus vector.

Now, given an equilibrium provider surplus vector $\hat{v}$, we show that the generalized transaction outcome $(x, \underline{p}, \overline{p})$ is ϵ-stable. Note that by construction, $(x, \underline{p}, \overline{p})$ is individually rational, and $\hat{v}_j \leq v_j$, where v is defined as in (G.6). Moreover, throughout the protocol, $\hat{c}_j \leq c_j + \epsilon$ for all $j \in J$.

Consider a type $\theta = (i,b,c)$ such that $\sum_{j\in J} x_j(\theta) < 1$, it must be that in some run of the protocol, $U_j(\hat{v},\hat{c}) < 0$ for all $j \in J$, which implies that $b_j - c_j - v_j \leq b_j - \hat{c}_j - \hat{v}_j + \epsilon < \epsilon$. Hence, we have (G.4) since $u(\theta) = 0$ for such a type θ.

Consider a type $\theta = (i,b,c)$ such that $\sum_{j\in J} x_j(\theta) = 1$, then $x_{j'}(\theta) > 0$ implies that

$$U_{j'}(\hat{v},\hat{c}) \geq b_j - \hat{c}_j - \hat{v}_j \qquad \text{for all } j \in J. \tag{G.12}$$

Taking the infimum over $j' \in \{j' : x_{j'}(\theta) > 0\}$, the left hand side becomes $u(i,b,c)$. The right hand side is at least $b_j - c_j - v_j - \epsilon$ as before, from which we get (G.4). In fact, this argument proves that the outcome is $(\epsilon/2)$-stable, which is a stronger result than ϵ-stable.

Finally, for the bound on the average per-agent communication cost. Note that the expected number of providers each customer contacts until she obtains an "yes" is at most $1/d^J$. Moreover, each interaction requires only $\lceil \log_2(n+1) \rceil + \lceil \log_2(1/\epsilon) \rceil + 1 = O(\log n)$ bits of communication. □

The above proof implies the existence of an ϵ-stable generalized transaction outcome for any $\epsilon > 0$. It also implies existence for $\epsilon = 0$ if we modify the protocol so that the provider communicates his actual cost c_{ij} instead of the estimate $\Delta_j := \lceil (c_j - \hat{c}_j)/\epsilon \rceil$.

G.2. Protocol 2 (providers search) in the dynamic model

Given $\epsilon > 0$, a batch size parameter $m \geq 1$, and a provider surplus vector $\hat{v} \in [0,1]^J$, consider the following dynamic matchmaking protocol, where the set of customers is labeled as $A = \{1, 2, \cdots, m\}$ without loss of generality and for each $a \in A$, her type (i_a, b_a, c_a) is defined as in Definition G.1.

1. **Initialization:** Initialize $\hat{b}_{aj} = \overline{b}_{i_a j} - \epsilon$, $\hat{u}_a = 0$, $J^* = J$. Initialize (x,p) to be an empty transaction outcome. Define $V_{aj}(\hat{u},\hat{b}) = \hat{b}_{aj} - \hat{u}_a - c_{aj}$.

2. **Providers search for customers:** If $|J^*|=0$, then terminate the protocol and return (x,p). Otherwise, let $j^*\in J^*$ be the provider with the smallest index. Define $A^*=\{a\in A: V_{aj}(\hat{u},\hat{b})\geq \hat{v}\}$. If $|A^*|=0$, the provider j^* sends the message "0," in which case he is removed from the set J^*, and the protocol repeats step 2. Otherwise, provider j^* communicates the identity of the customer $a\in A^*$ with the smallest index, and the customer responds whether $b_{aj^*}\geq \hat{b}_{aj^*}$.
   - If "yes," update $x_{aj^*}=1$ and $p_a=\hat{b}_{aj^*}-\hat{u}_a$, and increment $\hat{u}_a$ by ϵ. If a was previously matched to a different provider j', then update $x_{aj'}=0$ and add j' to the set J^*. Repeat step 2.
   - If "no," the customer communicates $\Delta=\lceil(\hat{b}_{aj}-b_{ij})/\epsilon\rceil$, and the protocol updates $\hat{b}_{aj}=\hat{b}_{aj}-\epsilon\Delta$. Repeat step 2.

PROPOSITION G.3 **(Performance guarantees of providers search)**. *Given $\epsilon>0$ and $m\geq 1$, define an equilibrium surplus vector $\hat{v}$ for the above protocol analogously as in Definition G.2. An equilibrium surplus vector $\hat{v}$ always exists, and given such a $\hat{v}$, the generalized transaction outcome returned by the above protocol is ϵ-stable. If $m=\Omega(d^I n/(\log n))$, then the average per-agent communication cost is at most $O((\log n)/d^I)$.*

Proof of Proposition G.3. The existence of the equilibrium surplus vector $\hat{v}$ follows from the same argument as in the proof of Proposition G.3. At such a $\hat{v}$, if $(x,\underline{p},\overline{p})$ is the generalized transaction outcome induced by the protocol, then the outcome is individually rational. Moreover, if v is defined as in (G.6), then $v_j\geq \hat{v}_j$ for every $j\in J$.

The construction of the protocol ensures that $\hat{b}_{aj}\geq b_{aj}-\epsilon$ throughout. Moreover, the termination condition of the protocol ensures that

$$\hat{v}\geq V_{aj}(\hat{u},\hat{b})\geq b_{aj}-\hat{u}_a-c_{aj}-\epsilon \qquad \text{for every } j\in J. \tag{G.13}$$

Now, the surplus of customer a is at least $\hat{u}_a-\epsilon$ upon termination. Combining with the above and taking the infimum over the customer surplus when the customer's type is (i,b,c), we get that for any type $(i,b,c)\in\Theta$ and any $j\in J$,

$$u(i,b,c)+v_j\geq u(i,b,c)+\hat{v}_j\geq b_j-c_j-2\epsilon, \tag{G.14}$$

so $(x,\underline{p},\overline{p})$ is ϵ-stable.

Finally, to show the bound on average per-agent communication cost, observe that the protocol requires n bits for each of the n providers to communicate that $|A^*|=0$. Furthermore, for each instance in which a provider contacts a customer, the number of bits of communication is at most $\lceil\log_2(m+1)\rceil+\lceil\log_2(1/\epsilon)\rceil+1$. Now, as in the proof of Proposition 4, each $\hat{u}_a$ can increment at most $\lceil 1/\epsilon\rceil$ times, and the customer will say "yes" with probability at least d^I each time, resulting in an increment. Therefore, the per-agent communication cost is at most

$$\left\lceil\frac{1}{\epsilon}\right\rceil\frac{\lceil\log_2(m+1)\rceil+\lceil\log_2(1/\epsilon)\rceil+1}{d^I}+\frac{n}{m}, \tag{G.15}$$

which is $O((\log n)/d^I)$ if $m=\Omega(d^I n/(\log n))$. □

G.3. Protocol 3 (both sides search) in the dynamic model

Given $\epsilon > 0$, a batch size parameter $m \geq 1$, a matrix $z \in [0,1]^{I\times J}$, and a provider surplus vector $\hat{v} \in [0,1]^J$, consider the following dynamic protocol: the protocol is identical to the version of Protocol 3 in Appendix C with the set of customers being A and set of providers J, except that $\hat{v}$ never increments during the protocol and is fixed at the original given value.

The following result is analogous to Propositions 3, 6 and 7.

Proposition G.4 **(Performance guarantees of both sides search)**. *Given parameters $\epsilon > 0$, $m \geq 1$ and z, an equilibrium surplus vector $\hat{v}$ for the above protocol (defined analogously as in Definition G.2) always exists. At such a $\hat{v}$, the generalized transaction outcome induced by the above protocol is ϵ-stable, and the average per-agent communication cost is bounded above by the following expression:*

$$1 + \frac{n}{m} + \kappa_1 \left(\sum_{i\in I, j\in J} \lambda_i \mathbb{P}_{y\sim C_{ij}}(y < z_{ij}) \right) + \kappa_2 \left(\sum_{i\in I} \frac{\lambda_i}{\min_{j\in J}\{\mathbb{P}_{y\sim C_{ij}}(y \leq z_{ij} + \epsilon | y \geq z_{ij})\}} \right), \tag{G.16}$$

where $\kappa_1 := \lceil \log_2 m \rceil + \lceil \log_2(1/\epsilon) \rceil$ and $\kappa_2 := \lceil \log_2(1+n) \rceil + \lceil \log_2(1/\epsilon) \rceil + 1$.

Suppose $\epsilon = \Omega(1)$, $\Omega(\sqrt{n}/(\log n)) \leq m \leq n^{O(1)}$, and $z = z^*$, where z^* is as defined in (22) *in Section 3.2, then the bound in* (G.16) *is at most $O(\sqrt{n}\log n)$.*

Proof of Proposition G.4. The existence of equilibrium surplus vector $\hat{v}$ for any given ϵ, m and z follows from the same argument as in the proof of Proposition G.2. The fact that the induced generalized transaction outcome induced is ϵ-stable follows also from the same argument as in the proof of Proposition G.2, after observing that at the beginning of Step 3 of the static protocol, $\hat{c}_{aj} \leq c_{aj} + \epsilon$ for any customer $a \in A$ and provider $j \in J$.

To show the bound on the average per-agent communication cost, note that one bit per customer is needed to terminate her search in Step 3; the second term of n/m comes from the fact that in Step 2 of Protocol 3 in Appendix C, every provider needs to send at least one bit, and the $|J| = n$ bits is divided by m when computing the average communication cost. Now, κ_1 is the number of bits needed each time a provider contacts a customer in Step 2, so the second term bounds the expected number of bits sent in Step 2 per customer. Similarly, κ_2 is the number of bits needed for each iteration of Step 3, and the last term bounds the expected number of bits sent in Step 3 per customer. These bounds arise from similar arguments as those in the proof of Proposition 6.

The bound of $O(\sqrt{n}\log n)$ when $z = z^*$ follows from the fact that $\mathbb{P}_{y\sim C_{ij}}(y < z^*_{ij}) \leq \lceil 1/\epsilon \rceil / \sqrt{n} = O(1/\sqrt{n})$ and $\mathbb{P}_{y\sim C_{ij}}(y \leq z^*_{ij} + \epsilon | y \geq z^*_{ij}) \geq \mathbb{P}_{y\sim C_{ij}}(y \in [z^*_{ij}, z^*_{ij} + \epsilon]) \geq 1/\sqrt{n}$. □

In the above protocol, providers reach out to customers first, which is the same as in Protocol 3 in the main body of the paper. In both the static and the dynamic models, it is possible to reverse the order of communication in "both sides search" so that customers reach out first to providers they especially like, and then providers search among customers and initiate contact. In the dynamic model, such a protocol would be a modification of the dynamic version of Protocol 2 in Section G.2, in which the estimates of benefits are initialized instead as $\hat{b}_{aj} = z_{ij} - \epsilon$, where i is the segment for customer a. Moreover, before providers initiate contact, there is an additional step in which customers whose benefit b_{ij} for a provider j exceeds the threshold z_{ij} proactively contact the provider and communicate their benefit. Suppose that $\epsilon = \Omega(1)$, $\Omega(n) \le m \le n^{O(1)}$, and z is defined as follows:

$$z^*_{ij} := \sup\left\{y \in [0, \bar{b}_{ij}] : \mathbb{P}(b_{ij} \in [y-\epsilon, y]) \ge 1/\sqrt{n}\right\}. \tag{G.17}$$

Then a similar analysis as the above implies that this protocol also terminates with an ϵ-stable generalized transaction outcome, and the average per-agent communication cost is at most $O(\sqrt{n}\log n)$. Note that the minimum batch size required by this performance bound is $\Omega(n)$, which is larger than the minimum batch size of $\Omega(\sqrt{n}/(\log n))$ when providers reach out to customers first. This suggests that having providers reach out first is more efficient when the batch size is small.

G.4. Protocol 4 (centralized matching) in the dynamic model

This section shows that a variant of Protocol 4 can approximate an ϵ-stable outcome in the dynamic model using relatively low amounts of communication, as long as d^I and d^J are both sufficiently large. The notion of approximation is that the proportion of customers who receive a near-optimal match is at least $1-\delta$ for some parameter $\delta > 0$, which can be made arbitrarily small. Here, the notion of a near-optimal match is based on the solution of an assignment LP with optimistic estimates of benefits and costs, similar to in Proposition 9.

Suppose that the providers can be partitioned into ex-ante homogeneous segments, $J = \dot{\bigcup}_{k\in K} J_k$, such that the preference distributions B_{ij} and C_{ij} depend only on i and the provider segment $k \in K$ that contains the provider j. Suppose also that the capacities of providers are identical within each provider segment, and that the minimum segment size is $l := \min_{k\in K}\{|J_k|\}$. The only difference from Definition 2 is that within each segment, the entire distributions B_{ij} and C_{ij} are identical, not only the supports $\bar{b}_{ij}$ and $\underline{c}_{ij}$.

Given $\epsilon > 0$. For each $i \in I$ and $k \in K$, define $\hat{b}_{ik} = \bar{b}_{ij} - \epsilon$, where $j \in J_k$. The exact choice of j does not matter for the definition because of the assumption that the distribution B_{ij} is identical within each provider segment. Similarly, define $\hat{c}_{ik} = \underline{c}_{ij} + \epsilon$. Let $\hat{x}$ be an optimal solution to the following LP, which corresponds to the optimization problem of matching customers to provider segments based on maximizing welfare using the above optimistic estimates of benefits and costs.

$$\text{Maximize:} \quad \sum_{i\in I, k\in K} \lambda_i(\hat{b}_{ik} - \hat{c}_{ik})\hat{x}_{ik} \tag{G.18}$$

$$\text{s.t.} \quad \lambda \sum_{k\in K} \hat{x}_{ik} \le \lambda_i \qquad \text{for each } i \in I, \tag{G.19}$$

$$\sum_{i\in I} \lambda_i \hat{x}_{ik} \le Q_k := \sum_{j\in J_k} q_j \qquad \text{for each } k \in K, \tag{G.20}$$

$$\hat{x} \ge 0 \tag{G.21}$$

Given $\hat{x}$, define $\hat{r}_i$ as the rate that a segment i customer is matched, and $\hat{q}_k$ as the rate that each segment k provider is matched:

$$\hat{r}_i := \sum_{k\in K} \hat{x}_{ik}, \tag{G.22}$$

$$\hat{q}_k := \frac{1}{|J_k|} \sum_{i\in I} \lambda_i \hat{x}_{ik}. \tag{G.23}$$

Moreover, let $(\hat{u},\hat{v})$ be an optimal solution to the dual LP:

$$\text{Minimize:} \quad \sum_{i\in I}\lambda_i\hat{u}_i+\sum_{k\in K}Q_k\hat{v}_k \tag{G.24}$$

$$\text{s.t.} \quad \hat{u}_i+\hat{v}_k\geq \hat{b}_{ik}-\hat{c}_{ik} \qquad \text{for each } (i,k)\in I\times K, \tag{G.25}$$

$$\hat{u},\hat{v}\geq 0 \tag{G.26}$$

Given $\epsilon,\delta>0$, consider the following dynamic matchmaking protocol with batch size 1. Let the customer's type be (i,b,c). Without loss of generality, label the provider segments as $K=\{1,2,\cdots,|K|\}$.

1. **Recommendation generation:** Sample a segment $k\in K\cup\{0\}$ based on the probability vector $(1-\hat{r}_i,\hat{x}_{i1},\hat{x}_{i2},\cdots,\hat{x}_{i|K|})$. If $k=0$, then terminate the protocol without matching the customer to anyone. Otherwise, generate a random subset $S\subseteq J_k$ by including each provider $j\in J_k$ with probability a_k, where
$$a_k:=\min\left(1,\frac{\ln(1/\delta)}{|J_k|d^Id^J}\right). \tag{G.27}$$
Each provider $j\in S$ is said to be a recommended provider for the customer. Note that the expected number of recommended providers is $a_k|J_k|\leq(\ln(1/\delta))/(d^Id^J)$.

2. **Preference elicitation:** For each recommended provider $j\in S$, ask the customer whether $b_j\geq\hat{b}_{ik}$, and ask the provider whether $c_j\geq\hat{c}_{ik}$. If both answers are "yes," then the recommended provider is said to be accepted. If no provider $j\in S$ is accepted, then terminate the protocol without matching the customer to anyone.

3. **Match determination:** Match the customer to a uniformly random recommended provider who was accepted, with payment $\hat{v}_k+\hat{c}_{ik}$.

Step 3 can be interpreted as either the natural outcome of preference heterogeneity among customers when breaking ties among accepted providers, or as the result of a conscious effort by the platform to evenly split demand among similar providers so as to maximize the overall match rate.

Proposition G.5 **(Performance gurantees of centralized matching)**. *Given any $\epsilon>0$ and $\delta>0$, the above is a matchmaking protocol with an average per-agent communication cost of $2(\ln(1/\delta))/(d^Id^J)$. Suppose that $d^Id^J\geq\min(1,(\ln(1/\delta))/l)$, where l is the minimum segment size, then for each segment $i\in I$, the proportion of customers who are matched is at least $(1-\delta)\hat{r}_i$; conditional on being matched, each customer's surplus is at least $\hat{u}_i$. Similarly, for each provider j in segment k, the rate at which he is matched is at least $(1-\delta)\hat{q}_k$, where $\hat{q}_k$ is defined in (G.23). The market surplus per customer is at least $W^*-2\epsilon-\delta$, where W^* is the maximum surplus per customer from any feasible generalized matching.*

Proof of Proposition G.5. Note first that the randomization in steps 1 and 3 is allowed by Definition G.1. This is because one can equivalently think of the above as randomizing over deterministic communication protocols, each of which is associated with a deterministic choice of the set S for each customer segment $i\in I$, as well as a permutation for each customer segment that encodes a priority ordering for choosing among accepted providers in step 3. The only communication in each deterministic protocol is in step 2, which requires $2|S|$ bits.

The generalized transaction outcome induced by the above protocol is guaranteed to be individually rational, because whenever a customer from segment i is matched to a provider from segment k, $\hat{x}_{ik} > 0$, which implies by complementary slackness that the payment $\hat{v}_k + \hat{c}_{ik} = \hat{b}_{ik} - \hat{u}_i$. Due to the definition of what it means for a recommended provider is accepted, we have that conditional on being matched, the customer's surplus is at least $\hat{u}_i \geq 0$, and the provider's surplus from this transaction is at least $\hat{v}_k \geq 0$.

Now, if $d^I d^J = 1$, then every match recommendation is accepted. If $d^I d^J < 1$, then $l \geq (\ln(1/\delta))/(d^I d^J)$, so the chance that none of the recommended providers in S is accepted is upper-bounded by $(1 - a_k d^I d^J)^{|J_k|} \leq e^{-\ln(1/\delta)} = \delta$, and this holds within each segment of customers. The desired lower bound on the proportion of customers being matched follows from the fact that the probability that Step 2 is executed is equal to $\hat{r}_i$, and Step 2 moves to Step 3 with probability at least $1-\delta$. Similarly, for each $k \in K$, the rate at which some provider within J_k is matched to a customer is at least $(1-\delta)\sum_{i \in I} \lambda_i \hat{x}_{ik} = (1-\delta)|J_k|\hat{q}_k$. By symmetry, each individual provider within the segment is matched with rate at least $(1-\delta)\hat{q}_k$.

The above implies that the market surplus induced by the above protocol is at least $(1-\delta)W'$, where W' is the optimal objective to the assignment LP (G.18)-(G.21), as well as to the dual LP (G.24)-(G.26). If x' is any feasible generalized matching, then the market surplus induced by x' is

$$\mathbb{E}_{(i,b,c)\sim F}\left[(b-c)\cdot x'(i,b,c)\right] \leq 2\epsilon + \mathbb{E}_{(i,b,c)\sim F}\left[(b-c-2\epsilon)\cdot x'(i,b,c)\right] \leq 2\epsilon + W'. \tag{G.28}$$

The first inequality follows from the total rate of matching being upper-bounded by the total arrival rate of customers, which is normalized to 1. The second inequality follows from the definition of $\hat{b}$ and $\hat{c}$. Therefore, $W^* \leq 2\epsilon + W'$, and the market surplus induced by the above protocol is at least

$$(1-\delta)W' \geq (1-\delta)(W^* - 2\epsilon) \geq W^* - 2\epsilon - \delta, \tag{G.29}$$

where the last inequality comes from $W^* \leq 1$, as $\overline{b}_{ij} - \underline{c}_{ij} \leq 1$ and the total rate of matching is at most 1.

G.5. Near-optimality in the dynamic model

As in the static model, one can ask the following question in the dynamic model: for a given class of markets, in the worst case, how many bits of communication per customer is needed to achieve a good outcome? The answers are given in the following table, which is almost identical to Table 4 for the static model.

The upper bounds in Table 5 follow from Propositions G.2, G.3, G.4 and G.5. The lower bounds follow from the following proposition, which is analogous to Proposition 12.

PROPOSITION G.6 **(Lower bounds in the dynamic model)**. *In the dynamic model, define $M(n, d^I, d^J)$ as the market with one segment of customers and n providers, each with capacity at least $1/n$. For each provider $j \in J$, the benefit and cost associated with each customer are independently drawn from the following distribution,*

$$b_j = \begin{cases} .99 & \textit{with probability } d^I, \\ 0 & \textit{otherwise;} \end{cases} \tag{G.30}$$

$$c_j = \begin{cases} .01 & \textit{with probability } d^J, \\ 1 & \textit{otherwise.} \end{cases} \tag{G.31}$$

Table 5 The average per-customer communication cost needed in the worst case by any dynamic matchmaking protocol that guarantees the given notion of a good outcome in every market within the given class. All bounds assume a constant $\epsilon > 0$, which is incorporated into the constants in the big-O notation.

Space of protocols	Notion of a good outcome	Class of markets	Lower bound	Upper bound	Protocol achieving the upper bound
Unrestricted	ϵ-stability	$d^J \geq \Omega(1)$	$\Omega(\log n)$	$O(\log n)$	Customers search with batch size one.
		$d^J \geq 1/\sqrt{n}$	$\Omega(1/d^J)$	$O((\log n)/d^J)$	
		$d^I \geq \Omega(1)$	$\Omega(\log n)$	$O(\log n)$	Providers search with batch size $d^I n$.
		$d^I \geq 1/\sqrt{n}$	$\Omega(1/d^I)$	$O((\log n)/d^I)$	
		All markets	$\Omega(\sqrt{n})$	$O(\sqrt{n}\log n)$	Both sides search with batch size $\sqrt{n}$.
One-round protocols	Individually rational outcome that matches $1-\delta$ proportion of customers	$d^I d^J \geq \Omega(1)$, and the minimum segment size $l \geq \ln(1/\delta)/(d^I d^J)$.	$\Omega(\log(1/\delta))$	$O(\log(1/\delta))$	Centralized matching with batch size one.

a) *A dynamic matchmaking protocol is said to guarantee ϵ-stability in a given market if the generalized transaction outcome induced by the protocol is always ϵ-stable. Any dynamic matchmaking protocol that guarantees ϵ-stability in the market $M(n, d^I, d^J)$ requires an average per-agent communication cost of at least $\phi(n, d^I, d^J)$, where*

   i) $\phi(n, 1, 1/n) = (1 - e^{-1})\log_2 n$;

   ii) $\phi(n, d^I, d^J) = .007/d^J$ *if* $d^I d^J = 1/n$ *and* $d^J \in [1/\sqrt{n}, 1/2]$;

   iii) $\phi(n, 1, 1/n) = (1 - e^{-1})\log_2 n$;

   iv) $\phi(n, d^I, d^J) = .007/d^I$ *if* $d^I d^J = 1/n$ *and* $d^I \in [1/\sqrt{n}, 1/2]$;

b) *Any dynamic matchmaking protocol that only uses one-round communication protocols and matches at least $1-\delta$ fraction of customers in the market $M(n, 1/2, 1/2)$ requires an average per-customer communication cost of at least $2\log_2(1/\delta)$.*

Proof of Proposition G.6. Consider a matchmaking protocol with batch size m. We can think of the randomization over deterministic communication protocols as sampling a random variable R, and the protocol is deterministic conditional on R. Let Π denote the transcript of the protocol, which is the entire sequence of zero-one bits that are communicated before termination, and let $|\Pi|$ denote the length of the protocol, which is the total number of bits communicated. Let the set of customers be $[m]$ and let the benefit and cost vector of each customer $a \in [m]$ be b^a and c^a. The average per-customer communication cost is

$$\frac{1}{m}\mathbb{E}[|\Pi|] = \frac{1}{m}\mathbb{E}_R\left[\mathbb{E}\left[|\Pi|\middle|R\right]\right] \geq \frac{1}{m}H(\Pi|R) \geq \frac{1}{m}\sum_{a\in[m]}\mathcal{I}(b^a, c^a; \Pi|R). \tag{G.32}$$

Here, $H(X|Z)$ is the conditional entropy of X given Z, and $\mathcal{I}(X;Y|Z)$ is the conditional mutual information between X and Y given Z. For background on these concepts from information theory, see chapters 1 and 2 of Cover and Thomas (2012). The above implies that in order to lower bound the average per-customer communication cost, it suffices to lower bound $\mathcal{I}(b^a, c^a; \Pi)$ for any deterministic protocol Π. In the remainder of the proof, we consider one customer at a time, so we drop the superscript a.

As in the proof of Proposition 12, by the independence of preferences,

$$\mathcal{I}(b, c; \Pi) \geq \sum_{j=1}^{n} \mathcal{I}(b_j, c_j; \Pi). \tag{G.33}$$

For part a)-i), note that for any $j \in J$,

$$\mathcal{I}(b_j, c_j; \Pi) \geq \mathcal{I}(b_j; \Pi) \geq \mathbb{E}[x_j^\Pi] D(b_j|_{x_j^\Pi=1} \| b_j) = \mathbb{E}[x_j^\Pi] \log_2(n), \tag{G.34}$$

where x_j^Π is defined as whether the customer is matched to provider j when the transcript is Π, and the last inequality comes from the fact that $\mathbb{P}(b_j = .99 | x_j^\Pi = 1) = 1$, and the KL-divergence $D(Bernoulli(1) \| Bernoulli(1/n)) = \log_2(n)$. The desired result follows by combining (G.33) and (G.34) with the following expression for the probability that a customer in $M(n, 1/n, 1)$ is matched in an ϵ-stable generalized transaction outcome:

$$\sum_{j=1}^{n} \mathbb{E}[x_j^\Pi] = 1 - \left(1 - \frac{1}{n}\right)^n \geq 1 - e^{-1}. \tag{G.35}$$

A similar argument implies the bound in part a)-iii).

For part a)-ii), define the random binary variable P_j as whether provider j is a perfect match for the customer, which means that $(b_j, c_j) = (.99, .01)$ and $(b_{j'}, c_{j'}) \neq (.99, .01)$ for any $j' \neq j$. Note that for any $j \in J$,

$$\mathbb{E}[P_j] = \frac{1}{n}\left(1 - \frac{1}{n}\right)^{n-1} = \frac{1}{n} \Big/ \left(1 - \frac{1}{n-1}\right)^{n-1} \geq \frac{e^{-1}}{n}. \tag{G.36}$$

Now, since the protocol guarantees an ϵ-stability,

$$\mathbb{P}\big(x_j^\Pi = 1 | (b_j, c_j) = (.99, .01)\big) \geq \frac{\mathbb{E}[P_j]}{1/n} \geq e^{-1}, \tag{G.37}$$

$$\mathbb{P}\big(x_j^\Pi = 1 | (b_j, c_j) = (0, 1)\big) = 0. \tag{G.38}$$

Therefore, the argument in the proof of Lemma D.2 b) in Appendix D.5 implies that

$$\mathcal{I}(b_j, c_j; \Pi) \geq \frac{(e^{-1})^2/(2\ln 2)}{\left(\sqrt{\frac{1}{d^I(1-d^J)}} + \sqrt{\frac{1}{(1-d^I)d^J}} + \sqrt{\frac{1}{(1-d^I)(1-d^J)}}\right)^2}. \tag{G.39}$$

As in (D.47), we have that the denominator is at most $13nd^J$ under the assumptions of a)-ii), and it is at most $13nd^I$ under the assumptions of a)-iv). Combining with (G.33) yields the desired result. The symmetric argument implies the result for part a)-iv).

For part b), we follow a similar argument as in the proof of Proposition 12 c) in Appendix D.6. Let β be the message of the customer and β_j be the message of provider j. Define

$$G(\beta) := \Big\{ j : \mathbb{P}(b_j = .99 | \beta) = 1 \Big\}. \tag{G.40}$$

Since Definition G.1 requires the outcome to be individually rational, whenever the customer is matched to provider j, it must be that $j \in G(\beta)$. Now,

$$\mathcal{I}(b;\beta) \geq \sum_{j=1}^{n} \mathbb{E}_\beta \Big[D(b_j|\beta \| b_j) \Big] \geq \mathbb{E}\Big[|G(\beta)|\Big], \tag{G.41}$$

where the first inequality comes from the independence of each component of b and the identity relating mutual information and KL-divergence. The second inequality follows from $\mathbb{E}_\beta\Big[D(b_j|\beta\|b_j)\Big] \geq \mathbb{P}(j \in G(\beta))D(Bernoulli(1)\|Bernoulli(.5)) = \mathbb{P}(j \in G(\beta))$. Now, by the constraint that the protocol needs to match $1-\delta$ proportion of customers, we have

$$\delta \geq \mathbb{E}_\beta\left[2^{-|G(\beta)|}\right] \geq 2^{-\mathbb{E}[|G(\beta)|]}, \tag{G.42}$$

where the second inequality is by Jensen's inequality.

Combining (G.41) and (G.42), we get $\mathcal{I}(b;\beta) \geq \log_2(1/\delta)$. The symmetric argument yields $\mathcal{I}(c;\gamma) \geq \log_2(1/\delta)$. The desired result follows from $\mathcal{I}(b,c;\Pi) \geq \mathcal{I}(b;\beta) + \mathcal{I}(c;\gamma)$, which holds because (b,β) is independent of (c,γ). □

Appendix H: Model without transfer payments

Consider the following alternative model with non-transferable utility (NTU). Reinterpret b_{ij} as customer i's value of being matched to provider j, and c_{ij} as provider j's value of being matched to customer i, with higher values being preferred. These exogenously given values completely pin down agents' utilities in a match, and there is no longer an endogenously determined transfer payment. An unmatched agent receives zero value, and the bounded support assumption is modified to $b_{ij}, c_{ij} \in [-1,1]$, so that it is possible for an agent to prefer being unmatched over matching with certain partners. Each matching x induces agent utilities $u_i := \sum_{j\in J} b_{ij}x_{ij}$ and $v_j := \sum_{i \in I} c_{ij}x_{ij}$, and the matching is ϵ-stable if $u, v \geq 0$ and for any $(i,j) \in I \times J$, either $u_i \geq b_{ij} - \epsilon$ or $v_j \geq c_{ij} - \epsilon$.

The analogous research question is how to facilitate an ϵ-stable matching with high probability using minimal communication. This section defines the analogs of the four protocols and re-derives all of the results in the main body of the paper. The protocols all satisfy the formal definition of a communication protocol in Section 4.1. The following is analogous to the description of Protocol 3 in Appendix C.

NTU Protocol 3 *Given $\epsilon > 0$, as well as an $I \times J$ matrix z that depends only on the commonly known market information (I, J, B, C).*

1. **Initialization:** *Initialize the matching x to be empty, and the set $A = I$. For each $j \in J$, initialize $\hat{v}_j = 0$. For each $(i,j) \in I \times J$, initialize $\hat{c}_{ij} = z_{ij} - \epsilon$. Define $\hat{J}_i(\hat{v}, \hat{c}) := \{j \in J : \hat{c}_{ij} \geq \hat{v}_j$ and $b_{ij} \geq 0\}$.*
2. **Providers search for customers:** *Each provider $j \in J$ sends a message $(i, \Delta_{ij} := \lfloor (c_{ij} - \hat{c}_{ij})/\epsilon \rfloor)$ for each $i \in I$ such that $c_{ij} > z_{ij}$. The protocol updates $\hat{c}_{ij} = \hat{c}_{ij} + \epsilon\Delta_{ij}$. After each provider is finished contacting customers, he uses one bit to communicate that he is done.*
3. **Customers search for providers:**
   (a) **Customer selection:** *If $|A| = 0$, then terminate the protocol with the matching x. Otherwise, select the customer $i \in A$ with the smallest index and proceed to step 3(b).*

(b) **Provider interaction:** *The selected customer i sends the message "0" if the set $\hat{J}_i(\hat{v},\hat{c})$ is empty, in which case the protocol removes i from A and go back to step 3(a). Otherwise, customer i selects the provider*

$$j^* \in \arg\max_j \{b_{ij} : j \in \hat{J}_i(\hat{v},\hat{c})\} \tag{H.1}$$

with the smallest index and communicates the index of j^*. The provider j^* responds with one bit indicating whether $c_{ij^*} \geq \hat{c}_{ij^*}$.

- *If "yes," then the protocol increments $\hat{v}_{j^*}$ to $\hat{c}_{ij^*} + \epsilon$, updates $x_{ij^*} = 1$, and removes i from A. If j^* was previously matched to another customer i', then the protocol updates $x_{i'j^*} = 0$, and adds i' to A. Go back to step 3(a).*
- *If "no," then provider j^* communicates $\Delta_{ij^*} := \lfloor (c_{ij^*} - \hat{c}_{ij^*})/\epsilon \rfloor$. The protocol updates $\hat{c}_{ij^*} = \hat{c}_{ij^*} + \epsilon\Delta_{ij^*}$. Repeat step 3(b).*

Define NTU Protocol 1 as the special case of NTU Protocol 3 in which $z_{ij} \equiv \bar{c}_{ij}$, and NTU Protocol 2 as the mirror image of NTU Protocol 1.

NTU Protocol 4 *Given $\epsilon > 0$, and a set $S_i \subseteq J$ for each customer $i \in I$. The protocol is identical to the Protocol 4 in Section 4.1, except that all references to the price vector p are removed. Moreover, since high values of c_{ij} are preferred and preferences may be negative, the condition for accepting a recommended pair (i,j) is if $b_{ij} \geq \max(0, \bar{b}_{ij} - \epsilon)$ and $c_{ij} \geq \max(0, \bar{c}_{ij} - \epsilon)$.*

The following result is analogous to Proposition 1, except that instead of guaranteeing near-optimal average surplus, ϵ-stability in the NTU model guarantees an approximate version of Pareto optimality.

Proposition H.1 **(Approximate Pareto optimality)**. *Given any ϵ-stable matching x in the NTU model, it is impossible for another feasible matching x' to simultaneously improve everyone's utility by more than ϵ.*

Proof of Proposition H.1. Suppose on the contrary that such an x' exists, then any matched pair (i,j) would be an ϵ-blocking pair in x, which contradicts the ϵ-stability of x. On the other hand, if x' doesn't match anyone, then everyone's utility is zero, but everyone's utility is at least zero in x by ϵ-stability. □

The following result is analogous to Propositions 2 and 3.

Proposition H.2 **(Effectiveness of NTU Protocol 3)**. *Given any $\epsilon > 0$ and any choice of z, NTU Protocol 3 always terminates with a (2ϵ)-stable matching.*

Proof of Proposition H.2. The protocol always terminates, because in every iteration of step 3(b), either $\hat{c}_{ij^*}$ is updated to be within $[c_{ij^*} - \epsilon, c_{ij^*}]$, after which it will not be updated again, or v_{j^*} increments by at least ϵ. The matching x returned by the protocol is always individually rational, as no one is forced to be matched with an undesirable partner. To see that x is (2ϵ)-stable, observe that at the beginning of step 3, $c_{ij} \leq \hat{c}_{ij} + \epsilon$ for all $(i,j) \in I \times J$. Moreover, in step 3, each component of $\hat{c}$ can only decrease and each component of $\hat{v}$ can only increase. Let $\hat{c}$ and $\hat{v}$ refer to their final values upon the protocol's termination, and (u,v) be the agent surpluses associated with x. Note that $u, v \geq 0$ since no one is ever forced to be matched to an unacceptable partner. By the customers' optimization in (H.1) and the monotonicity of $\hat{c}$ and $\hat{v}$, we have

$$u_i \geq b_{ij} \qquad \text{for all } j \text{ such that } \hat{c}_{ij} \geq \hat{v}_j. \tag{H.2}$$

Moreover, step 3 guarantees that $v_j \geq \hat{v}_j - \epsilon$, where the "$-\epsilon$" comes from the increment in $\hat{v}_j$ after being matched. Therefore, for all j such that $\hat{c}_{ij} < \hat{v}_j$, we have

$$v_j \geq \hat{v}_j - \epsilon > \hat{c}_{ij} - \epsilon \geq c_{ij} - 2\epsilon. \tag{H.3}$$

Combining (H.2) and (H.3), we have that x is (2ϵ)-stable. □

Since high values of c_{ij} are desirable in the NTU model, we define the preference density parameter d^J in the NTU model as the largest value such that

$$\mathbb{P}(c_{ij} \geq \overline{c}_{ij} - \epsilon) \geq d^J \qquad \text{for all } (i,j) \in I \times J. \tag{H.4}$$

The definition of d^I is unchanged and is given in (15).

Proposition 4 as stated in Section 3.1 continues to hold in the NTU model, and the same proof applies. This is because the chance that the provider j^* responds in Step 3b) with a "yes" is still at least d^J in each iteration, independent of everything else. Moreover, for each provider j, $\hat{v}_j$ can still increment at most $\lceil 1/\epsilon \rceil$ times before no customer will reach out to the provider.

The following result is analogous to Proposition 5, in that it shows that allowing both sides to initiate contact can yield dramatic improvements if both d^I and d^J are very low.

Proposition H.3 **(Illustration of NTU Protocol 3)**. *Consider a market in which n is even and there are two segments of customers and providers: $I = I_1 \cup I_2$ with $|I_1| = |I_2|$, and $J = J_1 \cup J_2$ with $|J_1| = |J_2|$. Suppose that for any $(i,j) \in I_1 \times J_1$, both b_{ij} and c_{ij} are negative. For any $(i,j) \in I_2 \times J_2$, $b_{ij} = c_{ij} = .1$. For any $(i,j) \in I_1 \times J_2$, $c_{ij} = .2$ and*

$$b_{ij} = \begin{cases} 1 & \text{with probability } d^I := 1/n, \\ -1 & \text{otherwise} \end{cases} \tag{H.5}$$

Symmetrically, for any $(i,j) \in I_2 \times J_1$, $b_{ij} = .2$ and

$$c_{ij} = \begin{cases} 1 & \text{with probability } d^J := 1/n, \\ -1 & \text{otherwise} \end{cases} \tag{H.6}$$

In this market, for any $\epsilon < 1$, the expected number of interactions per customer is at least $.19n$ under NTU Protocol 1 or NTU Protocol 2, but is equal to 1.25 under NTU Protocol 3 if parameter $z_{ij} \equiv 0.2$.

Proof of Proposition H.3. Under NTU Protocol 1, each customer in I_2 would try to contact every provider in J_1 before settling with a partner in J_2. Let X be a geometrically distributed random variable with mean n. Each such customer would contact in expectation $\mathbb{E}[\min(X, n/2)]$ providers in J_1, where $q := (1-1/n)^{n/2} \leq e^{-0.5}$. Since $|I_2| = n/2$, the overall expected number of interactions per customers in this protocol is at least $n(1-q)/2 > .19n$. The same bound holds under NTU Protocol 2 by symmetry.

Under NTU Protocol 3 with $z_{ij} \equiv 0.2$, the expected number of interactions in Step 2 is $n/4$, as the $n/2$ providers in J_1 each reaches out to $1/2$ customers in I_2 in expectation. In Step 3, every customer would only contact at most one provider. So the total expected number of interactions per customer is 1.25. □

Proposition 6 in Section 3.2 continues to hold after accounting for the change that high values of c_{ij} are now preferred, so (21) becomes

$$N(z) := \left(\sum_{i\in I, j\in J} \mathbb{P}(c_{ij} > z_{ij}) \right) + \left\lceil \frac{1}{\epsilon} \right\rceil \left(\sum_{j\in J} \frac{1}{\min_{i\in I}\{\mathbb{P}(c_{ij} \geq z_{ij} - \epsilon | c_{ij} \leq z_{ij})\}} \right). \tag{H.7}$$

Similarly, Proposition 7 continues to hold if (22) is updated to account for the change in the direction of c_{ij}:

$$z^*_{ij} := \sup\{y \in [0, \bar{c}_{ij}] : \mathbb{P}(c_{ij} \in [y - \epsilon, y] \geq 1/\sqrt{n}\}. \tag{H.8}$$

The analogous proofs as those in Section 3.2 continue to apply, as the mechanics of NTU Protocol 3 is almost identical to that of Protocol 3.

The definition of providers being "relatively homogeneous" (Definition 2) needs to be modified in the NTU model so that it is the quantities $\bar{b}_{ij}$ and $\bar{c}_{ij}$ that depend only on the customer i and the segment k that contains provider j. This is because high values of c_{ij} are preferred in the NTU model, so an optimistic estimate of provider preference is $\bar{c}_{ij}$. The following result is analogous to Propositions 8 and 9.

PROPOSITION H.4 **(Effectiveness and performance of NTU Protocol 4)**. *Given $\epsilon > 0$ and a market in which the providers are relatively homogeneous with minimum segment size l (with the definition of "relatively homogeneous" being modified as explained above). Let $\hat{x}$ be a stable matching when preferences are given by $\hat{b}_{ij} := \bar{b}_{ij} - \epsilon$ and $\hat{c}_{ij} := \bar{c}_{ij} - \epsilon$. Consider NTU Protocol 4 in which the recommendation set S_i is defined as follows: for each customer i matched to a provider of segment k in $\hat{x}$, construct S_i to be a random subset of J_k such that each $j \in J_k$ is included with probability $a_k := \min(1, (3\ln n)/(|J_k| d^I d^J)$, independently across (i, j) pairs. For each customer i who is unmatched in $\hat{x}$, define $S_i = \emptyset$.*

If $d^I d^J \geq \min(1, (3\ln n)/l)$, then this protocol yields an ϵ-stable matching with probability at least $1 - 3/n^2$.

Proof of Proposition H.4. Let x be the matching returned by NTU Protocol 4, with utilities given by (u, v). Let $(\hat{u}, \hat{v})$ be the utilities associate with the matching $\hat{x}$ when preferences are given by the optimistic estimates $(\hat{b}, \hat{c}) = (\bar{b} - \epsilon, \bar{c} - \epsilon)$. Let $\hat{I}$ be the set of customers matched in $\hat{x}$. The proof of Proposition 9 in Appendix D.2 implies that with probability at least $1 - 3/n^2$, the set of matched customer in x is also equal to $\hat{I}$. When this occurs, we show that x is ϵ-stable.

Define $x(i)$ and $\hat{x}(i)$ to be the matched partner of customer i under matchings x and $\hat{x}$ respectively. For any subset $A \subseteq I$, define $x(A)$ and $\hat{x}(A)$ to be the respective set of matched partners to customers $i \in A$. Define $\sigma : J \to J$ as a bijection that maps each provider segment J_k to itself, such that if $j = x(i)$, then $\sigma(j) = \hat{x}(i)$. Moreover, if $j \in J_k \backslash x(\hat{I})$, then $\sigma(j) \in J_k \backslash \hat{x}(\hat{I})$. Such a bijection exists because within each segment J_k, the number of matched providers is the same in x and in $\hat{x}$. By construction, we have that $u_i \geq \hat{u}_i$ for all $i \in I$, and $v_j \geq \hat{v}_{\sigma(j)}$ for all $j \in J$. Since $\hat{u}, \hat{v} \geq 0$, we have that $u, v \geq 0$. Moreover, we have by the definition of $\hat{x}$ that for every $(i, j) \in I \times J$,

$$\text{either } u_i \geq \hat{u}_i \geq \hat{b}_{ij} = \bar{b}_{ij} - \epsilon \geq b_{ij} - \epsilon, \tag{H.9}$$

$$\text{or } v_j \geq \hat{v}_{\sigma(j)} \geq \hat{b}_{i\sigma(j)} = \bar{c}_{ij} - \epsilon \geq c_{ij} - \epsilon. \tag{H.10}$$

Thus, the matching x is ϵ-stable. □

Propositions 10 and 11 in Section 4.2 continue to hold in the NTU model if one simply replaces the phrase "ϵ-stable outcome" with "ϵ-stable matching." This is because the performance guarantees for the four protocols derived above are almost identical to those derived for the transferable utility model in Section 3.

Proposition 12 in Section 4.3 also holds if one updates the definition of the market $M(n, d^I, d^J)$ as having preferences drawn from the following distributions:

$$b_{ij} = \begin{cases} 1 & \text{with probability } d^I, \\ -1 & \text{otherwise;} \end{cases} \tag{H.11}$$

$$c_{ij} = \begin{cases} 1 & \text{with probability } d^J, \\ -1 & \text{otherwise.} \end{cases} \tag{H.12}$$

For any $\epsilon < 1$, any ϵ-stable matching for the above market must only match profitable pairs (i, j), for which $(b_{ij}, c_{ij}) = (1, 1)$, and cannot leave both agents in a profitable pair unmatched. This is the same property as that used in the proof of Proposition 12 except that a profitable pair (i, j) in the transferable utility case is one in which $(b_{ij}, c_{ij}) = (.99, .01)$. After accounting for this minor difference in definitions, the same proofs as in Appendices D.5 and D.6 will go though, and we obtain the same lower bounds as in Proposition 12 with the phrase "ϵ-stable outcome" everywhere replaced by "ϵ-stable matching."

The following result is analogous to Proposition A.1 in Section A.

PROPOSITION H.5 **(Incentive guarantees for centralized matching in the NTU model)**. *Under the assumptions of Proposition H.4, with probability at least $1 - 3/n^2$, NTU Protocol 4 returns an ϵ-stable matching, and no customer can unilaterally deviate and improve her utility by more than ϵ. Moreover, no provider of segment k can unilaterally deviate and obtain a utility higher than*

$$\max\{\overline{c}_{ij} : \hat{x}(i) \in J_k\}, \tag{H.13}$$

where $\hat{x}(i)$ is the provider matched to customer i under the matching $\hat{x}$ as defined in Proposition H.4.

Note that the incentive guarantees for providers is weaker in Proposition H.5 than in Proposition A.1. This is because in the NTU model, if a provider j is recommended to two customers i and i', where $c_{ij} >> c_{i'j}$, then he may be tempted to accept i even if $c_{ij} < \overline{c}_{ij} - \epsilon$ and reject i' even if $c_{i'j} \geq \overline{c}_{i'j} - \epsilon$, thus deviating from the protocol. In the transferable utility model, such an incentive issue does not arise, because the payments are constructed such that $p_i - \underline{c}_{ij} = p_{i'} - \underline{c}_{i'j}$ whenever a provider j is recommended to both customers i and i', so that the provider is essentially indifferent between the two if his costs for both are close to their lower support.

Proof of Proposition H.5. Let $\hat{I}$ be the set of customers who are matched in $\hat{x}$. As in the proof of Proposition H.4, with probability $1 - 3/n^2$, the matching returned by NTU Protocol 4 is ϵ-stable. When this happens, every customer $i \in \hat{I}$ obtains a utility of at least $\overline{b}_{i\hat{x}(i)} - \epsilon$, and the maximum possible utility she can obtain is under the protocol is $\overline{b}_{i\hat{x}(i)}$, since $\overline{b}_{ij}$ is the same for all providers $j \in S_i$ whom she might be matched to. For a customer $i \notin \hat{I}$, she is not offered the chance to be matched to any providers, and so obtains zero utility regardless of what she does. For each provider j of segment k, he can only be recommended to a customer in $\{i \in I : \hat{x}(i) \in J_k\}$, so his maximum possible utility is as given in (H.13). $\square$