Pricing Responses to Platform Leakage: Optimal Platform Design When Matches Are Irrevocable*

Isaías N Chaves

(Preliminary)

July 29, 2018

Abstract

I analyze a two-period model of a monopolist platform running a two-sided one-to-one matching market between workers and firms, where (i) agents can contract privately off the platform once they are matched (matches are irrevocable); (ii) they differ in how likely they are to want further business with a match (their durability); and (iii) firms have private information about their durability. I characterize the profit-maximizing matching and pricing policies. I show that the platform typically wants to impose a novel kind of distortion. In the first period, the platform wants to give firms either extremely high durability matches, or extremely low durability ones, but it would rather leave more firms unmatched than give any of them "medium" durability matches. This is in contrast to a seemingly related static model where a platform matches agents with privately known taste for quality to goods of different quality: in that case, the platform might price the bottom tier of agents out of the market, but it would never sell a low quality good before a medium quality one.

*Many thanks to Andy Skrzypacz for guidance throughout the project, and to Kyle Bagwell, Fuhito Kojima, Maciej Kotowski, Shengwu Li, Bobby Pakzad-Hurson, and seminar participants at Stanford for helpful comments.

1 Introduction

Many for-profit matching platforms face the problem that the matches they make are *irrevocable*: they can be given, but not taken away. Consider, for instance, a freelancing platform like UpWork (formerly Elance-oDesk), though the argument applies to growing online companies such as Handy and TaskRabbit, and to business-to-business platforms more generally. Once UpWork introduces a freelance developer to a firm that needs work on its website, the pair can exchange contact details, leave the platform, and contract privately again, while avoiding UpWork's fees. Practitioners call this *platform leakage*. If UpWork knew every agent's *durability* (how likely the firm was to demand further business from a worker), it could charge each agent up front for the expected value of the entire long-term relationship. But since the platform often cannot observe agents' durability, it must design the market in a way that accounts for agents' ability to strategically misrepresent themselves. Moreover, as it considers which firms should receive high durability workers, the platform must deal with the following tradeoff: more durable firms are willing to pay more for a durable match, but are more likely to leave the platform and side-deal with their match. Knowing this, a profit maximizing platform may well try to prevent these socially efficient matches between highly durable agents, in order to keep the largest number of interactions within its marketplace and paying fees.

This paper addresses three related questions. How should a platform react when agents can produce additional surplus without its consent, but it cannot observe which agents are more likely to produce this additional surplus? Can it design non-linear prices and menus of matchings to mitigate the effect on profits? Finally, what kinds of distortions will the platform try to induce as it tries to maximize revenue?

I analyze a two-period model of a monopolist platform running a two-sided one-to-one matching market, where (i) agents can contract privately off the platform once they are matched (matches are irrevocable); (ii) they differ in how likely they are to want further business with a match (their durability); and (iii) one side of the market (say, the firms in a freelancing platform like UpWork/Elance-oDesk) has private information about its durability. I characterize the revenue-maximizing mechanism that a platform should use in this environment. The optimal first-period matching pairs the highest durability firms with the highest durability workers assortatively, medium durability firms with the very lowest durability workers (also assortatively), and leaves low durability firms and medium durability workers un-matched.

I also explain how a model of privately known durability differs from a seemingly related Becker (1973) model of privately known taste for quality. While in a static model with privately known taste for quality, the platform would just price the bottom tier of firms and workers out of the market (thus selling a single premium product), with privately known durability, the platform not only prices out the bottom tier of firms, but it also sells two dramatically different products to the firms: one premium sub-market of high durability workers, at prices that only high durability firms would be willing to pay, and an inferior discount sub-market of very low durability workers, at prices that draw the middle-range firms away from the superior product but are too high for low durability firms.

I show that the platform creates two distinct distortions: mis-matching (giving agents lower-durability matches than desirable, even higher durability options are available) and under-matching (leaving agents un-matched even when there are available partners). Interestingly, I find that matching frictions (coordination costs incurred at the beginning of run long relationship) have contradictory effects on the amount of distortions: when frictions increase, the platform induces more under-matching distortions but lessens mis-matching distortions. For the case with two-sided incomplete information, monotonicity constraints can bind even in highly regular environments where all type distributions have increasing hazard rates.

To see how private information may be more severe on one side and fix ideas, consider an online platform matching firms to freelance workers. A straightforward way in which these one sided informational asymmetries could arise would be if firms had a number of different tasks they could have labor demand for, each requiring separate skills, while workers had skill portfolios that gave rise to different levels of flexibility across tasks. It would be easy for a firm to mis-represent the nature of its labor demand, while the worker's flexibility or skill would be observable from her CV. As a result, the firm would be privately informed about its need for repeat business from any given worker, but the worker's ability to create future business with any given firm would be common knowledge.

While I only focus here on one specific economic force, the paper takes an initial step toward four larger contributions. First, it introduces a substantive dynamic consideration into the models of profit-maximizing two-sided platforms, which have been overwhelmingly static. Second, the paper adds to a nascent literature on market-design (as opposed to purely pricing-based) considerations in models of profit-seeking two-sided platforms. Third, the paper introduces revenue-maximization into models of dynamic matching (the few recent papers on dynamic matching consider environments without money, and look only at implementing efficiency). Fourth, I show that, unlike the optimal auctions or good allocation problems, mapping from efficiency to revenue maximization in this matching environment is not simply a matter of replacing true values with virtual values. The reason is that the behavior of match marginal revenues and marginal surplus can differ dramatically as a function of match types, even in the regular case. Thus, unlike the auction design or trading environments, the form of the profit-maximizing rule in matching markets with incomplete information need not follow readily from the efficient rule.

The rest of the paper is as follows. In section 2, I give a brief description of relevant literature. Section 3 describes the economic environment. Section 4 formally states the platform's design problem, categorizes the different distortions the platform can use, and presents a version of virtual-surplus maximization for the case when monotonicity constraints are slack. The analysis is presented for the case when only one side of the market has private information about repeat business, but by and large the results in that section continue to hold when both sides are privately informed. Section 5 presents a full solution to the platform's problem for the case of one-sided private information, and describes issues with developing a full solution in the two-sided incomplete information case. Section 6 concludes and outlines steps for future research.

2 Related Literature

I briefly describe three strands of literature directly relevant to this proposal. The first, mostly in industrial organization, studies two-sided platforms. Prominent papers in this literature, such as Rochet and Tirole (2006) and Weyl (2010), analyze the optimal distribution of fees across the two sides of the market, and take the platform design as given. In particular, these papers model platform effects in reduced form. For instance, they both assume as a primitive of the model that agents' utilities are increasing in the number of cross-side agents that participate in the platform, and they do not model the matching structure inherent in the platform's problem. Rochet and Tirole (2006) had noted the irrevocable matches problem I am address in this paper: they write that "Buyers and suppliers may find each other and trade once on a [business-to-business] exchange, and then bypass the exchange altogether for future trade (p. 651)." However, to the best of my knowledge, no work has modeled this issue formally.

This paper also adds to a nascent literature on price discrimination in matching markets with transferable utility. These papers usually assume a continuum of agents on each side, all with private information on match values. Damiano and Li (2007) studied a static, one-to-one matching platform, like a dating website, and considered the different kinds of bunching that profit maximization could induce. In contrast, Gomes and Pavan (2015) study the technically challenging many-to-many problem, and show that cross-subsidization arises endogenously, with certain agents acting as inputs and some as consumers. The profit-maximizing allocation in their model may assign inefficiently small matching sets to certain agents, but unlike this paper, it features no gaps.

Finally, a very recent strand of literature concerns market design in fully dynamic environments. The two prominent examples are Akbarpour et al. (2014) and Baccara et al. (2015). This literature has not studied profit-maximization and screening, and is geared more closely to the kidney-exchange and adoption agency examples. In particular, neither model allows for transfers, so they do not address pricing issues. Akbarpour et al. (2014) study a continuous time matching model with stochastic arrivals of perishable agents. Agents have randomly drawn compatible matches in the existing pool. The authors evaluate the performance of different online algorithms, and conclude that waiting to thicken the market overwhelms any gains that could accrue from finding the best static matches at every point. Baccara et al. (2015) consider a discrete-time two-sided matching problem with discrete types where agents arrive sequentially and face a linear waiting cost. The authors solve explicitly for the first-best allocation and the decentralized allocation; they show that the decentralized allocation has inefficiently long queues.

3 Model

To fix ideas and terminology, I consider throughout the example of an online free-lancing platform that matches firms to workers, but the model applies more generally to other business-to-business platforms. In particular, similar incentives will apply so long as (i) producing match surplus requires communication, so the platform cannot anonymize the transaction, and once matched agents are able to coordinate on leaving the platform; (ii) agents can produce additional match surplus without the platform's consent; and (iii) there is some level of standardization across tasks performed in a match, such that horizontal differentiation across agents does not overwhelm this model's mechanism.

Players, Actions, Timeline A profit-maximizing monopoly platform matches workers w to firms f across two periods $t = 1, 2$ and charges them fees. Let F (W) denote the initial pool of firms (workers) at the beginning of period 1. Both F and W are finite sets.

Each firm $f \in F$ needs one $t = 1$ task, and an additional $t = 2$ task with probability q_f. Likewise, each worker $w \in W$ is available for one $t = 1$ task, and an additional $t = 2$ task with probability p_w. I assume nature draws q_f iid from a distribution G^F, which has as smooth density g^F that is strictly positive over $[0, 1]$. The tasks are indivisible and require only one worker. When a firm needs a second task or a worker is available for a second period, I say that the agent "survives." Denote by s_i and $\tilde{s}_i$, respectively, the indicators for whether agent i and i's $t = 1$ match partner survive. Note that the s_i's are drawn independently conditional on $\mathbf{p} = (p_w)_{w \in W}$ and $\mathbf{q} = (q_f)_{f \in F}$.

At the beginning of period 1, the platform can commit to any lottery over feasible matching policies and transfers across both periods, but agents can leave the platform before each match is announced and must be given incentives to stay. Figure 1 illustrates the timeline of events. After observing the platform's proposed mechanism, agents choose

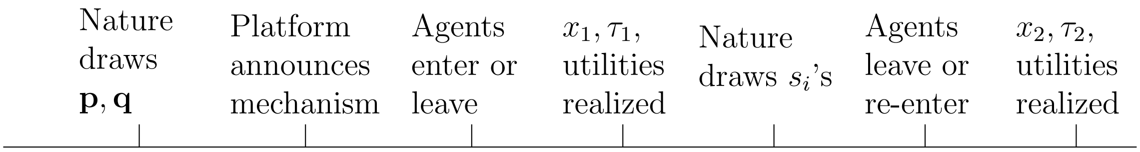

Figure 1: Timeline of Events

to enter the platform for the first period, or leave. For simplicity I assume that agents who decline at that point cannot re-enter in period 2. After observing the agents that accept the initial offer, the platform executes a lottery over matchings x_1 and a fee schedules τ_1. The fee schedule specifies a payment $\tau_1^i \in \mathbb{R}$ for each $i \in W \cup F$. Meanwhile, the first period matching is a function $x_1 : W \cup F \to F \cup W$ satisfying the usual properties: (i) $x_1^2(i) = i$, (ii) $x_1(i) \in W$ whenever $x_1(i) \neq i, i \in F$, and (iii) $x_1(i) \in F$ whenever $x_1(i) \neq i, i \in W$. $x_1(i) = i$ denotes that agent i is unmatched.

Next, having observed x_1 and τ_1 and whether they and their match partner survived, agents choose simultaneously whether to leave the platform before the second period allocation, or re-enter. The platform, observing which agents chose to remain on the platform, implements once again a (possibly random) matching x_2 and fee schedule τ_2. I assume that agents with $s_i = 0$ leave immediately, so they cannot be matched or charged, but to keep notation compact in the description of mechanism, I incorporate this by specifying that $x_2(i) = i, \tau_2^i = 0$ and i's utility is identically 0 whenever i perishes. Agent pairs with $s_i = \tilde{s}_i = 1$, on the other hand, can continue to transact off platform with no additional cost: that is, matches are irrevocable.

I make the following stark assumption:

Assumption (E)**.** There is enough entry in $t = 2$ that all survivors can be matched.

Assumption (E) has two effects. First, it simplifies the description of feasible matchings in $t = 2$. Under this assumption, the platform can always give an agent a new partner in $t = 2$ without preventing any other surviving agent from matching with its surviving $t = 1$ partner; that is, feasibility never requires the platform to "break" ongoing matches. Since, in addition, all matches other than its old partner generate the same utility to an agent, it suffices to describe a matching x_2 as a function from $W \cup F$ into $(O, N, \emptyset)$, where $x_2(i) = O$ means i is matched with its $t = 1$ partner, $x_2(i) = N$ means i gets any other partner, and $x_2(i) = \emptyset$ meaning i is un-matched in period 2.

Second, given the assumption that all agents still in the system can be matched at $t = 2$, it is harmless to ignore the profit from new arrivals. Since arrivals in the last period have no information rents, the platform can always fully extract their match surplus without affecting the incentives of $t = 1$ agents. Including the profit from these last-period arrivals therefore only adds a non-negative constant to the revenue from $t = 1$ agents' second period matches; this would not change the analysis, while tilting the results in favor of inefficiency by making it more desirable profitable for the platform to have $t = 1$ agents arrive unmatched to the second period.

Preferences Firms and workers have quasilinear preferences, with the non-monetary component given by the following. A match between a firm and a worker produces a gross surplus of $2v$ per period. Before this surplus can be produced, firms and workers incur a one-time matching friction $2c$ to form a new match with $c < v$. The net match surplus (gross of fees to the platform) is then split equally between firm and worker every period.[1] Formally, preferences for an agent i over matchings in periods 1 and 2 are as follows:

$$\mathbf{v}_1^i(x_1) = \begin{cases} v - c, & \text{if } x_1(i) \neq i \\ 0 & \text{otherwise} \end{cases}$$

$$\mathbf{v}_2^i(x_1, x_2) = \begin{cases} v, & \text{if } x_2(i) = x_1(i) \neq i \\ v - c, & \text{if } x_2(i) \notin \{x_1(i), i\} \\ 0 & \text{if } x_2(i) = i \end{cases}$$

Moreover, all agents' preferences are quasilinear in money.

I motivate the cost c with the fact that there are often frictions in the match-formation process. To produce surplus in each task requires the two parties to coordinate and undertake certain specific investments, such as training and explaining performance standards in the freelancing case. This friction ensures that, all else equal, it is more socially efficient to match long-term types together: two tasks done by the same worker-firm pair together produce more surplus than if the tasks were done by a firm with two different workers consecutively.

Outside Options There are two kinds of outside options in the model: the one that applies when an agent leaves unmatched (which the platform must match in an individually rational mechanism), and the one that applies when matched pair of agents choose to leave together. The former is specified as follows: an agent who leaves the platform in any period without a match receives a continuation utility of 0. Thus, an agent has a reservation value of 0 both at the beginning of period 1, at the beginning of period 2 when she survived but her $t = 1$ match partner has perished ($s_i = 1, \tilde{s}_i = 0$), and when she survived and was unmatched in $t = 1$ ($s_i = 1, x_1(i) = i$); this will be the relevant reservation value that the platform must deliver in order to satisfy an agent's individual rationality constraint after such histories.

The $t = 2$ outside option after histories with $s_i = \tilde{s}_i = 1$ (surviving agents with surviving partners) is more subtle. After such histories, a surviving pair (f, w) has an attractive alternative in the form of leaving the platform together, in which case they both receive utility v, their full share of the second period match surplus. If, on the other hand, either f or w leaves while their surviving partner is staying on the platform, the agent leaving only receives the continuation value of an un-matched agent, i.e., the

[1]This is possibly the result of Nash bargaining between them, but I do not model this explicitly. Moreover, if there were such a bargaining process, one would expect that the surplus split would depend on the composition of types within the match. One possible modification I do not explore is to treat the surplus split as a parameter that the platform can control by changing the space of contracts that firms and workers can sign. For instance, before the merger of Elance and oDesk, oDesk only gave firms warranties on work paid by the hour, while Elance would also guarantee fixed-rate "milestone" contracts.

Note also that the condition that firm-worker pairs split the surplus gross of fees is equivalent to assuming that the platform's fees do not affect the bargaining between firms and workers.

outside option is worth 0. Thus, what agent i can expect from leaving the platform before the $t = 2$ match depends on whether or not i's first-period partner is also choosing to leave.

How tight a constraint this outside option will be for the platform depends on whether we only consider *individual* or *pairwise* deviations from a mechanism's suggested path of play. The alternative where a surviving pair (f, w) leaves together is very attractive for both if they can exercise the outside option together, but is worthless otherwise. As explain in Section 4.1.1, it is essential to consider joint deviations in this model. Not only is this in line with a long tradition the matching literature (where pairwise stability with respect to reported preferences is a key desirable property of any matching mechanism), but without the possibility for joint, coordinated exit, one cannot possibly capture the idea that players can communicate in a way that hurts the platform's bottom line. In fact, I show in Section 4.1.1 that if one considers mechanisms that only rule out single-agent deviations at the re-entry stage, all information rents in the model vanish, but in the resulting mechanism surviving pairs play equilibrium strategies that are Pareto dominated for them.

Information It is common knowledge among all the agents and the platform that nature draws each q_f iid from the common distribution G^F, and that the s_i's are drawn independently conditional on $(\mathbf{p}, \mathbf{q})$.[2] Below I consider the case with one-sided incomplete information, where only f knows q_f at the start of the game, while the vector $\mathbf{p}$ is publicly observed. (In section 5.2, I briefly discuss what happens when, in addition, only w knows p_w at the beginning of the game, but is is common knowledge that $p_w \overset{iid}{\sim} G^W$, where G^W has the same regularity properties as G^F). The size of the initial pools W and F are common knowledge, but the active $t = 2$ pools are not. Agents know whether they and their match partner (if there was any) survived and whether they re-entered, but do not observe others' survivals or decisions to re-enter the platform; meanwhile, the platform observes re-entry decisions but does not observe survival outcomes.

4 Solving the Platform's Problem

4.1 Preliminaries

I assume that the platform can commit to any mechanism, and as discussed in the previous section, the platform can replicate any outcome that the agents are able to achieve off-platform with appropriate choices of allocations and transfers. Therefore, by the Revelation Principle and the argument in Myerson (1986), without loss of generality I restrict attention to direct revelation mechanisms such that (i) agents report types truthfully at the beginning; (ii) they choose to stay in the platform in both periods (the mechanism satisfies participation constraints); and (iii) they are not told anything beyond their own allocations as they happen.

Normally, (ii) would only require checking that each agent individually is promised a utility at least matching the value of the outside option (i.e., verifying individual rationality). However, in the present paper, surviving pairs have a very attractive outside option, but one that can only be exercised through a joint deviation. It therefore requires

[2] As shown in Myerson (1981) for the case of optimal auctions, asymmetries among bidders only increase a profit-maximizing designer's incentives to induce inefficient allocations. Therefore, I study only the same-side symmetric case, to focus attention on the new sources of inefficiency introduced by the current model.

		w In	*w* Out
f	In	U^f_{both}, U^w_{both}	U^f_{alone}, 0
f	Out	0, U^w_{alone}	v, v

Figure 2: Re-entry Game for Surviving Matches

additional work to ensure that agents do stay on the platform along the equilibrium path; as I explain in the following section, the mechanism should satisfy "pairwise rationality."

4.1.1 Pairwise Rationality

As described in Section 3, what agent i can expect from leaving the platform before the $t = 2$ match depends on whether or not i's first-period partner is also choosing to leave. Formally, any mechanism used by the platform induces a reduced-form simultaneous-move game between surviving partners (f, w) at the start of period 2, when both have to choose whether to re-enter the platform or leave; f's outside option along the equilibrium path of a mechanism depends on the strategy of w's strategy in that reduced-form game, and vice versa. This re-entry bi-matrix game is shown in Figure 2. Each agent in a surviving match receives v from leaving when her partner leaves, too, but 0 when her partner re-enters the platform instead (since in that case the agent is leaving without a match). Meanwhile, an agent's reduced-form utility from re-entering when her partner re-enters (leaves) is whatever the platform offers to returning agents with (without) returning partners. Let (U^f_{both}, U^w_{both}) be the promised utilities to the pair (f, w) when they both re-enter, while $(U^f_{alone}, U^w_{alone})$ are the promised utilities when either f or w re-enters alone.

Again note that by the arguments in Myerson (1986), without loss of generality we can restrict the search for optimal mechanisms to those where f and w choose to rejoin the platform along the equilibrium path. If one considers only single-agent deviations from such an equilibrium path (i.e., imposing only individual rationality on the mechanism), this requirement reduces to (In,In) being a Nash Equilibrium of the game in Figure 2, which holds whenever U^f_{both} and U^w_{both} are non-negative. Indeed, along the equilibrium path, one can condition on i's partner participating in the platform at $t = 2$, so if i left she would be leaving by herself.

Hence, when looking only at individual deviations, it would seem that the platform can extract the full efficient surplus. Since agents' $t = 1$ private information does not matter for their second-period preferences, and they have no new private information at the beginning of $t = 2$, the platform could always extract every agent's full $t = 2$ surplus: simply rematch each surviving pair, and charge each member of a returning pair an amount v; then find new matches and charge $v - c$ to all other returning agents (the platform can do this by Assumption (E)).[3] In that way, the platform collects all of $t = 2$ surplus, so $U^f_{both} = U^w_{both} = 0$ and (In, In) is indeed a Nash equilibrium of the re-entry game, i.e., no agent in a surviving pair wants to deviate individually from re-entering. Likewise, all agents arriving at $t = 2$ single are weakly indifferent between re-joining and staying off the platform, so assuming as usual that agents break ties in the mechanism designer's favor, these agents would also choose to rejoin. Finally, looking at the form of firm's $t = 1$ expected utility as a function of its true and reported types in (3), it is clear

[3] Note also that the platform can always distinguish these groups from each other.

that when an agent's expected value upon surviving to $t = 2$ is 0 (as is the case with these $t = 2$ policies) every firm's true type q_f drops out of its indirect utility. Therefore, the platform can now use the efficient $t = 1$ matching policy and choose transfers to fully extract the first period surplus of each agent without violating incentives for truthful revelation (this is possible since each agent's surplus now depends only on the assigned allocation and not on any hidden type). Altogether, the platform here would extract full surplus while respecting incentive compatibility and inducing entry in every period, and firms' private information becomes irrelevant.

While the argument in the previous two paragraphs would be in the spirit of partial implementation (i.e., some equilibrium of the mechanism exists that implements the desired social choice function), the equilibrium it corresponds to is highly dubious. First, note that the platform cannot discriminate between an agent whose partner perished and one whose partner survived but left, since these two could costlessly mimic each other. Therefore, since an agent with $s_i = 1, \tilde{s}_i = 0$ who re-enters can have non-transfer expected utility of at most $v - c$, (Out,Out) is also an equilibrium of the reduced form game unless the platform pays out strictly more than c to all agents who arrive at the second stage without partners. More importantly, the Nash Equilibrium (Out, Out) Pareto dominates (from the pair's perspective) all other equilibria unless the platform also re-matches all surviving pairs and pays them non-negative amounts in period 2. The practicality of a mechanism that relies on agents playing a Pareto-dominated Nash equilibrium is dubious at best— all the more so in the present case, where the motivation for irrevocable matches is that agents can communicate and transact on their own once introduced by the platform.

To capture the fact that matched pairs can communicate and leave together, and to avoid the implausible behavioral implications of looking only at individual deviations, one must therefore consider *pairwise* deviations. There are many different modeling assumptions that could be made here, corresponding to different communication technologies between matched pairs and on whether they can transfer utility to each other. In this paper, I model joint deviations from (In, In) by requiring that the mechanism satisfy a "pairwise rationality" condition. Concretely, I impose the following assumption on behavior, which will translate into an additional constraint for the platform:

Assumption (Pairwise Rationality (PR))**.** Surviving pairs will choose a pure Pareto dominant Nash Equilibrium of the Re-entry game, if one exists.

Therefore, if the platform wants to induce surviving pairs to re-enter, it must insure that (In, In) be not only a Nash equilibrium of the re-entry game, but that no other pure equilibria Pareto dominate it.[4] Imposing (PR) directly as constraint on the platform's problem leads to some considerable difficulties, since (as I show in the appendix), it implies a *system of mutually exclusive systems of inequalities*, many of which describe constraint sets that are not closed. Each mutually exclusive set captures the conditions under which another pure profile is an equilibrium, and the associated restrictions that prevent that alternative equilibrium from Pareto-dominating (In,In). Fortunately, this difficulty can be bypassed by conjecturing that the only relevant case is the one where (Out,Out) and (In,In) are the only Nash equilibria, but (In,In) weakly dominates. This will happen whenever $\min\{U^f_{both}, U^w_{both}\} \geq v$, which is easy to incorporate in the platform's problem as the agent-by-agent restriction that *both surviving agents simultaneously have*

[4]Assumption (PR) is practically the same as requiring that (In, In) be a coalition-proof Nash equilibrium.

outside option v. The conjecture is then proven true by showing that in the solution to the platform's problem imposing only $\min\{U^f_{both}, U^w_{both}\} \geq v$ and dropping the other (PR) constraints, the former binds while all the other systems of inequalities, corresponding to cases that are excluded by the present one, are violated.

4.2 Formulating the Platform's Problem

I introduce some notation that helps state the mechanism design problem compactly. Let $\mathbf{h}$ and h_f denote, respectively, the entire history observed by the platform and the history observed by agent i at the beginning of period 2; also let $\mathcal{H}$ denote the space of possible platform period 2 histories. $\mathbf{h}$ includes the realized allocations x_1 and τ_1 and identities of re-entering agents, reports about realized survivals $\mathbf{s} = (s_i)_{i \in W \cup F}$, as well as the history that the platform had observed at the beginning of period 1, i.e., the vector of reported and publicly observable survival types $(\mathbf{p}, \mathbf{q})$. h_i contains only realized $x_1(i), \tau_1^i, s_i, \tilde{s}_i$, re-entry decisions by i and its partner, publicly observable survival types, and, if i's side of the market has private information, i's survival type. Also, note that after the educated guess about the form of (PR) described in the previous section, both the (PR) and individual rationality constraints at the start of $t = 2$ are equivalent to delivering the following history-dependent individual reservation values to each agent:

$$r(h_i) = \begin{cases} v & \text{if } x_1(i) \neq i,\ s_i = \tilde{s}_i = 1 \\ 0 & \text{otherwise} \end{cases}.$$

A (possibly stochastic) direct mechanism is then a pair of functions (χ_1, τ_1) mapping $(\mathbf{p}, \mathbf{q})$ into lotteries over feasible matchings and transfers in the first period, and functions (χ_2, τ_2) mapping $\mathbf{h}$ into lotteries of feasible matchings and transfers in the second period. Period 1 transfer choices are functions $\tau_1 : [0,1]^{|W|} \times [0,1]^{|F|} \to \Delta(\mathbb{R}^{|W|} \times \mathbb{R}^{|F|})$. To describe first-period allocations, note that by the Birkhoff-von-Neumann Theorem, any χ_1 that satisfies, for all $\mathbf{p}, \mathbf{q}$, the natural constraints

$$\begin{aligned} 0 \leq \chi_1^{f,w}(\mathbf{p}, \mathbf{q}) \leq 1, \forall f, w \\ 0 \leq \sum_{w \in W} \chi_1^{f,w}(\mathbf{p}, \mathbf{q}) \leq 1, \forall f \\ 0 \leq \sum_{f \in F} \chi_1^{f,w}(\mathbf{p}, \mathbf{q}) \leq 1, \forall w \end{aligned} \tag{χ_1-F}$$

indeed corresponds to a lottery over feasible matchings x_1, where the (f, w) entry of $\chi_1(\mathbf{p}, \mathbf{q})$ is the probability that firm f is assigned to worker w. Thus, for the initial allocation it suffices to specify $\chi_1 : [0,1]^{|W|} \times [0,1]^{|F|} \to M_{[0,1]}(|W|, |F|)$ satisfying (χ_1-F), where $M_{[0,1]}(|W|, |F|)$ denotes a $|F| \times |W|$-matrix with entries in $[0,1]$. For future reference, when parts of χ_1 or τ_1's argument are omitted, this denotes that they have been integrated out: $\chi_1^{f,w}(q_f) = \mathbb{E}_{\mathbf{p},\mathbf{q}_{-f}|q_f}[\chi_1^{f,w}(\mathbf{p}, \mathbf{q}_{-f}, q_f)] = \mathbb{E}_{\mathbf{p},\mathbf{q}_{-f}}[\chi_1^{f,w}(\mathbf{p}, \mathbf{q}_{-f}, q_f)]$ (the equality follows by independence), and so forth.

In principle, the period 2 stochastic allocation $\chi_2(\mathbf{h})$ should be an object with the same structure as χ_1, satisfying feasibility constraints akin to (χ_1-F). However, as discussed in the previous section, by Assumption (E), the platform can always give an agent a new partner in $t = 2$ without preventing any other surviving agent from matching with its surviving $t = 1$ partner; that is, feasibility never requires the platform to "break" ongoing matches. Since, in addition, all matches other than its old partner generate the same utility to an agent, it suffices to specify a stochastic allocation χ_2 as a set of pairs $(\chi_2^i(\mathbf{h}|P), \chi_2^i(\mathbf{h}|N))$ for each i, where the two elements are the probabilities that after history $\mathbf{h}$ i is matched to its previous partner or a new partner, respectively. Thus χ_2 must satisfy, for each $\mathbf{h}$ and each i ,

$$\begin{aligned} &0 \le \chi_2^i(\mathbf{h}|\cdot) \le 1 \\ &0 \le \chi_2^i(\mathbf{h}|N) + \chi_2^i(\mathbf{h}|P) \le 1 \\ &\chi_2^i(\mathbf{h}|P) = 0 && \text{if } \tilde{s}_i = 0 \text{ or } x_1(i) = i \\ &\chi_2^i(\mathbf{h}|P) = \chi_2^i(\mathbf{h}|N) = 0 && \text{if } s_i = 0 \end{aligned} \tag{χ_2-F}$$

Meanwhile, any $\tau_2 : \mathcal{H} \to \Delta(\mathbb{R}^{|W|} \times \mathbb{R}^{|F|})$ that satisfies $\tau_2^i(\mathbf{h}) = 0$ if $s_i = 0$ describes a physically feasible $t = 2$ transfer rule.

With the above notation, Firm f's continuation value when re-entering the platform at $t = 2$ is

$$U_2^f(h_f) = \mathbb{E}_{\mathbf{h}|h_f}^{\chi_1,\tau_1}\left[v \cdot \chi_2^f(\mathbf{h}|P) + (v-c) \cdot \chi_2^f(\mathbf{h}|N) - \tau_2^f(\mathbf{h})\right]. \tag{1}$$

Meanwhile, its expected continuation utility conditional on survival, from the perspective of period 1, is

$$V_S^f(q_f) = \mathbb{E}_{\mathbf{h}|s_f=1,q_f}^{\chi_1,\tau_1}\left[\max\{U_2^f(h_f), r(h_f)\}\right]. \tag{2}$$

(Expressions for V_S^w, U_2^w are identical). The difference between U_2^f and V_S^f is that in the former, f conditions on all of h_f, which includes not just q_f, s_f, but also the realizations of $t = 1$ allocations and its match partner's survival state. The superscript on $\mathbb{E}$ highlights that the measure over $\mathbf{h}$ depends endogenously on the platform's allocation rule, in two ways: first, because x_1 is part of $\mathbf{h}$ and is drawn according to χ_1; second, because the realized allocation $x_1(i), \tau_1^i$ can "leak" information to i about $-i$'s types, so that i's $t = 2$ posterior beliefs over $\mathbf{p}$ are no longer given by the product measure $(G^F)^{|F|}$.

f's expected utility from the mechanism when other agents are truthful but it reports $\hat{q}_f$ when it has true type q_f is then

$$U_1^f(\hat{q}_f|q_f) = \sum_{w \in W} \chi_1^{f,w}(\hat{q}_f)(v-c) + q_f V_S^f(\hat{q}_f) - \tau_1^f(\hat{q}_f). \tag{3}$$

Let $U_1^f(q_f) \equiv U_1^f(q_f|q_f)$, and make identical definitions for workers. With this notation, a mechanism induces entry (is individually rational and pairwise rational for surviving match pairs) if it is feasible[5] and

$$U_1^f(q_f) \ge 0,\, U_1^w(p_w) \ge 0 \,\forall f, w, q_f, p_w \tag{IR-1}$$

$$U_2^i(h_i) \ge r(h_i) \forall i, h_i \text{ such that } s_i = 1 \tag{IR-2, PR}$$

The following standard lemma simplifies the characterization of incentive compatibility.

Lemma 1. *A direct mechanism $(\chi_1, \chi_2, \tau_1, \tau_2)$ is incentive compatible iff*

$$U_1^f(q_f) = U_1(0) + \int_0^{q_f} V_S^f(\tilde{q})d\tilde{q} \quad \forall f \in F, q_f \in [0,1] \tag{ICFOC-F}$$

$$V_S^f(q_f) \text{ non-decreasing} \tag{M-F}$$

[5] It is unnecessary to check individual rationality at histories where i does not survive.

The monotonicity requirement is on the composite object V_S, rather than on the allocations $\chi_1^f(q_f)$ or $\chi_2^f(\mathbf{h})$ and so on, as would be the case in a typical linear-trading scenario. The platform, by choosing τ_2 appropriately, could induce truth-telling with non-monotone or even decreasing allocations.

The platform's problem is therefore

$$\max_{\chi_1,\chi_2,\tau_1,\tau_2} \mathbb{E}\left[\sum_{i\in W\cup F} \tau_1^i(\mathbf{p},\mathbf{q}) + \tau_2^i(\mathbf{h})\right]$$

$$\begin{aligned} &\text{subject to} \\ &\qquad (\chi_1\text{-F}), (\chi_2\text{-F}), \\ &\qquad (\text{ICFOC-F}), (\text{M-F}), \\ &\qquad \text{and} \\ &\qquad (\text{IR-1}), (\text{IR-2, PR}). \end{aligned} \tag{P}$$

I refer the above program without (M-F) as the *relaxed problem*.[6] Also, I refer to (P) with the additional constraint that the platform does not charge workers ($\tau_1^w(\mathbf{p},\mathbf{q}) = \tau_2^w(\mathbf{h}) = 0$ for all $(w,\mathbf{p},\mathbf{q},\mathbf{h})$ as is often true of freelancing or job-search platforms) as the "one-sided profit" case, studied in section 5.1.

As usual, from (3), monotonicity and (IR-1) holding for the lowest type imply that (IR-1) holds for all higher types. Hence, increasing τ_1^f, τ_1^w by a constant until (IR-1) binds at the bottom preserves incentives and increases revenue, and we can replace (IR-1) with $U_1^f(0) = U_1^w(0) = 0$.

Using lemma 1, the expression (3), and iterated expectations on $\mathbb{E}[\tau_1^f(\mathbf{p},\mathbf{q})]$, we can eliminate τ_1^i from the platform's objective. Then after the usual integration by parts and applying iterated expectations to $\mathbb{E}[\chi_1^{f,w}(q_f)] = \mathbb{E}[\mathbb{E}_{\mathbf{p},\mathbf{q}_{-f}|q_f}[\chi_1^{f,w}(\mathbf{p},\mathbf{q}_{-f},q_f)]]$, f's contribution to expected revenue becomes

$$\mathbb{E}\left[\sum_{w\in W} \chi_1^{f,w}(\mathbf{p},\mathbf{q})(v-c) + J^F(q_f)V_S^f(q_f) + \tau_2^f(\mathbf{h})\right]$$

By (χ_2-F), $\mathbb{E}^{\chi_1,\tau_1}_{\mathbf{h}|s_f=0,q_f}[\tau_2^f] = 0$ (the platform cannot charge agents who expire). So using iterated expectations again, rewrite each of the $\mathbb{E}[\tau_2^f(\mathbf{h})]$ terms in the above display as $\mathbb{E}\left[q_f\mathbb{E}^{\chi_1,\tau_1}_{\mathbf{h}|s_f=1,q_f}[\tau_2^f(\mathbf{h})]\right]$. Then substitute the expression for V_S^f in (2) into the above display,[7] and group together the terms with τ_2^f to obtain the following expression for

[6] To obtain the platform's problem in the two-sided incomplete information case, it suffices to add the additional worker-side incentive constraints (ICFOC-W) and (M-W).

$$U_1^w(p_f) = U_1(0) + \int_0^{p_w} V_S^w(\tilde{p})d\tilde{p} \quad \forall w\in W, p_w\in[0,1] \tag{ICFOC-W}$$

$$V_S^w(q_w) \text{ non-decreasing.} \tag{M-W}$$

[7]Note that $V_S^f(q_f) = \mathbb{E}^{\chi_1,\tau_1}_{\mathbf{h}|s_f=1,q_f}\left[U_2^f(h_f)\right]$ in an individually rational mechanism.

revenues from f:

$$R_f \equiv \mathbb{E}\left[\begin{array}{c}\sum_{w\in W}\chi_1^{f,w}(\mathbf{p},\mathbf{q})(v-c)+J^F(q_f)\mathbb{E}^{\chi_1,\tau_1}_{\mathbf{h}|s_f=1,q_f}\left[v\cdot\chi_2^f(\mathbf{h}|P)+(v-c)\cdot\chi_2^f(\mathbf{h}|N)\right]\\ +\quad \frac{1}{h^F(q_f)}\mathbb{E}^{\chi_1,\tau_1}_{\mathbf{h}|s_f=1,q_f}[\tau_2^f(\mathbf{h})]\end{array}\right]. \tag{4}$$

These expressions used $q_f - J^F(q_f) = 1/h^F(q_f)$.

To find the worker-side contributions to revenue, note that, since workers have no private information, all of their $t = 1$ IR constraints must bind in the optimal mechanism. Therefore, substituting into the platform's objective and using iterated expectations,

$$\begin{aligned} R_w &\equiv \mathbb{E}\left[\begin{array}{c}\sum_{f\in F}\chi_1^{f,w}(\mathbf{p},\mathbf{q})(v-F)+p_w\mathbb{E}^{\chi_1,\tau_1}_{\mathbf{h}|s_w=1,p_w}\left[v\cdot\chi_2^w(\mathbf{h}|P)+(v-c)\cdot\chi_2^w(\mathbf{h}|N)-\tau_2^w(\mathbf{h})\right]\\ +\quad p_w\mathbb{E}^{\chi_1,\tau_1}_{\mathbf{h}|s_w=1,p_w}[\tau_2^w(\mathbf{h})]\end{array}\right]\\ &= \mathbb{E}\left[\sum_{f\in F}\chi_1^{f,w}(\mathbf{p},\mathbf{q})(v-F)+p_w\mathbb{E}^{\chi_1,\tau_1}_{\mathbf{h}|s_w=1,p_w}\left[v\cdot\chi_2^w(\mathbf{h}|P)+(v-c)\cdot\chi_2^w(\mathbf{h}|N)\right]\right]. \end{aligned} \tag{5}$$

4.3 Solution to the Relaxed Problem

To make further progress, by the standard approach, I conjecture that (M-F) is slack at the solution to (P), solve the relaxed problem, and then verify conditions under which the solutions to (P) and its relaxation are identical.

Dropping monotonicity constraints allows an easy characterization of $t = 2$ allocations and transfers.

Proposition 1. *In any solution to the relaxed problem,*

1. *Every surviving agent whose partner perished receives a new match :*$\chi_2^i(\mathbf{h}|N) = 1$ *for* $s_i = 1, \tilde{s}_i = 0$.
2. *Every surviving pair is re-matched on-platform:* $\chi_2^i(\mathbf{h}|O) = 1$ *for* $s_i = 1, \tilde{s}_i = 1$.
3. *Transfers satisfy* $\tau_2^i(\mathbf{h}) = v - c$ *if* $s_i = 1, \tilde{s}_i = 0$*, and 0 otherwise.*

Allocations and transfers at all other histories are pinned down by (χ_2-F)*).*

Proof. From (4) and (5), it is clear that, holding χ_1, τ_1 fixed, the platform wants to increase all χ_2^i's and τ_2^i's as much as possible, subject to the constraints. By the discussion in section 4.1, problem (P) is simply

$$\max_{\chi_1,\chi_2,\tau_1,\tau_2} \sum_{i\in W\cup F} R_i$$

subject to

$$(\chi_1\text{-F}), (\chi_2\text{-F}), (\text{M-F}),\ \text{and (IR-2, PR)}.$$

with R_i given by (4) and (5). After removing (M-F), there are no other cross-type or cross-history constraints, so for any χ_2, the platform will optimally increase each $\tau_2^i(\mathbf{h})$ pointwise until (IR-2, PR) binds for at h_i's with $s_i = 1$. Therefore,

$$\mathbb{E}^{\chi_1,\tau_1}_{\mathbf{h}|h_i}[\tau_2^i(\mathbf{h})] = \mathbb{E}^{\chi_1,\tau_1}_{\mathbf{h}|h_i}\left[v\cdot\chi_2^i(\mathbf{h}|P)+(v-c)\cdot\chi_2^i(\mathbf{h}|N)\right] - r(h_i).$$

Look first at $i \in F$. By iterated expectations and the above display,

$$\mathbb{E}^{\chi_1,\tau_1}_{\mathbf{h}|q_f,s_f=1}[\tau^f_2(\mathbf{h})] = \mathbb{E}^{\chi_1,\tau_1}_{h_i|q_f,s_f=1}\left[\mathbb{E}^{\chi_1,\tau_1}_{\mathbf{h}|h_i}\left[v \cdot \chi^i_2(\mathbf{h}|P) + (v-c)\cdot \chi^i_2(\mathbf{h}|N)\right] - r(h_i)\right].$$

Apply the same procedure for $i \in W$ and substitute into R_i in (4) and (5) to conclude that R_i is strictly increasing in $u(h_i)\cdot\chi^i_2(\mathbf{h})$. Moreover, each term $v\cdot\chi^i_2(\mathbf{h}|P)+(v-c)\cdot\chi^i_2(\mathbf{h}|N)$ is feasibly maximized by setting $\chi^i_2(\mathbf{h}|O) = 1$ for histories with $s_i = \tilde{s}_i = 1$ and $\chi^i_2(\mathbf{h}|N) = 1$ for histories with $s_i = 1, \tilde{s}_i = 0$. (This respects (χ_2-F), and is physically feasible by Assumption (E)).

Finally, note that the platform can replace any $\tau^i_2(\mathbf{h})$ that varies with h_{-i} with its expectation $\mathbb{E}^{\chi_1,\tau_1}_{\mathbf{h}|h_i}[\tau^i_2(\mathbf{h})]$ without loss. □

Proposition 1 also simplifies the verification of (M-F).

Lemma 2. *The solution to the relaxed problem satisfies (M-F) iff the average expected match types* $\mathbb{E}_{\mathbf{p}_{-w},\mathbf{q}}[\sum_{f\in F} q_f \chi^{w,f}_1(\mathbf{p}_{-f}, p_w, \mathbf{q})]$ *are monotone non-decreasing in* q_f.

The straightforward proof, which I omit, consists merely in substituting χ_2, τ_2 from Proposition 1 into (2) and using the same iterated expectation/conditional independence tricks as in the proof of Theorem 1.

Proposition 1 implies that the platform cannot gain from treating different types q_f differently in $t = 2$. This is somewhat surprising, since agents who are more likely to survive to $t = 2$ value the future more, so it would have seemed possible that the platform could use this continuation utility to incentivize agents to its own benefit. The proposition also means that the platform chooses a "sequentially" efficient $t = 2$ rule , i.e., efficient taking as given the $t = 1$ matches and the set of surviving agents, and it fully extracts this conditional continuation surplus. Of course, this does not mean the platform implements the efficient rule, since as I show below, it will distort the set of initial matches away from the social optimum.

An obvious but important consequence of full surplus extraction in period 2 is that, in the current model, all $t = 2$ participation/rational entry constraints bind. Meanwhile, in dynamic mechanism design problems with outside options, a common intermediate step involves showing that, for all periods other than the first, individual rationality constraints are slack. While the $t = 2$ participation constraint here comes in part from pairwise rationality, it has the same form as an individual rationality constraint, so the source of binding participation constraints must lie elsewhere. The difference, it turns out, stems from the possibility of agents' "dying" in this model. The usual approach holds that the designer can always ask agents to post a bond up front, to be repaid only if the agent stays in the mechanism for the duration. If the posted bonds and repayments are type-independent, incentive constraints will still hold while the principal becomes weakly better off. This line of argument meets two obstacles in the present model. First, the feasibility of such schemes depends on unrestricted transfers. In contrast, here agents can only be charged (or paid) in $t = 2$ when they survive, which introduces a restriction on the feasible τ_2 that will defeat the standard argument. Second, it is impossible to devise type-independent bond repayment schemes here, since different types value $t = 2$ transfers differently.

4.4 First-period Matching Policies

Before studying the platform's optimal policy, let consider first what the ex post efficient $t = 1$ would be (the second period efficient policies are just as in Proposition 1). A brief derivation in the appendix shows the social planner chooses probability shares χ_1 so as to solve the following optimal assignment problem, as in Koopmans and Beckmann (1957) and Shapley and Shubik (1971):

$$\max_x \sum_{w\in W}\sum_{f\in F} x^{w,f} a^E(p_w, q_f)$$
$$\text{s.t.} \quad \begin{aligned} & x^{w,f} \in [0,1] \\ & \sum_{w\in W} x^{w,f} \in [0,1] \\ & \sum_{f\in F} x^{w,f} \in [0,1] \end{aligned}$$

where

$$a^E(p_w, q_f) = 2[1 - p_w q_f](v-c) + 2p_w q_f v = 2(v-c) + 2p_w q_f v, \tag{6}$$

the planner's weight on $x^{w,f}$, is a firm-worker specific marginal match surplus. It is easy to prove, looking at the $t = 1$ objective $\sum_{w\in W, f\in F} x^{w,f} a^E(p_w, q_f)$, that the planner always makes as many matches as possible, and always matches workers and firms assortatively according to their probabilities of survival, starting from the top and continuing until it runs out of unmatched agents on either side.

With the $t = 2$ allocations in hand, the platform's choice of χ_1 reduces to a static problem parallel to the one solved by the social planner. Specifically, the platform finds the optimal $t = 1$ matching rule by solving a fictitious optimal assignment problem, where the entries of the payoff matrix entries are given by "marginal revenues" rather than marginal surplus. The following definition is therefore central to the rest of the paper.

Definition (Marginal Revenues)**.**

$$a^F(p_w, q_f) = (1 - p_w q_f)(v - c) + J^F(q_f) p_w v$$

is the *firm-side marginal revenue from the match* (w, f). Analogously, call

$$a^W(p_w, q_f) = a_1^W(p_w, q_f) = (1 - p_w q_f)(v-c) + p_w q_f v \tag{7}$$

the *worker-side marginal revenue from the match* (w, f), and $a(p_w, q_f) \equiv a^W(p_w, q_f) + a^F(p_w, q_f)$ the *full marginal revenue from the match* (w, f).

I use the term "marginal revenue" here in analogy to the discussion of optimal auctions in Bulow and Roberts (1989); it is not exactly these authors call marginal revenue (which refers in their paper to the virtual type function J^F), but it plays an analogous role.

With this definition in hand, I can state the first main result, which further simplifies the relaxed problem.

Theorem 1. *In any solution to the relaxed problem,* $\chi_1(\mathbf{p}, \mathbf{q})$ *solves the following linear* *program:*

$$
\max_x \sum_{w\in W}\sum_{f\in F} x^{w,f} a(p_w, q_f)
$$

subject to

$$
\begin{aligned}
& x^{w,f} \in [0,1] \\
& \sum_{w\in W} x^{w,f} \in [0,1] \\
& \sum_{f\in F} x^{w,f} \in [0,1]
\end{aligned}
\tag{LP}
$$

Note 1. Note that $x^{i,i}$, the probability of i going unmatched the first period, enters implicitly in this objective with a coefficient of 0.

I relegate the proof to the appendix. It hinges on rewriting R_f with the help of Proposition 2, and it relies heavily on the independence of private types and of survival outcomes (conditional on types).

Remark 1 (Comparison to Efficient Benchmark)**.** To gain intuition on the form of a, and on the effects that information rents and irrevocable matches have on the platform's incentives, it is useful to compare $a(p_w, q_f)$ to the coefficient on the (w, f) match in the objective of a fully informed social planner.

Comparing a^E to $2a^W$ (the coefficient used by an informed monopolist) leads to two broad conclusions. First, the perfectly informed monopolist executes the same matches as the social planner. Match irrevocability on its own generates no allocative inefficiencies in this model, and it misaligns the platform's incentives and social objectives only when combined with private information about survival or repeat business.

A second, more subtle conclusion is that *agents' information rents only factor in through their valuation for a continuing match.* After all, it is not surprising that somewhere in the platform's objective true types get replaced by virtual types. What is surprising, however, is that, even though a firm's privately known q_f affects her valuation of all $t = 2$ outcomes, the only part of the information rent that matters to the platform's $t = 1$ choices is the part corresponding to matches where both partners survive and leave the market together. A rough intuition for why this is the case is as follows. Whenever a firm's partner perishes, the platform gets to extract its full $t = 2$ surplus $v - c$, so the firm's expected utility in the second period on that event is 0 regardless of its type.

4.5 Three Distinct Kinds of Distortion

As mentioned above, the solution to the social planner's first period problem is to match the agents assortatively, starting from the top. I describe in detail the exact properties of a^E that lead to this solution and how a can fail to satisfy them; this will elucidate the nature of the distortions generated by the platform's profit maximization.

The form of the efficient solution relies on three properties of $a^E(p_w, q_f)$:

(A1) it is supermodular in p_w, q_f;

(A2) it is non-negative (in particular, it is always larger than the weight on $x^{f,f}$ and $x^{w,w}$); and,

(a) Cross-Matching (b) Under-matching (c) Mis-matching

Figure 3: Three Forms of Distortion

(A3) it is strictly increasing in p_w and q_f.

The failure of each property leads to a distinct kind of distortion, which I now discuss. For the discussion below, note that a^W, the marginal revenue contribution from the side with no private information, has all the "nice" properties [A1-A3] of a^E. Therefore, violations of [A1-A3] always come from a^F.

Failure of (A1): Cross-matching distortions. Supermodularity ensures that all agents not matched to themselves get matched assortatively by rank (as in Becker (1973)). For a pair of workers $\bar{p} > \underline{p}$ and firms $\bar{q} > \underline{q}$, does the designer want to match $\bar{p}$ to $\underline{q}$ *and* $\underline{p}$ to $\bar{q}$? If a were supermodular, flipping this match to $(\underline{p}, \underline{q}), (\bar{p}, \bar{q})$ increases the objective. (Note, however, that without (A2) and (A3), neither $(\underline{p}, \underline{q})$ nor $(\bar{p}, \bar{q})$ are necessarily part of the planner's solution). Therefore, failures of (A1) might motivate the platform to create *cross-matching* distortions, where it matches pairs anti-assortatively. This is illustrated in Figure 3(a). Squares denote firms, circles are workers, and arrows denote that two agents are matched, and both sides are arranged in decreasing order of type from the top.

Note that a matching policy with anti-assortative segments would tend to would violate monotonicity.

The cross-partial of a in p_w, q_f is

$$a_{pq} = [J^{F\prime}(q_f)v - (v - c)] + c,$$

where the first term corresponds to a^F_{pq} and the second to a^W_{pq}. Thus a could violate (A1) even in the regular case (Myerson, 1981) where type densities have full support and J's are non-decreasing.

The second and third properties rule out different kinds of "gaps" in the matching.

Failure of (A2): Under-matching distortions Given that the platform would never form matches (w, f) where $a(p_w, q_f) < 0$, if the weight function were supermodular and increasing, but ever negative, then the matching where all agents are paired assortatively from the top could be dominated by one where all matches (w, f) with $a(p_w, q_f) < 0$ are dropped. Profit maximization therefore incentivizes the platform to create *under-matching* distortions: it fails to create socially beneficial matches, even when these are feasible. Figure 3(b) illustrates under-matching distortions. Depending on what types go unmatched, this kind of distortion could lead χ_1 to violate (M-F).

Since virtual types J^F can be negative, a will in general violate (A2) for some types.

Failure of (A3): Mis-matching distortions A different kind of gap appears if (A3) is violated, so that that a is decreasing, in say, p. Now, conditional on a set of matches

$x_1(F \setminus \{f\})$ for other firms, the platform always matches f, if at all, to

$$\tilde{w} \in \arg \max_{w \in W \setminus \{x_1(F \setminus \{f\})\}} a(p_w, q_f).$$

So even starting from the planner's "assortative from the top allocation" for all firms ranked higher than f, $\tilde{w}$ will not be the worker with the same rank as f, but strictly lower. Suppose firms are the short side of the market: $|F| < |W|$. Then, even if a is supermodular and non-negative, the matching where all $|F|$ firms are matched in order to the top $|F|$ workers could be defeated by one where the lowest firm were matched to the lowest worker. The failure of (A3) thus incentivizes the platform to create *mis-matching* distortions, where an agent is matched to someone lower (or higher) than the agent on the other side of the market with the same rank. Figure 3(c) illustrates mis-matching distortions and how they are different from cross-matching distortions (from Figure 3(a)). In the latter, the second firm from the top is not matched to its rank partner, but given the top firm's allocation, no better match is available for it. In Figure 3(c), however, the lower-q firm is matched to worker 3, even though 2 is available.

It is clear from the picture that mis-matching distortions would make the monotonicity constraints bind on the side that getting "skipped." However, by Lemma 2, they can be consistent with monotonicity (and therefore optimal) in the one-sided incomplete information case, since higher durability firms always receive higher durability matches ex post.

Crucially, property (A3) will also fail to hold quite generally. Take a^F, for instance. Then

$$\begin{aligned} a^F_p &= J^F(q_f)v - q_f(v-c) \\ a^F_q &= p_w[J^{F\prime}(q_f)v - (v-c)] \end{aligned}$$

Since, again, J^F can be negative, a^F can be fail to be increasing in p_w. Troublingly, it can fail to be increasing in q_f even in the regular case (Myerson, 1981) of increasing virtual types. Since full marginal revenue is $a = a^W + a^F$, it can also violate (A3).

The above discussion makes clear that the usual regularity condition (virtual types are increasing) falls quite short in this current model, leading to different solutions in (P) and the relaxed problem. For the remainder, I impose the stronger, but still natural assumption on types that will give the problem sufficient structure:

Assumption (H)**.** The hazard rates $h^F = g^F/(1-G^F)$ is monotone non-decreasing.

Assumption (H) is sufficient to rule out cross-matching distortions.

Lemma 3. *If (H) holds, the platform never uses cross-matching distortions in the relaxed problem. Moreover, a^F and a are strictly increasing in q_f.*

Proof. I show the first claim, since the second follows immediately. It suffices to check the sign of the cross-partial of a. I show the result for the two-sided incomplete-information, two-sided profit case, the other cases being essentially identical.

$$a_{pq} = [J^{F\prime}(q_f)v - (v-c)] + c \geq 2c > 0$$

The inequality follows from $J^{F\prime}(q_f) = 1 + h^{F\prime}(q_f)/(h^F(q_f))^2 \geq 1$, which (H) guarantees. □

Remark 2 (Comparative Statics)**.** Using Assumption (H), one finds also that the matching friction has interesting effects on the platform's incentives. On the one hand, a_p, a_q, and $a_{p,q}$ (subscripts indicate partial derivatives) are all increasing in c, so higher match frictions reduce the platform's propensity to induce mis-matching and cross-matching. On the other, a is decreasing in c, so as match frictions increase, the platform is more likely to induce under-matching.

The intuition for why under-matching goes down as c increases is straightforward: expected surplus decreases in c, and accordingly the expected revenue that the platform can capture must decrease. Thus, it is more likely that the platform finds it unprofitable to match a given agent at all when c is high. To see why mis-matching distortions decrease in c, recall the trade-off mentioned in the introduction between a firm's value for an initial match and its incentive to leave. When the platform matches a firm "properly" (with a high-survival worker) at least it can capture a share of that firm's value for an irrevocable surviving match, after giving up some information rents. This benefit from matching a firm well does not depend on c. On the other hand, by matching a firm to a good worker, the platform loses the full surplus it could extract from the firm by giving it a low-survival worker who perishes and then re-matching it to someone else in $t = 2$. This cost from proper matching is increasing in c, so of course as c grows the platform is less likely to mis-match.

5 Full Solution for One-sided Incomplete Information

Assignment problems like (LP) seldom have explicit solutions, but by using the special structure of the assignment matrix $a(p_w, q_f)$, I can give a full solution for the case where workers' probabilities of survival are publicly observed. Below, I show derivations for the case where the platform collects fees from both sides of the market, pointing out at the end how not charging workers would change the analysis.

I make a few convenient definitions, which will help frame the analysis of properties [A1-A3] in a^W, a^F, and a.

Definition 1. Let $b^F(q_f) = J^F(p_f)v - q_f(v - c)$. Say a firm f is *pro-assortative* if $b^F(q_f) + q_f c \geq 0$ (and anti-assortative) otherwise, and let F^+ denote the pro-assortative subset of F, with $F^- = F \setminus F^+$. Let $P(q) \equiv \{p : a(p, q) \geq 0\}$ be the set of *viable matches* for a firm.

The terminology stems from the fact that $a(p_w, q_f)$, the marginal revenue from matching f to w, is increasing in match type p_w iff $b^F(q_f) + q_f c \geq 0$. Thus, looking at the case with one-sided incomplete information, whenever q_f is pro-assortative, the platform wants to match f to the most durable worker possible, insofar as it wants to match it at all. Seen a different way, (A3) always holds for pro-assortative types.

The following elementary lemma then characterizes the functions b^F and sets $P(q)$.

Lemma 4. *Under assumption (H),*

1. *b^F is strictly increasing, so there is a unique $\hat{q}$ such that f is pro-assortative iff $q_f \geq \hat{q}$, and anti-assortative otherwise.*

2. *Moreover, $P(q)$ is always of threshold form $\{p \in [0,1] | p \leq \tilde{p}(q)\}$, where the highest viable match $\tilde{p}(q)$ has $\tilde{p}(q_f) < 1$ only for anti-assortative firms, and is strictly increasing whenever that is the case.*

Identical results apply to b^W, $Q(p)$.

Proof. Under assumption (H), b^F is strictly increasing in q_f: $b^{F\prime}(q_f) = J^{F\prime}(q_f)v - (v-c) \geq v - (v-c) > 0$, where I use $J^{F\prime}(q_f) \geq 1$ (see the proof of Lemma 3). So $b^F(q) + qc$ is increasing.

Suppose first that $a(p, q_f) \geq 0$ for all $p \in [0,1]$; then $P(q_f)$ has threshold structure with $\tilde{p}(q_F) = 1$. So assume instead that $a(p, q_f) = 0$ has an interior solution. Recall that $a^F = (v-c) + p_w(b^F(q_f) + q_f c)$, so if q_f is pro-assortative, then $a(p, q_f) > 0$ for all p. Thus in the case with an interior solution, $b^F(q_f) + qc < 0$ and $a(p, q_f)$ is strictly decreasing in p, which implies that $p \mapsto a(p, q_f)$ crosses 0 only once and from above and the solution is unique when it exists. Therefore $\tilde{p}(q)$ is well-defined. Moreover, by the implicit function theorem, when $\tilde{p}(q) \in (0,1)$,

$$\tilde{p}'(q_f) = -\frac{a_q(\tilde{p}(q_f), q_f)}{a_p(\tilde{p}(q_f), q_f)}. \tag{8}$$

Since q_f is anti-assortative in this case, from part 1 of the lemma we know the denominator above is negative, and by Lemma 3, a_q in the numerator is strictly positive. Therefore $\tilde{p}$ is strictly increasing whenever it is in $(0,1)$. □

Remark 3. Refusing to match pairs (w, f) with $a(p_w, q_f) < 0$ is analogous to how a revenue-maximizing auctioneer never sells an object to a bidder with negative marginal revenue. Indeed, $\tilde{p}^{-1}(p_w)$ acts more or less as a w-specific optimal reserve price: the platform only matches firm f to w if $q_f \geq \tilde{p}^{-1}(p_w)$, and $\tilde{p}^{-1}(p_w)$ solves $a(p_w, \tilde{p}^{-1}(p_w)) = 0$, in the same way that an optimal reserve price solves $MR = 0$ in the auction setting. To see this, note that $\tilde{p}$ is invertible, since it is increasing. And since a is increasing in q and $a(p_w, \tilde{p}^{-1}(p_w)) = 0$ by definition, w is viable for f if and only if $q_f \geq \tilde{p}^{-1}(p_w)$.

With the help of Lemma 4, I can give a full description of the optimal allocation when only firms are privately informed.

Theorem 2. *The solution to (LP) sets $x^{w,f} = 1$ if (w, f) are matched by the following algorithm (and $x^{w,f} = 0$ otherwise):*

1. *In decreasing order of q, sequentially match pro-assortative firms F^+ to the highest un-matched worker. Continue until either exhausting F^+ or W, whichever happens first. If W is exhausted, stop, and leave all other firms unmatched.*

2. *If F^+ is exhausted first, then, in decreasing order of q, sequentially match each anti-assortative firm f to the highest viable un-matched worker (the available worker with highest p_w subject to $p_w \leq \tilde{p}(q_f)$). If there are no viable un-matched workers left, leave all subsequent firms in F^- un-matched. Continue until either running out of firms, or running out viable un-matched workers.*

3. *In* increasing *order of q, sequentially rematch all firms F^- that were matched in Step 2 to the* lowest *un-matched worker, and continue until exhausting all firms in F^- matched at Step 2.*

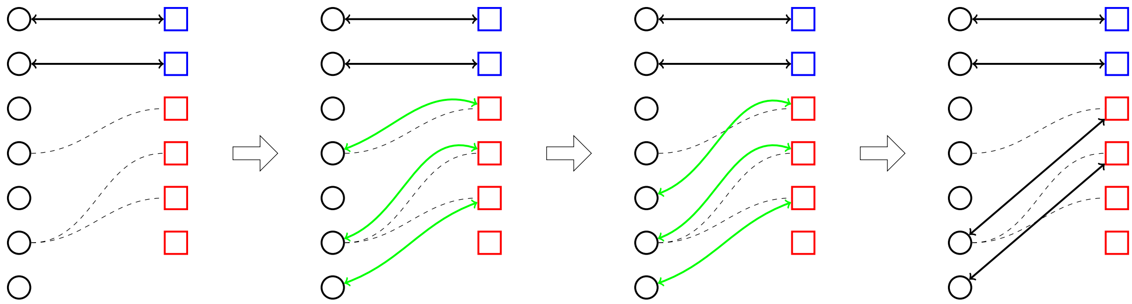

Figure 4: Algorithm Steps 1-4. Circles are workers and squares are firms, sorted so higher types are on top. Red indicates anti-assortative, blue pro-assortative. Dotted lines indicate the highest viable match for a firm. Solid lines green lines denote temporary matches, while solid black lines with denote final matches.

4. *Discard the lowest matched firm in F^-, and shift all matched F^- down by one rank. Repeat this operation while the objective in (LP) increases.*

A mechanism consisting of $\chi_1^{w,f} = x^{w,f}$ for all w, f and (χ_2, τ_2) as in Proposition 1, satisfies (M-F)*, so together with $\tau_1(q)$ given by* (ICFOC-F)*, it is indeed the solution to (P).*

Before proving the theorem, it is helpful to look at a worked example of how the algorithm operates. Figure 4 illustrates this for a case with six workers and five firms. Circles represent workers and squares are firms. Red indicates anti-assortative (firms 3-6, from the top), blue pro-assortative (firms 1 and 2). Agents are sorted so higher types are on top. Dotted lines indicate the highest viable match for a firm, so 4 and 5 share worker 6 as the highest viable match (either one generates negative marginal revenue with workers 5 and above), while none of the workers is viable for firm 6. Solid green lines denote temporary matches, while solid black lines denote final matches.

During Step 1, pro-assortative firms get matched to the workers for their own rank, and these matches are permanent. In Step 2, firm 3 is first tentatively matched to worker 4 and firm 4 is tentatively matched to worker 6. Then, since worker 6 is unavailable, firm 5 is tentatively matched to worker 7, the highest worker viable for firm 5 with the exception of worker 6. Firm 6, which has no viable matches in this pool of workers, receives no tentative match. In Step 3, firm 3 is re-matched with the *lowest* available worker (worker 5). Since this move stacks anti-assortative firms without gaps at the bottom, the algorithm transitions to Step 4. Finally, in Step 4, firm 5's match is dropped, and firms 3 and 4 are shuffled downwards one rank, which in this example is assumed to increase (LP).[8]

Remark 4 (Distortions and Implications for Product Design)**.** The right-most diagram in Figure 4 also illustrates the categorization of distortions in section 4.5. First, no solid

[8]Note that, given how I have drawn the highest viable match curves, a second iteration of Step 4 would never be profitable. Assume it were, so that the algorithm drops firm 4 and leaves firm 3 matched with worker 7. But then, worker 6 would be un-matched, and since she is viable for firm 4, (LP) would increase yet again by matching her to firm 4, who had just been dropped. This would produce a cross-matching distortion (firms 3 and 4 matched to workers 7 and 6 respectively), which by the supermodularity of a is never optimal; the platform could improve (LP) further by switching the places of firms 3 and 4. But this would only return us to the matching that prevailed at the end of Step 3 of the algorithm. Since the first iteration of Step 4 was (LP)-profitable by assumption, this is a contradiction.

arrows cross, so the platform uses no cross-matching distortions. Second, firms 5 and 6 gets matched to themselves, even though workers 3-5 are available in the final match, and both firms would produce positive surplus with either worker. That is, the platform induces "under-matching" distortions. Third, firm 3 get matched to worker 6 and firm 4 gets matched to worker 7, rather than to workers 3 and 4 respectively, so the platform induces mis-matching.

Mis-matching and under-matching distortions together produce allocations with a "gap" on the uninformed side. These distortions have a useful interpretation in terms of the platform's product design. With privately known durability, the platform not only prices out the bottom tier of firms, but it also sells two dramatically different products to the firms: one premium sub-market of high durability workers, at prices that only high durability firms would be willing to pay, and an inferior discount sub-market of very low durability workers, at prices that draw the middle-range firms away from the superior product but are too high for low durability firms.

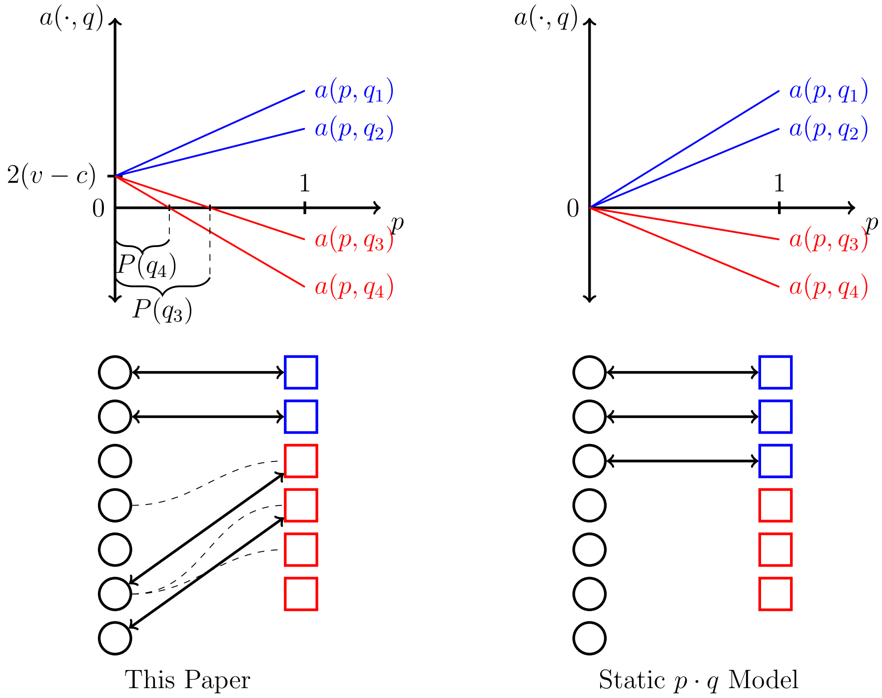

Figure 5: Comparison of marginal revenue curves and optimal matchings between a model of (i) heterogeneous durability and irrevocable matches, and (ii) heterogeneous quality and taste for quality (optimal auctions with heterogeneous goods and unit demand, or price discrimination in static Becker assortative matching model with match preferences $u = p \cdot q$).

Remark 5 (Comparison to Price Discrimination in Static One-to-One Allocation Problems with Privately Known Taste for Quality)**.** To understand the economic forces that generate the more novel "mis-matching" distortions and polarization on the uninformed side, consider the contrast between the current dynamic matching model and two related static problems. One is an optimal auction of heterogeneous goods, with unit-demand bidders that have preferences that are linear in quality and value and private iid values. To keep notation consistent, again let p_w denote good w's quality and q_f denote bidder

f's value for quality. Similar derivations to the ones in Theorem 1 show that marginal revenue from assigning good w to bidder f is given by $a^{static}(p_w, q_f) = p_w J^F(q_f)$. The second static case is that of a monopolist platform matching firms and workers one-to-one, but with preferences over matches as in Becker (1973) or Damiano and Li (2007): both firms f and workers w obtain non-monetary utility $u(p_w, q_f) = p_w \cdot q_f$ from a match, and once again only firms have private iid types. Again, with similar techniques one can show that marginal revenue from matching f and w is simply $a^{static}(p_w, q_f) = p_w(J^F(q_f) + q_f)$ (I use the same notation for both static cases, since they are qualitatively identical for the purpose of this discussion). As this remark will show, the new incentives in the present model arise from the fact that matches in this paper are not heterogeneous in quality (as in the static models), but heterogeneous in *durability*.

Since marginal revenue in the irrevocable matches case can be written

$$a(p_w, q_f) = 2(v - c) + p_w(b^F(q_f) + q_f c),$$

then from Lemma 4 it follows that, whenever the hazard rate h^F is non-decreasing, marginal revenue in this paper behaves similarly in many respects to marginal revenue in the static cases: it is strictly increasing in q_f, supermodular in (p_w, q_f), and increasing in p_w for all types q_f above some cutoff (and strictly decreasing for q_f below the cutoff).[9] These properties imply that, in both the current model and the static model, decreasing q rotates the marginal revenue curves $p \mapsto a^{static}(p, q)$ and $p \mapsto a(p, q)$ clockwise in (p, a) space *around the point at which they cross the y-axis.*

Figure 5 illustrates marginal revenue curves in (p, q) space for the present case (left) and the static cases (right). As shown in the diagram, the only substantive difference between my case and the standard static one is that in the former, the curves $p \mapsto a(p, q)$ are "shifted" upwards relative to the latter. They intersect the y-axis in (p, a)-space at a common point strictly above $(0, 0)$. In the latter case, all $p \mapsto a^{static}(p, q)$ pass through the origin in (p, a) space.

This fact has two important consequences:

- There exist anti-assortative firm types q for whom $p \mapsto a(p, q)$ is decreasing, but who nonetheless generate positive marginal revenue with some worker types p. In other words, the viable matches $P(q_f) = \{p : a(p, q_f) \geq 0\}$ for firms f with $a_p(p, q_f) < 0$ can be non-empty.

- Relatedly, the curves $p \mapsto a(p, q)$ intersect the x-axis in (p, a) space at different points, so the set of viable matches for an anti-assortative firm f changes as q_f changes.

In contrast, in both static scenarios, any bidder/firm with anti-assortative type has an empty viable match set, since marginal revenue for these types lies entirely below the x-axis in (p, a) space. Therefore, in a static model with privately known taste for quality, the platform would just price the bottom tier of firms and workers out of the market (thus selling a single premium product), and it would never create the two highly differentiated products from the privately known durability case. The bottom of row of Figure 5 illustrates optimal policies in the two cases.

To find the real culprit for the new distortions I present, it suffices to know why $p \mapsto a(p, q)$ has a non-zero y-intercept, i.e., where does the term $2(v - c)$ that is independent

[9] Extending the paper's terminology in the natural way, call a bidder or man type q_f "anti-assortative" if $p \mapsto a^{static}(p, q_f)$ is decreasing.

of p come from? The upward shift in marginal revenue is due to the fact that workers in my model are heterogeneous not in quality, but in *durability*. Specifically, changing a worker's durability changes a firm's $t = 2$ expected utility from matching to that worker, but it does not affect the utility that f receives from that match in period 1. Therefore, the $t = 1$ interaction adds a non-zero term to marginal revenue that will not vary with p. Moreover, since the firm's $t = 1$ preferences are also independent of its own private information, the y-intercept is the same for all firm types. In contrast, changing a worker's quality would change the utility a firm receives in every period that it remains matched to that worker. Seen another way, when matches are heterogeneous in quality, as in the Becker-style model, the worker's type always interacts with the matching allocation χ_1 in the utility function. Since match type comes into firm utility linearly, this will imply that marginal revenue curves in (p, a) space must go through the origin.[10]

The economic intuition for why this "intercept" in the utility function leads the platform to offer two sharply differentiated products (as opposed to a single premium product) is as follows. In a static, heterogeneous-quality $p \cdot q$ model, giving very low p matches to "medium" q firms (the high anti-assortative types) puts some downward pressure on the prices that the platform can charge to the high q firms. Moreover, the low p matches are worth very little to the medium firms, so the platform generates little additional profit from mis-matching the medium firms, compared to leaving them unmatched. To be able to extract significant surplus from the medium q firms, the platform would have to give them medium p matches, but this would put an inordinate amount of downward pressure on the prices that the platform can charge high q (pro-assortative) firms. Meanwhile, in a heterogeneous durability model, the very low p matches are worth a non-trivial amount to medium q firms, since they still enjoy the full $v - c$ in utility from being matched to those workers in the first period. At the same time, giving these very low p matches to the medium q firms generates relatively little downward pressure on the prices that can be charged to high q firms, since these firms have stronger preferences for durability. Therefore, the platform can gain from offering the second inferior product in a way that would be impossible in the static $p \cdot q$ case.

To further emphasize the different effects of heterogeneity in durability as opposed to quality, consider what would happen if we re-wrote the proof for Theorem 1 with $t = 1$ non-monetary utilities set to zero. Similar derivations would yield a coefficient of $p_w b^F(q_f)$ on $\chi_1^{f,w}(\mathbf{p}, \mathbf{q})$, so that once again marginal revenue curves in (p, a) space will go through the origin for all firm types, and no new distortions would arise relative to the static case.[11] Yet another way of seeing the durability-quality distinction is by thinking of an alternative preference structure for a static "quality"-based market that could produce the same kind of marginal revenue diagrams of a two period "durability"-based market. In particular, the upward shift of marginal revenue in (p, a) space would happen if the non-monetary utility of a firm with private type q had the following form:

$$u(p, q) = \begin{cases} a + p \cdot q & \text{if matched to worker of type } p \\ 0 & \text{otherwise} \end{cases}$$

[10]This last argument also explains why my own model with $t = 1$ non-monetary utility set to zero has all marginal revenue curves go through the origin in (p, a) space: the $t = 1$ allocation always interacts with worker type in expected $t = 2$ utility.

[11]Note that, in this hypothetical, one still needs to assume preferences are quasilinear over a $t = 1$ transfer, or else the model transforms into one of mechanism design without transfers at the information revelation stage.

with a some a positive constant.

Remark 6. Clearly, if workers have private information, it will be hard for the kind of allocation in Figure 4 to be incentive-compatible.

The allocation of workers 3 and 4 is not monotone in their own type reports–it increases as the they over-report past the types of workers 1 and 2, but also increases as they under-report part the types of workers 5 and 6. Of course, (M-W) depends on *average* match quality over all configurations, but the failure of χ_1 to be increasing in own type suggests monotonicity will bind for a large range of examples.

Proof of Theorem 2. Let $F^* = \{f : \exists w \text{ s.t. } x^{w,f} = 1\}$ denote the firms matched to workers by the algorithm. I make two claims. First, I claim the configuration of F^* dictated by $x^{f,w}$,

- pro-assortative types matched assortatively starting from the top, leaving no gaps;
- anti-assortative types matched assortatively starting from the bottom, leaving no gaps;

achieves higher value than any match that only switches which workers firms in F^* are matched with, but neither matches additional firms in $F \setminus F^*$ nor leaves any firms in F^* unmatched. That is, there is no *rearrangement* of F^* can yield an improvement to (LP). Second, I claim that F^* is in fact the optimal subset of firms to rearrange in this way.

Since firms in F^+ have marginal revenues strictly increasing in match type and those in F^- have marginal revenues strictly decreasing in match type, it is clear from the linear structure of (LP) that the optimal arrangement of any subset has all pro-assortative firms matched high as possible and all anti-assortative firms as low as possible. Moreover, by Lemma 3, a is supermodular, so conditional on firms from F^+ being matched toward the top, it is optimal to match them assortatively in $\mathbf{p}$ and $\mathbf{q}$. The same applies to firms in F^- matched at the bottom.

The argument above in fact establishes that the configuration from Steps 1 and 4 is the optimal rearrangement of any fixed set $F' \subset F$ of matched firms. It remains to show that a) the algorithm in fact produces this arrangment of matches within F^*, and b), F^* is in fact the optimal subset of firms to match.

To see a), note that Step 1 matches pro-assortative firms assortatively from the top by construction, while Step 3 and 4 never swap the order of matches for anti-assortative firms and always pack these firms at the bottom of the worker distribution. Therefore, the algorithm will produce the optimal re-arrangement if Step 2 always matches higher-q anti-assortative firms to higher-p workers. Part 2 of Lemma 4 ensures this is the case. Suppose otherwise that during Step 2, some firm with type $\underline{q}$ receives a tentative match of type $\bar{p}$ that is higher than the tentative match $\underline{p}$ given to some firm with type $\bar{q} > \underline{q}$. Then, since $\tilde{p}(q)$ is strictly increasing, and $\bar{p}$ was viable for $\underline{q}$, it must be viable for $\bar{q}$. Moreover, since Step 2 processes firms in decreasing order and $\bar{p}$ was un-matched when $\underline{q}$ was processed, it must also have been available when firm $\underline{q}$ was processed. Therefore, worker $\underline{p}$ was not the highest un-matched worker viable for firm $\bar{p}$, a contradiction.

It remains to be shown that F^* is indeed the optimal subset of firms to arrange optimally. For a fixed size $k \leq |F|$ of firms to have matched, the optimal subset is clearly the k highest, since a is increasing in q_f. Moreover, since Steps 1 and 2 matches firms in decreasing q order, F^* contains exactly the $|F^*|$ highest-q firms. The question then becomes, is there a differently sized subset to optimally arrange? By the construction in

Step 3 and the fact that $\tilde{p}(q)$ is strictly increasing (Lemma 4), the optimal subset can be no larger than F^*, since Step 3 finds the largest possible number of viable matches. To show no smaller subset improves on F^*, let r be the number of repetitions of Step 4. I claim that for any $(\mathbf{p}, \mathbf{q})$, and any optimally arranged subset $F' \subset F$, there is a unique critical number $r^*(F')$ such that for $r < r^*(F')$, the objective in (LP) increases with an additional repetition, and decreases for $r \geq r^*(F')$. Clearly, F^* is given by exactly $r^*(\tilde{F})$ repetitions of Step 4, where $\tilde{F}$ is the subset of firms identified by Steps 1 and 2, so if the claim is true, no smaller subset can improve on F^*. Each repetition of Step 4 has two effects on the (LP) objective. First, since the firms being re-shuffled are anti-assortative, there are gains from the firms getting less durable workers. Second, since by construction each of the firms matched by Step 2 are viable with the lowest rank worker (their Step 2 match satisfied $p \leq \tilde{p}(q_f)$, so $p_{|W|}$ will also), there are losses from un-matching this bottom firm. By supermodularity of a, the gains from shuffling each an anti-assortative firm down from w to $w+1$ strictly decrease the higher the firm's type. $a_p(p, q_f) < 0$ for $f \in F^-$ and $a_{pq} > 0$ always, so by the Fundamental Theorem of Calculus the gain $a(p_w, q) - a(p_{w-1}, q)$ must be decreasing. In addition, with each additional repetition, these re-shuffling gains are coming from fewer and fewer firms. Meanwhile, the losses from un-matching the currently lowest firm strictly increase with each repetition, since the firms were sorted in increasing order of q, and $a(p_{|W|}, q)$ (the loss from unmatching the bottom firm) is increasing in q. Therefore, the losses eventually outweigh the gains; call this point r^*.

To check monotonicity, note that firms within F^+ that get matched are matched assortatively within themselves, and likewise for F^-. Moreover, by Lemma 4, firms in F^+ always have higher types than firms in F^-. Hence, for every vector of opponent types, $\mathbf{q}_{-f}$, ex post match quality is increasing in q_f, and by Lemma 2, (M-F) is satisfied. □

The intuition behind the optimality of this matching algorithm is straightforward: Steps 1 and 3 find the optimal workers matches for any subset of matched firms, while Steps 2 and 4 find the best such subset of firms to match.

5.1 One-sided Profit

When the platform does not charge the uninformed worker side,

$$a(p_w, q_f) = a^F(p_w, q_f) = (v - c) + p_w b^F(q_f) \tag{9}$$

Therefore, full marginal revenue is increasing in p_w iff $b^F(q_f) \geq 0$. Thus pro-assortative firms are now those with b^F positive, and the definition of $\hat{q}$ changes slightly. But one obtains the same partition of firms into those whose marginal revenue increases in match and those for whom it decreases. Since a^F is always increasing in q, once again (8) and the implicit function theorem imply that $P(q)$ has threshold structure, with $\tilde{p}(q)$ strictly increasing and less than 1 only for anti-assortative firms. Suitably modified versions of Lemma 4 and Theorem 2 will then apply, and the optimal $t = 1$ allocation has the same structure as under two-sided profit.

Since $c \geq 0$, the set of firms put in F^-, and therefore matched to the bottom of the worker distribution, will expand once the platform cannot charge workers. Likewise, since a^W is always positive, the locus $a(\tilde{p}(q), q) = 0$ shifts downwards under one-sided profit. Altogether, both unmatching and mis-matching distortions will increase. The intuition is straightforward: when the platform is considering under-matching distortions under two-sided profit, it also loses out on the $t = 1$ fees that it would obtain from workers. Since workers have no information rents, from Remark 1, we know that the platform would want to allocate assortatively from the top, which mutes the platform's incentives to distort the market, relative to the one-sided profit case.

5.2 Two-sided Incomplete Information

When workers' survival probabilities are their private information, Proposition 1 and Theorem 1 hold mostly without modification, with the latter using the following new expression for worker-side marginal revenue (mirror image of a^F):

$$a^W(p_w, q_f) = (1 - p_w q_f)(v - c) + J^W(p_w) q_f v, \tag{10}$$

where $J^W(p) = p - \frac{1-G^W(p)}{g^W(p)}$. Notably, the first part of Lemma 3 continues to hold: the platform never uses cross-matching distortions in the relaxed problem.

The similarities with the one-sided incomplete information case end there, however. Lemma 2 must now include the requirement that $\mathbb{E}_{\mathbf{p}_w, \mathbf{q}_{-f}}[\sum_{w \in W} p_w \chi_1^{w,f}(\mathbf{p}_w, \mathbf{q}_{-f}, q_f)]$ also be monotone non-decreasing. But as mentioned in Remark 6, the kind of allocation found by Theorem 2 will probably not be incentive compatible when workers are privately informed, since it will lead to non-monotone average matches types for workers. The standard solution to failures of monotonicity would be ironing, as in Myerson (1981). Myerson proposed taking virtual values in quantile space, integrating them cumulatively, and taking the subgradient of its convex closure as the ironed virtual values; these will be increasing in true values by construction, so maximizing ironed virtual surplus will satisfy monotonicity. The issue with applying this technique here is that the very definition of the "virtual value" object (the marginal revenue term) depended on the slack monotonicity constraints. Unlike the optimal auctions setting, where the virtual surplus formula can be derived simply using the envelope formula ICFOC, my expression for marginal revenues (Definition 4.4) is only valid when when the planner achieves full conditional surplus extraction in period 2 (Proposition 1). I have been unable to prove that $t = 2$ efficiency and binding (IR-2) constraints would be optimal even when monotonicity binds. Thus, at best ironing would find the optimal mechanism subject to $t = 2$ conditional full surplus extraction.

In fact, the situation is even more dire, since with privately informed workers it is not even clear how to solve (LP) in the first place. For one thing, whether a firm's full marginal revenue is increasing in match type will now depend on who the firm is being matched to, since a^W is not increasing in q_f only for workers with $b^W(p_w) = J^W(p_w)v - p_w(v - c) \geq 0$. The proof of Theorem 2 relied heavily on Lemma 4, which ensured that a platform always wanted to match a given firm as high (or as low) as possible. With two-sided private information, depending on the remaining unmatched workers, the platform may be tempted to either shuffle a firm upwards or downwards. Second, the threshold structure of $P(q)$ ensured that, to match anti-assortative firms, one only needed to find the highest viable match since all matches below that one would also be viable. Now that marginal revenue can be non-monotonic in match type, $a(p, q) = 0$ can have multiple solutions in p, some increasing and some decreasing in q.

6 Conclusion and Next Steps

This paper has described a monopoly platform's optimal response to the problem of irrevocable matches. I solve for the optimal mechanism and illustrate that the platform induces two kinds of distortions: under-matching (matching agents to themselves even when there is excess supply on the opposite side) and mis-matching (matching agents to partners of lower rank than themselves). In general, the platform creates allocations with a "gap" on the uninformed side: mid-level agents on the uninformed side of the market are left unmatched, while agents on the informed side are matched to both extremes of the opposite-side distribution. In other words, the platform reacts to irrevocable matches and unobservable durability by offering two sharply differentiated products to the firms—a premium product of highly durable workers, at a a premium price, and a discount, inferior product of low durability firms—and pricing out the lower durability firms. This is in contrast to static models of price discrimination in platforms, which broadly speaking are models of heterogeneous quality and taste for quality. In that latter case, the platform would only create a single premium product and price out the bottom tier of firms.

To conclude, I discuss open questions and avenues for future research. The first major remaining issue is to understand what the "typical" allocation in the two-sided incomplete information case. As discussed above, it will likely involve ironing and randomization, but does there exist a clean algorithmic characterization, as in the one-sided case? Second, I have assumed throughout that there is no scarcity in period 2, which breaks part of the dynamic link across matches in both periods. What is the optimal mechanism in the case where Assumption (E) fails? The main complication this induces, since $\chi_2^i = 1$ for all agents is no longer feasible, is that $t = 2$ allocations are no longer easy to determine. Certainly the platform will match as many agents as possible and will not randomize conditional on $\mathbf{h}$, since the individual contributions to revenue are linearly increasing in χ_2^i. The coefficients on χ_2^i will depend on both i's type and χ_1, though, and the convenient separation between the two decisions is lost. Third, the current results lack a discussion of implementation. Is there a straightforward indirect mechanism that achieves the optimal allocation and transfers? Fourth, this paper has ignored value heterogeneity. If values are observable, a re-weighted version of the same framework will apply more or less unchanged, but how to induce truthful revelation of both probability of survival and match value is a much more difficult problem.

References

Akbarpour, Mohammad, Shengwu Li, and Shayan Oveis Gharan (2014), "Dynamic matching market design." Working paper, Stanford University.

Baccara, Mariagiovanna, SangMok Lee, and Leeat Yariv (2015), "Optimal dynamic matching." Working paper, California Insitute of Technology.

Becker, Gary S (1973), "A theory of marriage: Part i." *The Journal of Political Economy*, 813–846.

Bulow, Jeremy and John Roberts (1989), "The simple economics of optimal auctions." *The Journal of Political Economy*, 1060–1090.

Damiano, Ettore and Hao Li (2007), "Price discrimination and efficient matching." *Economic Theory*, 30, 243–263.

Gomes, Renato and Alessandro Pavan (2015), "Many-to-many matching and price-discrimination." Working paper, northwestern university.

Koopmans, Tjalling C and Martin Beckmann (1957), "Assignment problems and the location of economic activities." *Econometrica*, 53–76.

Myerson, Roger B (1981), "Optimal auction design." *Mathematics of operations research*, 6, 58–73.

Myerson, Roger B (1986), "Multistage games with communication." *Econometrica*, 323–358.

Rochet, Jean-Charles and Jean Tirole (2006), "Two-sided markets: a progress report." *The RAND Journal of Economics*, 37, 645–667.

Shapley, Lloyd S and Martin Shubik (1971), "The assignment game i: The core." *International Journal of game theory*, 1, 111–130.

Weyl, E Glen (2010), "A price theory of multi-sided platforms." *American Economic Review*, 1642–1672.

Appendix

Derivation of surviving pairs' reservation values. The proof proceeds by splitting the entry game into cases.
Case I: (Out, Out) is a Nash equilibrum This case applies whenever $U^f_{alone} \vee U^w_{alone} \leq v$. For (In, In) to be a Pareto-undominated Nash Equilibrium then requires (in addition to $U^f_{both} \wedge U^w_{both} \geq 0$) one of two conditions:

- $U^f_{both} \wedge U^w_{both} \geq v$, so that (In, In) weakly Pareto dominates (Out, Out); or
- $U^f_{both} \wedge U^w_{both} < v$ and $U^f_{both} \vee U^w_{both} > v$, so that (In, In) and (Out, Out) are Pareto-unranked.

Note that, under case I, (Out, Out) weakly Pareto Dominates (In, Out) and (Out, In), so we only need to check that (In,In) against (Out,Out).
Case II: (Out, Out) is not Nash, but (In, Out) is.
This case applies whenever $U^f_{alone} > v$ and $U^w_{both} = 0$. Then (In, Out) is Nash and Pareto dominates (In, In) unless $U^f_{both} \geq 0$.
Case III: (Out, Out) is not Nash, but (Out, In) is.
By symmetry with Case II, this case applies whenever $U^w_{alone} > v$ and $U^f_{both} = 0$. Then (Out,In) is Nash and Pareto dominates (In, In) unless $U^w_{both} \geq 0$.

Note that the inequalities for Cases II and III are mutually exclusive, so (In, In) cannot be Pareto undominated when both (In, Out) and (Out, In) are Nash but (Out, Out) is not. Thus the only remaining case is as follows:
Case IV: (In, Out), (Out, In), and (In, Out) are not Nash Equilibria This case requires any of the mutually exclusive systems of inequalities to apply:

- $U^f_{alone} > v > U^w_{alone}$ and $U^w_{both} > 0$.
- $U^w_{alone} > v > U^f_{alone}$ and $U^f_{both} > 0$.
- $U^w_{alone} \wedge U^f_{alone} > v$ and $U^f_{both} \wedge U^w_{both} > 0$

Whenever any of these systems is satisfied, (In, In) is the unique Nash equilibrium of the game, therefore Pareto undominated.
Conjecture and Relaxed Form of (PR) Using the history-dependent $t = 2$ reservation values

$$r(h_i) = \begin{cases} v & \text{if } x_1(i) \neq i,\ s_i = \tilde{s}_i = 1 \\ 0 & \text{otherwise} \end{cases} \tag{11}$$

in the participation constraints imposes $U^f_{both} \wedge U^w_{both} \geq v$ on the platform's problem. As shown in Proposition 1, the platform's policy under this relaxation of (PR) makes this constraint bind and sets $U^i_{alone} = 0$ for all $i \in W \cup F$. Therefore, with this form of reservation value, (PR) as a whole is satisfied, (In,In) and (Out,Out) are the only Nash equilibria of the Re-entry game, and (In,In) is undominated.

□

Proof of Theorem 1. I prove the theorem for the two-sided incomplete information case, since the other cases are nearly identical. Let $\tilde{u}(h_f)$ denote the non-transfer utility received by a surviving firm f in period 2 according allocations from Proposition 1. That is, $\tilde{u}(h_f) = v$ is $s_f = \tilde{s}_f = 1$, and $v - c$ if $\tilde{s}_f = 0$. Substituting in the $t = 2$ allocations from Proposition 1, the last two terms in R_f become

$$\begin{gathered}\mathbb{E}\left[J^F(q_f)\mathbb{E}^{\chi_1,\tau_1}_{\mathbf{h}|s_f=1,q_f}\left[\tilde{u}(h_f)\right] + \frac{1}{h^F(q_f)}\mathbb{E}^{\chi_1,\tau_1}_{\mathbf{h}|s_f=1,q_f}[\tilde{u}(h_f) - r(h_f)]\right] \\ = \\ \mathbb{E}\left[q_f\mathbb{E}^{\chi_1,\tau_1}_{\mathbf{h}|s_f=1,q_f}\left[\tilde{u}(h_f)\right] - \frac{1}{h^F(q_f)}\mathbb{E}^{\chi_1,\tau_1}_{\mathbf{h}|s_f=1,q_f}[r(h_f)]\right]\end{gathered} \tag{12}$$

To simplify, split the integral $\mathbb{E}^{\chi_1,\tau_1}_{\mathbf{h}|s_f=1,q_f}[\cdot]$ into three events:

- E_1, where f is unmatched in period 1 ($x_1(f) = f$), so $\tilde{u}(h_f) = v - c$, $r(h_f) = 0$.
- E_2, where f was matched and its partner survived to period 2 ($\tilde{s}_f = 1, x_1(f) \neq f$), so $\tilde{u}(h_f) = r(h_f) = v$; and
- E_3, where f was matched but its partner perished ($\tilde{s}_f = 0, x_1(f) \neq f$), so $\tilde{u}(h_f) = v - c$, $r(h_f) = 0$.

I show how to calculate the conditional probability of E_2; the other two are very similar. This conditional probability becomes $\sum_{w\in W}\mathbb{P}\left(\tilde{s}_f = 1, x_1(f) = w|s_f = 1, q_f\right)$ after decomposing by f's match partner. By iterated expectations, each term in the sum is

$$\mathbb{E}_{\mathbf{p},\mathbf{q}_{-f}|s_f=1,q_f}\left[\mathbb{P}\left(s_w = 1, x_1(f) = w|\mathbf{p},\mathbf{q}, s_f = 1\right)\right] \tag{13}$$

Conditional on $(\mathbf{p},\mathbf{q})$, (x_1, s_{-f}) is independent of s_f. Indeed, the allocation cannot affect f's survival, and agents' survival status is drawn independently for a fixed type vector, so their only correlation is through q_f. Therefore, the term in square brackets in the last display is $\mathbb{P}\left(s_w = 1, x_1(f) = w|\mathbf{p},\mathbf{q}\right)$, which in turn equals

$$\mathbb{P}\left(s_w = 1|\mathbf{p},\mathbf{q}, x_1(f) = 1\right) \cdot \mathbb{P}(x_1(f) = w|\mathbf{p},\mathbf{q}) = p_w \cdot \chi_1^{f,w}(\mathbf{p},\mathbf{q}).$$

As before, the equality follows from the independence of x_1, s_w conditional on types (and the definition of χ_1). Finally, conditional on q_f, s_f and $(\mathbf{p},\mathbf{q}_{-f})$ are independent, so (13) becomes simply $\mathbb{E}_{\mathbf{p},\mathbf{q}_{-f}|q_f}\left[p_w \cdot \chi_1^{f,w}(\mathbf{p},\mathbf{q})\right]$, and the conditional probability of E_2 is $\mathbb{E}_{\mathbf{p},\mathbf{q}|q_f}\left[\sum_{w\in W} p_w \cdot \chi_1^{f,w}(\mathbf{p},\mathbf{q})\right]$. Similar calculations yield $\mathbb{E}_{\mathbf{p},\mathbf{q}|q_f}\left[1 - \sum_{w\in W}\chi_1^{f,w}(\mathbf{p},\mathbf{q})\right]$ for E_1 and $\mathbb{E}_{\mathbf{p},\mathbf{q}|q_f}\left[\sum_{w\in W}(1 - p_w) \cdot \chi_1^{f,w}(\mathbf{p},\mathbf{q})\right]$ for E_3.

Plugging these expressions into (12) and evaluating R_f yields, after minor algebra and another application of iterated expectations,

$$R_f = \mathbb{E}\left[q_f(v - c) + \sum_{w\in W}\chi_1^{f,w}(\mathbf{p},\mathbf{q})\left(p_w\, J^F(q_f)v + (1 - p_w q_f)(v - c)\right)\right] \tag{14}$$

with a perfectly symmetric expression for R_w. Ignoring the $q_f(v - c)$ and corresponding $p_w(v - c)$ terms that do not depend on the allocation, the platform's objective becomes

$$\max_{\chi_1}\mathbb{E}\left[\sum_{w\in W}\sum_{f\in F}\chi_1^{f,w}(\mathbf{p},\mathbf{q})\Big(2[1 - p_w q_f](v - c) + \left[J^F(q_f)p_w + J^W(p_w)q_f\right]v\Big)\right],$$

subject to (χ_1-F)

which can be solved pointwise in $(\mathbf{p}, \mathbf{q})$. Note that the coefficient on $\chi_1^{w,f}$ in (14) is precisely $a^F(p_w, q_f)$, and the coefficient in the last display is $a(p_w, q_f) = a^F(p_w, q_f) + a^W(p_w, q_f)$, with a^W given by (10) since the worker has private information. □

Derivation of Planner's Problem. In period 1, the match always generates $2(v-c)$. With probability $p_w q_f$, the match survives into period 2, and generates an additional $2v$. Since the planner always rematches efficiently in period 2, agents generate an additional $v - c$ when only one of them survives, with probability $q_f(1-p_w)+p_w(1-q_f)$. Thus the expected surplus from (w, f) is $S(p_w, q_f) = 2[1-p_w q_f](v-c)+[p_w q_f + q_f p_w]v + q_f(v-c)+p_w(v-c)$, and the planner's $t = 1$ payoff from a feasible match x is

$$\sum_{w \in W} \left(1 - \sum_{f \in F} x^{w,f}\right) p_w(v-c) + \sum_{f \in F} \left(1 - \sum_{w \in W} x^{w,f}\right) q_f(v-c) + \sum_{w \in W} \sum_{f \in F} x^{w,f} a^E(p_w, q_f)$$

where the first two terms account for surplus from agents who go unmatched in period 1. Rearrange and discarding the parts that do not depend on x yields a coefficient of

$$a^E(p_w, q_f) = 2[1 - p_w q_f](v-c) + 2p_w q_f v$$

on $x^{w,f}$. □