Pandora's Auctions: Dynamic Matching with Unknown Preferences

By DANIEL FERSHTMAN AND ALESSANDRO PAVAN*

Matching theory typically assumes that agents know their values for possible partners and confines attention to settings in which matching is either static, or driven by population dynamics. In many environments of interest, instead, dynamics originate in the agents learning their preferences through interactions with other agents. In this short paper, we illustrate how platforms can use appropriately designed auctions to account for the joint value of experimentation and cross-subsidization in dynamic matching markets. The model presented below is a stylized special version of the general one in Fershtman and Pavan (2016). We refer the reader to that paper for a more general treatment and to the Online Supplement to this article for a discussion of how the results in the present paper can be obtained from those in the other paper.

I. Environment

A platform mediates the interactions between agents from two sides of a market. There are two agents on side 1, x and y, and a single agent on side 2, z. Agents are uncertain about the utility, or profit, they derive by interacting with potential partners, but perfectly learn such values after the first interaction. The platform can match at most a pair of agents in each period. As the agents' perceived, as well as true, values are their own private information, the platform uses an auction to determine which match to implement in each period. We allow the platform's objective to be a weighted combination of profits and welfare, but then illustrate how the results specialize in the case of pure profits and pure welfare maximization.

From the perspective of agent $i \in \{x, y, z\}$, the true quality ω_{ij} of the match with agent j can be either "good" ($\omega_{ij} = G$) or "bad" ($\omega_{ij} = B$). Importantly, while the match between i and j may be a bad one from the perspective of agent i, it may be a good one from the perspective of agent j (that is, $\omega_{ij} = B$ may be consistent with $\omega_{ji} = G$). A good match yields agent i a gross flow payoff θ_i, whereas a bad match yields the same agent a gross flow payoff equal to $-L\theta_i$, where $L > 0$, and where the values θ_i are drawn independently across agents, and independently from the horizontal values ω_{ij}, from a uniform distribution with support $[\underline{\theta}, \overline{\theta}]$, with $\overline{\theta} > \underline{\theta} > \overline{\theta}/2 > 0$. Upon joining the platform, each agent i receives a private signal $\sigma_{ij} \in \{G, B\}$, $ij \in \{xz, yz, zx, zy\}$, about the quality of each potential partner. A signal $\sigma_{ij} = B$ perfectly reveals to agent i that $\omega_{ij} = B$, whereas a signal $\sigma_{ij} = G$ reveals to the individual that $\omega_{ij} = G$ with probability $\lambda \in (0, 1)$. After the first interaction with agent j, agent i perfectly learns ω_{ij}. Given i's period-t information $\mathcal{I}_i^t$, i's period-t (expected) match value for interacting with j is thus equal to $\theta_i \cdot \varepsilon_{ij}^t$, where

$$\varepsilon_{ij}^t \equiv \Pr(\omega_{ij} = G|\mathcal{I}_i^t) - L\Pr(\omega_{ij} = B|\mathcal{I}_i^t).$$

The probability that agent i receives a favorable signal $\sigma_{ij} = G$ about agent j is equal to $q \in (0, 1)$. The signal agent i receives about the quality of his match with agent j is independent of the signal the same agent, or any other agent, receives about the quality of any other match. We assume q is large enough that $q\lambda > (1 - q\lambda)L$, which guarantees that, ex-ante, agents are sufficiently optimistic about the value of each potential interaction. To facilitate the analysis, we also focus on the limit in which $\lambda \to 1$ and $2\underline{\theta} \to \overline{\theta}$. Agents maximize the expected discounted sum of their flow payoffs us-

* Fershtman: Department of Economics, Northwestern University, 2001 Sheridan Road, Andersen Hall, Evanston, IL 60208-2600 (dfershtman@u.northwestern.edu). Pavan: Department of Economics, Northwestern University and CEPR, 2001 Sheridan Road, Andersen Hall, Evanston, IL 60208-2600 (alepavan@northwestern.edu). This short paper has been prepared for the 2017 *American Economic Review* Papers and Proceedings. We are grateful to the 2017 ASSA participants, and especially to our discussant, Dirk Bergemann, for helpful comments and suggestions. The usual disclaimer applies.

ing the common discount factor $\delta \in (0,1]$.

II. Pandora's Auctions

The platform uses the following auction procedure. At $t = 0$, after privately learning θ_i, each agent i decides whether or not to join the platform. If he joins, agent i then chooses a membership status $\beta_i = \beta(\hat{\theta}_i)$ and pays a fee p_i^0. Membership statuses are conveniently indexed by the agents' "vertical" types $\hat{\theta}_i \in [\underline{\theta}, \overline{\theta}]$, with the function $\beta : [\underline{\theta}, \overline{\theta}] \to \mathbb{R}$ strictly positive, non-decreasing, and bounded. Higher status grants an agent more favorable treatment in the auctions. Upon joining the platform and seeing which agents are on board, each agent i receives a signal $\sigma_{ij} \in \{G, B\}$ about the quality of the match with each possible partner.

At each subsequent period $t \geq 1$, each agent i then submits a bid b_{ij}^t for each possible partner j from the opposite side. Bids can be either positive or negative, reflecting the idea that certain agents may dislike certain interactions and ask to be compensated by the platform. The received bids, along with the agents' membership statuses, determine which match is implemented, according to the rules described below. Let $m_{ij}^t = 1$ denote the decision to match agents i and j in period t and $m_{ij}^t = 0$ the decision to leave the pair unmatched, and denote by $m^{<t}$ the period-t history of all past matches. Period-t bids that are inconsistent with the above processes are automatically replaced with a bid equal to 0.[1] Each matched agent pays p_i^t to the platform, whereas each unmatched agent pays nothing. After the first interaction with agent j, agent i perfectly, but privately, learns ω_{ij}. Unmatched agents receive no further information.

The selection of the matches is based on the following "scores." Given the agents' period-t bids, b^t, and the agents' statuses, β, the platform assigns a score S_{ij}^t to each possible match. If the pair (i,j) was matched in the past, the score is equal to $S_{ij}^t(b^t; \beta, m^{<t}) = \beta_i b_{ij}^t + \beta_j b_{ji}^t$. If, instead, the pair (i,j) has never been matched, the score is equal to

$$S_{ij}^t\left(b^t; \beta, m^{<t}\right) = \left(\beta_i b_{ij}^t + \beta_j b_{ji}^t\right)\left(\frac{1 - \delta + \frac{\delta\lambda^2}{\lambda - (1-\lambda)L}}{1 - \delta + \delta\lambda^2}\right)$$

if $b_{ij}^t, b_{ji}^t \geq 0$, it is equal to

$$S_{ij}^t\left(b^t; \beta, m^{<t}\right) = \beta_l b_{lk}^t \left(\frac{1 - \delta + \frac{\delta\lambda}{\lambda - (1-\lambda)L}}{1 - \delta + \delta\lambda}\right) + \beta_k b_{kl}^t$$

if $b_{lk}^t \geq 0$ and $b_{kl}^t < 0$, for $lk, kl \in \{ij, ji\}$, $kl \neq lk$, and it is equal to $S_{ij}^t(b^t; \beta, m^{<t}) = -1$ if $b_{ij}^t, b_{ji}^t < 0$.

In each period, the match with the highest non-negative score is implemented. If all scores are negative, no match is formed.[2]

The description of the matching auction is completed by specifying the payment scheme. The period-t payment of agent $i \in \{x, y, z\}$ when matched with agent $j \neq i$ in period t is equal to

$$p_i^t = \frac{(1-\delta)W_{-i}^t - \beta_j b_{ji}^t}{\beta_i}, \tag{1}$$

where $W_{-z}^t = 0$, whereas for $i, k \in \{x, y\}$, $k \neq i$, W_{-i}^t is the continuation weighted surplus associated with the match (k, z); that is, the expected weighted continuation surplus generated by matching the pair (k, z) from period t till this continuation value turns negative. Formally, for any $d \in \mathbb{R}$, let $[d]^+ \equiv \max\{0, d\}$. If the pair

[1]Formally, for any pair of agents who have interacted in the past, any bid $b_{ij}^t \notin \left[-L\overline{\theta}, -L\underline{\theta}\right] \cup \left[\underline{\theta}, \overline{\theta}\right]$ is automatically replaced with a bid equal to zero. Likewise, for any pair of agents who have never interacted in the past, any bid $b_{ij}^t \notin \left[-L\overline{\theta}, -L\underline{\theta}\right] \cup \left[(\lambda - (1-\lambda)L)\underline{\theta}, (\lambda - (1-\lambda)L)\overline{\theta}\right]$ is automatically replaced with a bid equal to zero. Finally, the bids b_{zx}^t, b_{zy}^t by agent z from side 2 are replaced with bids equal to zero if they reveal they are inconsistent with the same vertical type $\theta \in [\underline{\theta}, \bar{\theta}]$, that is, if $b_{zx}^t/k \neq b_{zy}^t/\hat{k}$ for all $k, \hat{k} \in \{-L, \lambda - (1-\lambda)L, 1\}$.

[2]The reader is referred to the Online Supplement for a discussion of how the above scores relate to the indexes in Fershtman and Pavan (2016), and for how the latter in turn capture the trade-offs between exploitation and experimentation in dynamic matching markets.

(y,z) was never matched, then

$$W_{-x}^{t} = \left(\beta_y b_{yz}^t + \beta_z b_{zy}^t\right) + \frac{\delta\lambda^2\left(\beta_y b_{yz}^t + \beta_z b_{zy}^t\right)}{(1-\delta)\left(\lambda-(1-\lambda)L\right)} + \frac{\delta\lambda(1-\lambda)}{(1-\delta)\left(\lambda-(1-\lambda)L\right)} \times \sum_{lk,kl\in\{yz,zy\},kl\neq lk}\left[\beta_l b_{lk}^t - \beta_k b_{kl}^t L\right]^+$$

if $b_{yz}^t, b_{zy}^t \geq 0$,

$$W_{-x}^{t} = \left[\left(\beta_l b_{lk}^t + \beta_k b_{kl}^t\right) + \frac{\delta\lambda}{1-\delta}\left(\frac{\beta_l b_{lk}^t}{\lambda-(1-\lambda)L} + \beta_k b_{kl}^t\right)\right]^+$$

if $b_{lk}^t \geq 0$ and $b_{kl}^t < 0$ for some $lk, kl \in \{yz, zy\}$, $kl \neq lk$, and $W_{-x}^t = 0$ if $b_{yz}^t, b_{zy}^t < 0$. If, instead, the pair (y,z) was matched in the past, then

$$W_{-x}^{t} = \frac{\left[\beta_y b_{yz}^t + \beta_z b_{zy}^t\right]^+}{1-\delta}.$$

The value W_{-y}^t is defined analogously.

Let $\hat{\boldsymbol{\theta}} \equiv (\hat{\theta}_x, \hat{\theta}_y, \hat{\theta}_z)$ denote a profile of membership statuses (recall that the latter are conveniently indexed by the vertical types, θ), and, similarly, let $\hat{\boldsymbol{\theta}}_{-i}$ denote a profile of membership statuses of all agents excluding agent i. Let $D_i(\hat{\boldsymbol{\theta}})$ denote agent i's expected "match quality" from participating in the auction, given $\hat{\boldsymbol{\theta}}$, and let $K_i \in \mathbb{R}$.[3] The period-0 membership fee charged to each agent i is equal to

$$p_i^0 = \hat{\theta}_i D_i(\hat{\boldsymbol{\theta}}) - \int_{\underline{\theta}}^{\hat{\theta}_i} D_i(\hat{\boldsymbol{\theta}}_{-i}, y)dy - \mathbb{E}\left[\sum_{t=1}^{\infty}\delta^t p_i^t | \hat{\boldsymbol{\theta}}\right] - K_i. \tag{2}$$

We refer to the above matching auction as a *Pandora β-auction*. Note that the auction is fully transparent – previous matches, payments, bids and membership choices are all public.

[3]See the Online Supplement for a precise definition of the D_i functions.

A perfect Bayesian equilibrium (PBE) of the above matching auction is *truthful* if each agent $i \in \{x,y,z\}$ (a) participates in period 0 and selects the membership status $\beta(\theta_i)$ designed for his true vertical type θ_i, and (b) in each subsequent period $t \geq 1$ submits for each potential partner j a bid equal to his true (expected) match value, $b_{ij}^t = \theta_i \varepsilon_{ij}^t$, given the information $\mathcal{I}_i^t$.

PROPOSITION 1: *For any positive, non-decreasing, and bounded function $\beta(\cdot)$, there exist $K_i \in \mathbb{R}$, $i = x,y,z$, such that the Pandora β-auction described above admits a truthful equilibrium.*

The above scores are the analogs of Weitzman's (1979) reservation prices in a "Pandora's boxes search problem". Matching agents according to the above procedure guarantees that the weighted continuation surplus is maximized at each history, provided that all agents, irrespective of their past behavior, in the continuation game that starts with period $t \geq 0$, follow truthful strategies. Note that, in a truthful equilibrium, once a pair of agents is matched for the second time, it is then matched in all subsequent periods as well, and payments remain the same in all subsequent periods.

Arguments similar to those in Bergemann and Välimäki (2010), along with arguments similar to those in Pavan, Segal and Toikka (2014) then imply that, starting from any period-t history, $t \geq 1$, irrespective of the agents' own past behavior and of their beliefs about other agents' past and current types, it is optimal for all agents to bid their true expected flow match values in each period of the continuation game. Allowing agents to adjust their bids at each period, including those following periods in which agents are expected to receive no new private information, guarantees that truthful strategies constitute a *periodic ex-post equilibrium*, starting from any period-t history, $t \geq 1$.

The proof is then completed by showing that, when the membership fees are as in (2), it is optimal for each type θ_i of each agent i to join the platform in period $t = 0$, and select the membership status $\beta(\theta_i)$ designed for θ_i. Specifically, the fee (2) is added to the subsequent payments (1) so that each agent's intertemporal equilibrium payoff, as perceived in period $t = 0$, satisfies an envelope condition relating the period-0 (interim) expected payoff to the expected discounted match quality. Matching dynamics under the truthful equilibria of the proposed auction can then be shown to satisfy an appropriate average monotonicity condition by which match quality is increasing in (a) each agent's period-0 membership status, for given true vertical type θ_i, and (b) each agent's true vertical type, for given period-0 membership statuses. Together with the aforementioned envelope conditions, such an average monotonicity property then implies that there exist K_i large enough, $i = x, y, z$, such that participating and selecting a membership status $\beta(\theta_i)$ is optimal for each type θ_i of each agent $i = x, y, z$ in period $t = 0$, again irrespective of the agent's beliefs about other agents' period-0 information.

III. Profit- vs Welfare-Maximizing Auctions

We now turn to the distortions in matching dynamics brought in by profit maximization. First, we argue that the above auctions admit as special cases both auctions that maximize welfare, as well as auctions that maximize the platform's profits, over all possible mechanisms. We then compare matching dynamics under profit maximization with their counterparts under welfare maximization.

PROPOSITION 2: *The Pandora β-auction with weights $\beta(\theta) = 1$, all $\theta \in [\underline{\theta}, \bar{\theta}]$, and payments given by (1) and (2), with K_i sufficiently large, $i = x, y, z$, maximizes welfare over all possible mechanisms. The Pandora β-auction with weights $\beta(\theta) = \left(2\theta - \overline{\theta}\right)/\theta$, all $\theta \in [\underline{\theta}, \bar{\theta}]$, and payments given by (1) and (2), with $K_i = 0$, all $i = x, y, z$, maximizes profits over all possible mechanisms.*

That the Pandora β-auction with weights $\beta(\theta) = 1$ all $\theta \in [\underline{\theta}, \bar{\theta}]$, and payments given by (1) and (2) with K_i large enough to induce participation by all agents, maximizes welfare over all possible mechanisms follows from the fact that the induced matches maximize the sum of all agents' continuation payoffs, starting from each history. That the β-auction with weights $\beta(\theta) = \left(2\theta - \overline{\theta}\right)/\theta$, all $\theta \in [\underline{\theta}, \bar{\theta}]$, and payments given by (1) and (2), with $K_i = 0$, maximizes profits is a consequence of the following observations. Given any mechanism Γ, and any Bayes Nash equilibrium (BNE) of the game induced by it, the platform's profits are equal to the sum of all agents' expected *weighted* payoffs, net of the payoff expected by each agent's lowest vertical type, $\underline{\theta}$, where the weights β correspond to the agents' "virtual vertical types" $\beta(\theta) = \left(2\theta - \overline{\theta}\right)/\theta$, all $\theta \in [\underline{\theta}, \bar{\theta}]$. It is then easy to see that, in the proposed β-auctions, when all agents follow truthful strategies, the (state-contingent) matches implemented in equilibrium maximize the above expected weighted surplus function after any period-t history, $t \geq 1$. Furthermore, when $K_i = 0$, the participation constraint of each agent's lowest vertical type $\underline{\theta}$ binds. The monotonicity of the average match quality mentioned above, along with the fact that the period-0 membership fees satisfy the envelope formula in (2) in turn guarantee that participation is optimal also for all types $\theta > \underline{\theta}$. Combined, the above properties imply that the platform's profits under the truthful equilibria of the proposed auctions are higher than under any BNE of any other mechanism, thus establishing the result.

Below, we show that the order of experimentation under profit maximization may be inefficient. Furthermore, matching dynamics under profit maximization may involve upwards distortions: matches that are proved to be socially undesirable (in the sense that the sum of the match values of the involved agents is negative) may be implemented for arbitrarily long horizons.

PROPOSITION 3: *Profit maximization may lead to inefficiency in the order of search, as well as to the implementation of matches that are socially undesirable for arbitrarily long horizons.*

To illustrate the distortions in Proposition 3, let $1 < L < \overline{\theta}/\underline{\theta}$, and consider the following sequence of realizations of the agents' private information: $\theta_x = \overline{\theta}$, $\theta_z = \underline{\theta}$, and $\theta_y = \underline{\theta} + \nu$, where $0 < \nu < (L-1)\underline{\theta}$. Upon joining, the agents receive the signals $\sigma_{xz} = \sigma_{yz} = \sigma_{zy} = G$, whereas $\sigma_{zx} = B$. That is, with the exception of agent z, who receives a bad signal about his match with agent x, all other agents are optimistic about the quality of each potential interaction. The true quality of each match is given by $\omega_{yz} = G$ and $\omega_{xz} = \omega_{zx} = \omega_{zy} = B$.

When $\lambda \to 1$ and $2\underline{\theta} - \overline{\theta} \to 0$, under welfare maximization, $S^1_{yz} > S^1_{xz} > 0$, whereas, under profit maximization, $S^1_{xz} > S^1_{yz} > 0$. The first match implemented under welfare maximization is therefore (y,z), whereas the first match implemented under profit maximization is (x,z). The period-2 scores under welfare maximization are then such that $S^2_{xz} = S^1_{xz} > 0 > S^2_{yz}$, whereas the period-2 scores under profit maximization are such that $S^2_{yz} = S^1_{yz} > 0 > S^2_{xz}$. Therefore, in period 2, the platform implements a match different from the one implemented in period 1, irrespective of whether it maximizes welfare or profits. Finally, for any $t \geq 3$, the scores under welfare maximization are such that $S^t_{xz}, S^t_{yz} < 0$, whereas the scores under profit maximization are such that $S^t_{xz} < 0 < S^t_{yz}$. Therefore, from period 3 onward, a profit-maximizing platform matches the pair (y,z) in each period, despite the fact that all matches are known to be undesirable from a welfare standpoint.

The reason why a profit-maximizing platform may "lock" a pair of agents whose payoff sum is negative into arbitrarily long interactions is that this may serve as a device to discourage agents with a high vertical type to underbid, thus mimicking lower types.

IV. Conclusions

Matching auctions similar to those introduced above can be used in richer environments in which (a) agents learn the attractiveness of partners gradually over time (i.e., through multiple, possibly infinite, interactions), (b) matching is many-to-many (i.e., the same agent may be matched to multiple agents from the opposite side), (c) exogenous shocks may alter the desirability of existing matching allocations, (d) the platform may incur (possibly history-dependent) costs in implementing the different matches, and (e) intermediate capacity constraints may permit the platform to implement more than a single match but not all possible matches in each period. In this case, the scores are more sophisticated than the ones presented here but often continue to take the form of "indexes" similar to those considered in the experimentation literature. We refer the reader to Fershtman and Pavan (2016) for details.

REFERENCES

Bergemann, Dirk, and Juuso Välimäki. 2010. "The dynamic pivot mechanism." *Econometrica*, 78(2): 771–789.

Fershtman, Daniel, and Alessandro Pavan. 2016. "Matching auctions." Mimeo, Northwestern University.

Pavan, Alessandro, Ilya Segal, and Juuso Toikka. 2014. "Dynamic mechanism design: A myersonian approach." *Econometrica*, 82(2): 601–653.

Weitzman, Martin L. 1979. "Optimal search for the best alternative." *Econometrica: Journal of the Econometric Society*, 641–654.