

The Efficiency of Open Access in Platforms for Networked Cournot Markets

John Z. F. Pang*, Hu Fu†, Won I. Lee*, and Adam Wierman*

*Department of Computing and Mathematical Sciences

California Institute of Technology, Pasadena, California 91125

Email: jzpang@caltech.edu, wilee@caltech.edu, adamw@caltech.edu

†Department of Computer Science

University of British Columbia, Vancouver, British Columbia V6T 1Z4, Canada

Email: hufu@cs.ubc.ca

Abstract—This paper studies how the efficiency of an online platform is impacted by the degree to which access of platform participants is open or controlled. The study is motivated by an emerging trend within platforms to impose increasingly fine-grained control over the options available to platform participants. While early online platforms allowed open access, e.g., Ebay allows any seller to interact with any buyer; modern platforms often impose matches directly, e.g., Uber directly matches drivers to riders. This control is performed with the goal of achieving more efficient market outcomes. However, the results in this paper highlight that imposing matches may create new strategic incentives that lead to increased inefficiency. In particular, in the context of networked Cournot competition, we prove that open access platforms guarantee social welfare within 7/16 of the optimal; whereas controlled allocation platforms can have social welfare unboundedly worse than optimal.

I. INTRODUCTION

Over the last decade, online marketplaces have reshaped whole industries, with internet platform companies leading the way [29]. Platform companies such as Facebook, Uber, Amazon, Ebay, etc. now make up a 3 trillion dollar market in the US alone, and it is growing quickly [1].

The rise of the platform economy [29] brings with it a wide variety of engineering, economic and social challenges. Historically, markets have been slow to evolve, and finding the “right” trading partners has been a daunting task. However, the integration of networks and information technology into marketplaces has led to complex platforms that facilitate matches among participants. There is now an unprecedented level of control over the operation of these markets. Companies are engineering platforms to control the flow of information, recommend matches, and enforce prices and terms of trade. As such, the decisions made in the design of the platforms create intricate and subtle interactions between computational constraints, network constraints, and market outcomes.

At the heart of platform design is the design of the matching algorithm that determines the matches between firms (sellers) and consumers (buyers). Platforms today have a wide variety of approaches for matching. Some platforms, such as Etsy, Airbnb, Ebay, and Upwork, follow an *open access* model – they provide information on all candidate matches, allowing the sellers and the buyers to make their own choices about how

to match [32, 23, 22, 25]. On the other extreme, platforms like Uber follow a *controlled allocation* model — they provide no information about the candidate matches, only presenting a specific opportunity for a match [15, 13, 38, 28]. In between, there are *discriminatory access* platforms, such as Amazon, which impose constraints on the firms limiting which markets they can enter; e.g., only sellers with low enough prices and high enough reviews are eligible to be shown in the Buy Box, the default seller on a product’s front page. Since about 80% of customers do not look beyond the Buy Box, in effect, they are only accessible to the eligible sellers [16, 37, 41, 5].

These contrasting design choices highlight interesting trade-offs in terms of the interaction between incentives and optimization. In particular, over the past decade there has been a noticeable shift from open access designs toward discriminatory access and controlled allocation designs. Uber, which uses a controlled allocation design, is a good exemplar of the current trends in platform design. Uber uses a highly optimized matching algorithm to determine matches between drivers and riders in addition to the price paid for the route. Riders and drivers are not alerted to candidate matches, only to the specific match that Uber determines is “optimal” [15, 13, 38, 28].

It is easy to understand the intuition leading to this shift. Platforms have an incentive to optimize social welfare in order to create strong positive network effects to feed growth, e.g., increasing customer satisfaction leads to growth in the number of customers, which makes it more desirable for firms to join the platform, and vice versa. Thus, it is highly desirable to ensure that matches are made so as to maximize social welfare, and thus network effects. Another recent work in this area [8] looked at a similar problem, in a slightly different setting, with the aim of maximizing transaction volume.

It is natural to assume that open access platforms fail at this goal. Allowing firms and consumers to make their own matches means that inefficient matches are possible, due to the strategic incentives of each participant. Naturally, platforms feel that they can match more optimally, and so choose to either limit the matches available (discriminatory access) or enforce specific “optimal” matches (controlled allocation).

While the above discussion is intuitive, it misses one crucial

point. The “optimal” matches a platform computes are not truly optimal because the fact that matches are limited or enforced impacts the incentives of the firms in the marketplace. In particular, this may cause it to be less desirable to produce, and thus there could be less participation in the marketplace as a result, which leads to social welfare loss.

Contributions of this paper: This paper addresses the question of whether platforms that control access lead to more efficient markets or not. That is, how does the inefficiency created by strategic choices of firms given an open access platform compare with that of a controlled allocation platform?

The *main results in this paper* show, perhaps counterintuitively, that open access designs are near optimal, achieving social welfare no worse than $7/16$ of the social optimal assignment; whereas controlled allocation designs can be unboundedly worse than the social optimal assignment.

More specifically, the results in this paper contrast open access, discriminatory access, and controlled allocation platforms in the context of networked Cournot competition, a classical model of competition in networked markets. We adopt the model of [2, 9] to characterize competition and introduce novel models of platforms within this framework.

In the context of this model, this paper makes three main contributions. First, in Section III, we study open access platform designs and prove our main result: that open access designs achieve social welfare no worse than $7/16$ of the social optimal assignment (Theorem 3), regardless of the size and makeup of the market. This result is the main technical contribution of the paper and is the first “price of anarchy” bound for networked Cournot competition. It extends the price of anarchy bound of Johari and Tsitsiklis [27] for a single Cournot market to multiple, networked markets using a novel technique based on bounding the production costs using terms proportional to the consumer welfare.

Second, in Section IV, we study discriminatory access platform designs. Clearly, the optimal discriminatory access design can be no worse than the open access design, so the question is what improvement is possible given the extra complexity involved. Our analysis highlights two key points: (i) determining the optimal discriminatory design is a very challenging optimization problem, and (ii) the gain of the optimal design over the open-access design is typically small, less than 3% in the simulations we perform; theoretically, our efficiency bound for the open access platforms also imply that introducing discriminatory access cannot improve the welfare over the open access by more than a factor of $\frac{16}{7}$.

Third, in Section V, we study controlled allocation platforms and show that, in the worst case, controlled allocation can lead to equilibrium outcomes with loss in social welfare that grows linearly in the number of markets in the system (Theorem 8). Thus, the efficiency loss in controlled allocation platforms can be unbounded.

The combination of these three results highlights an inherent danger associated with the current trend in platform design toward exerting more control over matches. While this seemingly allows fine-grained optimization over the matches, as a

byproduct it creates incentives that impact the participation of firms in terms of production, leading to a possibility of increased inefficiency if not done carefully. Open access platform designs avoid this danger.

Related work: Our work builds on, and contributes to, two related literatures: a) works studying the design of platforms and b) works studying competition in networked markets.

a) *Platform design:* The proliferation of online platforms has spurred on a new literature focused on understanding the impacts of the design of platforms on their success. Work in this literature has focused on a variety of topics, including pricing [42], insulation [43] and competition [7], to list a few. Particularly relevant to our work are recent empirical findings that show significant price dispersion in online marketplaces [20], which in turn drives platforms to focus on differentiated products in order to create distinct consumer markets [19]. Such results highlight the need to study platforms in the context of *networked* competition.

Our focus is on distinguishing platforms based on the level of access provided. While some platforms encourage open access, e.g. Airbnb, others, e.g., Uber, use some sort of black-box mechanism, trusting that their mechanism brings out the most value. Algorithmic matching, exemplified by Amazon’s Buy Box, leads to discriminatory access, a middle ground between the two extremes.

To this point, no analytic work has studied the impact of strategic incentives created by discriminatory access and controlled allocation in online platforms. This paper provides the first such study, using the classical model of networked Cournot competition as the setting.

b) *Competition in networked settings:* Models of competition in networked settings have received considerable attention in recent years. These models come in various forms, including networked Bertrand competition, e.g., [24, 6, 14, 4], networked Cournot competition, e.g., [9, 2, 26], and various other non-cooperative bargaining games where agents can trade via bilateral contracts and a network determines the set of feasible trades, e.g., [21, 18, 33, 3, 31].

Our paper fits squarely into the emerging literature on networked Cournot competition. The model of networked competition we study has been considered previously, beginning with [11] and continuing through [26, 9, 2, 10, 12]. The contribution of our work comes in two forms. First, the papers above focus on characterizing the existence and uniqueness of Nash equilibria [26, 9, 10, 12] and, more recently, the complexity of computing equilibrium outcomes [2]. In contrast, our work focuses on characterizing the economic efficiency of the equilibrium outcomes as compared to the social welfare maximizer, i.e., characterizing the so-called price of anarchy. Additionally, with the exception of [12] (which does not study the economic efficiency of the market), the literature on networked Cournot competition focuses on situations where firms operate independently, without governance, while we consider situations where a platform may exert control over the matches between firms and markets.

The technical work in our paper is most closely related to Johari and Tsitsiklis [27], who studied the economic efficiency of Cournot competition in a single market setting, as opposed to the networked market setting we consider. Due to the relative simplicity of a single market, [27] was able to exhaustively search the parameter space and solve for the worst case bound; this approach is not viable in more complex multi-market settings.

II. PRELIMINARIES

In this paper we contrast the efficiency of three different approaches (illustrated in Figure 1) for matching and pricing in the design of online platforms:

- (i) *Open access platforms* allow any firm to access any market, without constraint. These are exemplified by platforms such as Etsy, Airbnb, Ebay, and Upwork.
- (ii) *Discriminatory access platforms* limit the set of markets accessible by each firm as a result of properties of the firms or markets, e.g., the quality or cost of the firm in that market. These are exemplified by the Buy Box at Amazon, which shows only sellers with low enough cost and high enough quality [16].
- (iii) *Controlled allocation platforms* enforce a strict matching between firms and markets. These are exemplified by the automatic matching performed by Uber, which directly matches drivers with riders without providing a list of alternatives [15].

To contrast these three styles of platform design, we use a classic economic model of competition in networked markets: networked Cournot competition. In the following, we first introduce networked Cournot competition and then describe the three platform models we consider.

A. Networked Cournot Competition

This paper adopts the model of networked Cournot competition introduced by Bimpikis et al. [9] and Abolhassani et al. [2]. This model generalizes classical Cournot competition, in which firms compete in a single market. The defining characteristic of Cournot competition is that firms compete by choosing *quantities*, and the eventual price is determined by the aggregate supply. This is in contrast to price-competition, a.k.a. Bertrand competition.

Our interest is in Cournot competition in a *networked* setting. Specifically, we consider a set M of m markets and a set F of n firms that are connected by a set of edges $\mathcal{E} \subseteq F \times M$, where $(i, j) \in \mathcal{E}$ represents firm i having access to market j .

For $(i, j) \in \mathcal{E}$, the *production* $q_{ij} \geq 0$ is the supply from firm i to market j . The total production of a firm i is $s_i = \sum_{j:(i,j) \in \mathcal{E}} q_{ij}$, and the total supply at market j is $d_j = \sum_{i:(i,j) \in \mathcal{E}} q_{ij}$. (One can consider $q_{ij} = 0$ for all $(i, j) \notin \mathcal{E}$, and from this point on we will drop the range $(i, j) \in \mathcal{E}$ for summations.) We use q to represent the vector of productions $(q_{ij})_{(i,j) \in \mathcal{E}}$. Additionally, we use s_{-i} to denote the vector $(s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_n)$, and similarly for $q_{-i,j}$.

The price in a market is determined by the total supply through a demand function. Following [9], we focus on linear inverse demand throughout this paper. Specifically, the price p_j at market j is given by $\alpha_j - \beta_j d_j$, where α_j and β_j are market specific parameters. The *demand curve* of market j characterizes the price p_j as a function of demand d_j . As in [9], a firm experiences a production cost that is quadratic in its total production, i.e., firm i 's cost is $c_i s_i^2$, with c_i being a firm specific parameter. Note that analytic characterizations outside of linear demand and quadratic production costs are typically difficult. See [2] for a discussion.

The efficiency of a market is measured by the *social welfare*. Given production q , the social welfare is

$$\text{SW}(q) = \sum_j d_j \left(\alpha_j - \frac{\beta_j d_j}{2} \right) - \sum_i c_i s_i^2. \quad (1)$$

In (1), the first term for each market j , sometimes called the consumer welfare, is the area under the demand curve from 0 to d_j . If production is costless, this term would be the social value created by the production.

A production vector q is *socially optimal* if it solves the optimization problem $\max_q \text{SW}(q)$. We use q^* to denote a solution to the problem. $\text{SW}(q^*)$ is the maximum welfare achievable under a central planner that can control all the productions.

At an optimal solution, the marginal welfare with respect to each positive production q_{ij}^* should be 0. Note that the marginal consumer welfare is just the price at the market. Where q_{ij}^* is 0, it should be that the price at market j is lower than the marginal cost of firm i . Since $\text{SW}(q)$ is concave in q , these first order conditions fully characterize the solution to the maximization problem, giving the following lemma.

Lemma 1. *A production vector q^* is socially optimal if and only if the following is true: either $q_{ij}^* = 0$ and $p_j \leq 2c_i s_i$, or*

$$0 = \frac{\partial \text{SW}}{\partial q_{ij}} = p_j(d_j) - 2c_i s_i \quad (2)$$

B. Price of Anarchy

The optimal welfare $\text{SW}(q^*)$ is a benchmark, but in the settings we consider there is no social planner that can control the behavior of the firms. Instead, firms have the freedom to choose their production levels so as to maximize their individual profits. Concretely, given production vector q , the profit of a firm i is $\pi_i(q) = \sum_j q_{ij} p_j(d_j) - c_i s_i^2$. We will consider *Nash equilibria* of the platform: in an equilibrium, no firm can unilaterally improve its profit when the productions of the other firms remain the same.¹

In general, the social welfare in an equilibrium is less than the optimal. Following a large literature in algorithmic game theory [e.g. 30, 36, 35, 17], we quantify the inefficiency of the equilibria by the notion of *Price of Anarchy*. The price of anarchy of a platform is the ratio between the optimal welfare

¹Note that we focus on pure Nash equilibria only, since a mixed Nash equilibrium where firms randomize their supplies is unnatural.




Fig. 1. From left to right: (a) Open Access Platforms allow firms to participate in any of the markets at any quantity they desire. (b) Discriminatory Access Platforms constrain the markets in which firms can participate (the red edges). (c) Controlled Allocation Platform match firms’ productions and the markets. The firms still decide their total supply.

$SW(q^*)$ and the smallest welfare of any Nash equilibrium. The closer to 1 this ratio is, the more efficient the platform is. We study the price of anarchy of platforms under three types of access control.

C. Platform Models

We consider three styles of platform design in this paper: open access, discriminatory access, and controlled allocation.

1) *Open Access Platforms*: An open access platform allows access of all markets to all firms, and each firm i is able to decide its production q_{ij} for each market j . When all firms make this decision to maximize their individual profits, a game is formulated among them. Bimpikis et al. [9] showed that, in general, this game has a unique pure Nash equilibrium. In this paper we focus on this solution concept.

A vector of production levels q is at a Nash equilibrium iff each firm i ’s productions maximize its profit given all other firms’ productions. The marginal profit of firm i for any market j for which $q_{ij} > 0$ is $\frac{\partial \pi_i}{\partial q_{ij}} = \alpha_j - 2\beta_j q_{ij} - 2c_i s_i$. At the Nash equilibrium, this marginal profit should be 0; also, for any market j for which $q_{ij} = 0$, it must be that the price p_j is smaller than the marginal cost $2c_i s_i$. In fact, these conditions fully characterize the Nash equilibrium:

Lemma 2 (Bimpikis et al. [9]). *A production vector q^N forms a Nash equilibrium if and only if the following is true: for each firm i and each market j , either $p_j \leq 2c_i s_i$ and $q_{ij}^N = 0$, or*

$$0 = \frac{\partial \pi_i}{\partial q_{ij}} = p_j(d_j) - \beta_j q_{ij} - 2c_i s_i. \quad (3)$$

2) *Discriminatory Access Platforms*: A discriminatory access platform differs from an open access one in that the platform imposes constraints on the set of markets accessible by each firm. In other words, the platform can choose the set of edges in \mathcal{E} . With \mathcal{E} fixed, each firm i can still choose production levels q_{ij} for each market j it can access. The uniqueness of equilibrium shown by [9] still applies, and so we focus on pure Nash equilibria. The characterization of the Nash equilibrium for the open access platform in Lemma 2 also carries over, simply by considering only markets that are accessible by each firm i .

The value of the discriminatory access platform design is that the platform can limit firms to choosing only from markets where a matching is “efficient”. Thus, the platform can

optimize the set of edges in \mathcal{E} in order to maximize the social welfare and so discriminatory access can outperform open access if the set of edges is designed appropriately. However, the results in this paper give evidence that this optimization problem is nontrivial. Additionally, perhaps surprisingly, our results show that the complete graph, i.e., an open access platform, generates welfare that approximates that of the optimal discriminatory access design.

3) *Controlled Allocation Platforms*: A controlled allocation platform exerts complete control over the matching between firms and markets. In particular, firms choose their total production levels but the platform allocates the productions to the individual markets in order to maximize the social welfare.

More specifically, given the total productions s_1, \dots, s_n , the platform chooses q to maximize $SW(q)$ subject to the constraints $\sum_j q_{ij} = s_i$ for any i . The same reasoning as in Lemma 1 shows that the platform will supply to a subset of markets among which an equal price is maintained; when the total productions increase, this subset will expand to include more markets. The demand over all markets can be aggregated to form a platform-wide inverse demand curve that is piecewise linear, decreasing and convex. The reallocation significantly alters the firms’ incentives. The equilibrium production levels will be characterized by first order conditions similar to those in Lemma 2, but with respect to the aggregate demand curve. Unlike the previous two models, controlled allocation platforms may have multiple pure Nash equilibria in general.

III. THE EFFICIENCY OF OPEN ACCESS PLATFORMS

Intuitively, the open access design suffers from inefficiency due to the freedom given to the firms — the production q is determined by firms that maximize their own profits rather than the social welfare. However, our main result in this section shows that the inefficiency created by profit-maximizing behaviors in open access platforms is bounded by a relatively small constant, paralleling what has been observed in many other network games, e.g., routing games and facility location games [36, 40].

Theorem 3. *The social welfare at a Nash equilibrium of an open access platform is at least 7/16 of the optimal social welfare, i.e., open access platforms have a price of anarchy of at most 16/7.*

This theorem is the first bound on the market efficiency of networked Cournot competitions. This bound holds regardless of the potential complexities in the market parameters (i.e., the number of firms, markets, or the parameters of costs and demands).

We also point out that proving Theorem 3 requires overcoming a technical difficulty that is unusual compared to prior literature on the price of anarchy. In particular, in the literature, price of anarchy bounds that focus on settings where the objective is a summation of quantities are almost exclusively concerned with problems where all the terms in the summation are of the same sign, e.g., the congestion in routing games [36] and the welfare in auctions (the summation of bidders' valuations) [17, 39]. Our setting includes terms of different signs, which adds considerable complexity. A rare prior result for such a setting is Johari and Tsitsiklis [27], who studied the inefficiency of a single Cournot market with multiple firms. As a result, [27] could not rely on common tools and pursued an exhaustive analysis, which we find difficult to extend to the networked Cournot setting.

To highlight this issue directly, note that, even though the problem is an *ordinal potential game* [2], the lack of a relationship between the potential function and the welfare (due to the mixed signs of their terms) makes it difficult to apply standard techniques for bounding inefficiency in congestion games [36]. The game is also not *smooth*, as defined by Roughgarden [34]. As a result, our analysis pursues a fundamentally different approach. In particular, a key idea in our proof is to remove the production costs from the welfare expression and to carefully compensate that with a fraction of the consumer welfare.

We now describe the proof idea of Theorem 3. The first idea is to show that, in the Nash equilibrium, no market can have significantly less total demand compared with the optimal (Lemma 5). If there were no production costs, this alone would imply a PoA bound of $\frac{4}{3}$. However, when the costs are different across the firms, the distribution of production among the firms can be important as this affects the summation of the cost terms in (1). In general, this distribution (proportion of production among the firms) is different in the Nash equilibrium from that in the socially optimal production. To address this, the second idea of the proof is to observe that, for a firm that best responds to market demands, its cost can be accounted for by a fraction of the consumer welfare it generates (Lemma 6).

These ideas are fleshed out below. Recall that we denote by q^* the optimal production, and q^N the Nash production. Denote by S_N and S_* the sets $\{j : d_j^N > 0\}$ and $\{j : d_j^* > 0\}$, respectively.

Lemma 4. *There exists firm i , such that $s_i^N \leq s_i^*$.*

Proof. Suppose the lemma is not true, i.e., $\forall i, s_i^N > s_i^*$, we derive contradiction in two cases.

Case 1: $S_N \subseteq S_*$. Since $s_i^N > s_i^*$ for all i , there exists $j \in S_N, d_j^N > d_j^*$; given $S_N \subseteq S_*$, this implies that $p_j^* > p_j^N$

for some market j . On the other hand, for any firm i where $q_{ij}^N > 0$, by the first order condition we have

$$\alpha_j > p_j^N = 2c_i s_i^N + \beta_j q_{ij} \geq 2c_i s_i^N > 2c_i s_i^*.$$

First order condition of the optimal configuration gives $2c_i s_i^* \geq p_j^*$, a contradiction to $p_j^* > p_j^N$.

Case 2: There exists market $j \in S_N \setminus S_*$. In this case, for any i such that $q_{ij}^N > 0$, first order conditions give $2c_i s_i^N < \alpha_j \leq 2c_i s_i^*$, a contradiction to $s_i^N > s_i^*$. \square

Lemma 5. *For any market j , we have $d_j^N \geq \frac{d_j^*}{2}$.*

Proof. Suppose, for the sake of contradiction, that there is some j such that $d_j^N < \frac{d_j^*}{2}$. We show that this contradicts Lemma 4. Recall that $j \in S_*$ implies $q_{ij}^* > 0$ for all i . For any firm i such that $q_{ij}^N = 0$, we have $2c_i s_i^N \geq p_j^N > p_j^* = 2c_i s_i^*$, i.e., $s_i^N > s_i^*$. On the other hand, if $q_{ij}^N > 0$, we have

$$\begin{aligned} 2c_i s_i^N &= \alpha_j - \beta_j d_j^N - \beta_j q_{ij}^N > \alpha_j - \beta_j \frac{d_j^*}{2} - \beta_j \frac{d_j^*}{2} \\ &= \alpha_j - \beta_j d_j^* = 2c_i s_i^*. \end{aligned}$$

Therefore, $s_i^N > s_i^*$ for all i , a contradiction to Lemma 4. \square

We introduce the following short-hand notation to ease the presentation of the rest of the proof.

Definition 1. *For market j with demand rate β_j , let $\text{SWCL}_j(\alpha, d) := d(\alpha - \frac{\beta_j}{2}d)$ be the additional consumer welfare generated by costless production of quantity d when the price starts from α .*

Lemma 6. *Given a set of markets, where each market j has demand curve $p_j(d_j) = \alpha_j - \beta_j d_j$, if a single firm is best responding, and its production level is d_j^N at each market j , then its total cost is at most $\sum_j \text{SWCL}_j(\alpha_j - \beta_j \cdot \frac{d_j^N}{2}, \frac{d_j^N}{2})$.*

In plain language, the lemma states that, if we reduce the firm's production by half in each market but eliminate all its production cost, the remaining social welfare contributed by the firm is no less than the social welfare at its best-response production level with cost.

Proof. Since $\frac{\partial \text{SWCL}_j(\alpha_j, d_j)}{\partial d_j} = p_j(d_j)$, we have

$$\text{SWCL}_j(\alpha_j - \beta_j \cdot \frac{d_j^N}{2}, \frac{d_j^N}{2}) = \int_{\frac{d_j^N}{2}}^{d_j^N} p_j(d) dd.$$

By the Mean Value Theorem, this is equal to $\frac{d_j^N}{2} \cdot p_j(\hat{d})$ for some $\hat{d} \in [\frac{d_j^N}{2}, d_j^N]$. But

$$p_j(\hat{d}) \geq p_j(d_j^N) = 2c_i s_i^N + \beta_j d_j^N \geq 2c_i s_i^N,$$

where the inequality follows by first order conditions (Lemma 2). Therefore, $\text{SWCL}_j(\alpha_j - \beta_j \cdot \frac{d_j^N}{2}, \frac{d_j^N}{2}) \geq c_i s_i^N d_j^N$. The lemma immediately follows by summing over all markets. \square

Lemma 7. For any quadratic function of the form $f = bx - ax^2$, with $a, b > 0$, which takes its maximum value at $f(\frac{b}{2a}) = \frac{b^2}{4a}$, we have $f(\gamma \cdot \frac{b}{2a}) = [1 - (1 - \gamma)^2] \cdot \frac{b^2}{4a^2}$ for any $\gamma \in [0, 1]$.

Proof. One can do the substitution directly to check the statement. A more intuitive way to see this is to note that the image of f is a reversed and translated image of the function $f = ax^2$. \square

Proof of Theorem 3. At the Nash equilibrium, each firm i is best responding to the set of markets where each market j has a price starting from $\alpha_j - \beta_j q_{-i,j}^N$, where $q_{-i,j}^N$ denotes $\sum_{i' \neq i} q_{i',j}^N$. Therefore, by Lemma 6, the production costs at Nash equilibrium can be upper bounded by

$$\sum_j \sum_i \text{SWCL}_j \left(\alpha_j - \beta_j \left(d_j^N - \frac{q_{ij}^N}{2} \right), \frac{q_{ij}^N}{2} \right).$$

For each j , SWCL_j is convex in its first parameter, this is therefore at most $\sum_j \text{SWCL}_j(\alpha_j - \beta_j \cdot \frac{d_j^N}{2}, \frac{d_j^N}{2})$, and the social welfare at the Nash equilibrium is

$$\sum_j \text{SWCL}_j(\alpha_j, d_j^N) - \sum_i c_i (s_i^N)^2 \geq \sum_j \text{SWCL}_j \left(\alpha_j, \frac{d_j^N}{2} \right).$$

As long as the price on a market is positive (which is clearly the case for the Nash equilibrium), SWCL_j is an increasing function in its second parameter, and therefore by Lemma 6, this is at least $\sum_j \text{SWCL}_j(\alpha_j, \frac{d_j^*}{4})$.

So far we have shown that welfare at Nash is at least $\sum_j \text{SWCL}_j(\alpha_j, \frac{d_j^*}{4})$. Now consider the optimal production q^* . Suppose the productions $q_{i,j'}^*$ are fixed for each firm i and each market $j' \neq j$, then the optimal production must maximize the marginal social welfare from market j . More specifically, one can define a function $W_j(d_j | q_{-j})$ which is the maximum marginal social welfare from market j when the total production on j is d_j and the productions at the other markets are q_{-j} :

$$W_j(d_j | q_{-j}) = \max_{\hat{q}: \sum_i \hat{q}_{i,j} = d_j; \hat{q}_{-j} = q_{-j}} \text{SW}(\hat{q}) - \text{SW}(0, q_{-j}),$$

then d_j^* must maximize W_j . Since the cost functions are convex, the optimal social welfare is at most $\sum_j W_j(d_j^*)$. If one can show that W_j is a quadratic function of the form as in Lemma 7, then since $W_j(d_j)$ is pointwise dominated by $\text{SWCL}_j(\alpha_j, d_j)$, the theorem would follow from Lemma 7. W_j being quadratic is in fact easily seen, and can be written explicitly as

$$W_j(d_j | q_{-j}) := d_j(\alpha - \beta d_j) - \sum_i \left[c_i \left(\frac{d_j}{\sum_{i'} \frac{1}{c_{i'}}} + q_{i,-j} \right)^2 - c_i s_{i,-j}^2 \right],$$

where $s_{i,-j}$ is $\sum_{j' \neq j} q_{i,j'}^*$. This completes the proof. \square

IV. THE COMPLEXITY OF DISCRIMINATORY ACCESS PLATFORMS

One may hope to improve the performance of the open access platform design by limiting the access of firms to markets where participation would be “inefficient”, i.e., hurt social welfare. This motivation leads to the design of discriminatory access platforms. The question we ask in this section is: How much improvement can discriminatory access platforms provide over open access platforms?

To begin to address this question, we first need to find the optimal discriminatory access design, i.e., the optimal set of connections to allow between firms and markets. However, determining the optimal set of connections for maximizing social welfare requires optimization over a combinatorial set, the set of bipartite graphs between firms and markets. Thus, to be able to perform such an optimization efficiently it is necessary to find structural properties to exploit. Intuitively, one might expect the monotonicity in edges and costs to help make such an optimization tractable. However, such desirable properties do not hold.

A. The difficulty of optimizing discriminatory access platforms

To highlight the challenge in optimizing discriminatory access platforms we describe a variety of simple examples that violate desirable structural properties. Our first example shows that allowing more access does not always improve welfare at the Nash equilibrium. Surprisingly, this phenomenon occurs even when there is only one market.

Example 1 (Excluding a firm could increase welfare). Consider a three firms, one market system, with the following parameters:

$$\alpha = 100, \beta = 0.45, c = [0.01 \ 0.01 \ 1]^T$$

It can be checked that the Nash productions for this system

$$q = (68.6672, 68.6672, 13.1717),$$

yielding a social welfare of 9686.1. On the other hand, if the platform bars access of the last high-cost firm, the Nash production will be (72.9930, 72.9930), with the welfare improved to 9696.9.

Our second example shows that, not only is there non-monotonicity with respect to the network, there is non-monotonicity with respect to the costs of firms as well.

Example 2 (Nash welfare may increase as a firm’s cost increases). In Example 1, when one increases the high-cost firm’s cost while fixing the costs of the other two firms, the third firm will soon be included in the optimal set of firms. In particular, this shows that the Nash welfare can increase as a firm’s cost increases, even when all other market conditions are fixed.

Finally, our third example highlights that the optimal solution set itself is non-monotonic.

Example 3 (The optimal set of $k + 1$ edges may not contain the optimal set of k edges). Consider the following two firms, two markets system, with the following parameters:

$$\alpha = \begin{bmatrix} 100 \\ 70 \end{bmatrix}, \beta = \begin{bmatrix} 0.5 \\ 0.25 \end{bmatrix}, c = \begin{bmatrix} 0.25 \\ 0.24 \end{bmatrix}.$$

It can be checked that the optimal one-edge access would be to allow firm 2 to access market 1, whereas the optimal two-edge access would be to allow firm 1 and 2 to access market 1 and 2, respectively. This shows that a naïve “greedy” approach for optimizing a discriminatory approach is not possible.

These examples highlight the difficulty of optimizing the set of connections in a discriminatory access platform. However, they stop short of showing that the algorithmic problem is computationally hard in a formal sense. An open problem for future work is to either prove computational hardness formally or exhibit a tractable algorithm.

B. The efficiency of optimized discriminatory access platforms

Despite the difficulty exhibited by the examples above, we can still compute the optimal design for small examples. Thus, we can provide some insight into the efficiency gains of discriminatory access over open access.

To study this question we have performed numerical calculations for settings that have up to four firms and up to three markets. For each combination of markets and firms, e.g., 4 firms and 3 markets, we randomly choose the demand and cost function parameters – $\alpha \sim \text{Uniform}([0, 100])$, $\beta \sim \text{Uniform}([0, 1])$ and $c \sim \text{Uniform}([0, 1])$ – and compute the social welfare for the open access, optimal discriminatory access, and social optimal in 1000 random systems of firms and markets.

Across all these runs, the improvement of the optimal discriminatory access platform over the open access platform was never more than 3%. Further, the improvement of the social optimal over the open access platform was never more than 33% (which matches the inefficiency in the case of one market and one firm).

This is to be expected given Theorem 3, which can be interpreted as an “approximation guarantee” for optimizing the discriminatory access platform design. In particular, using the complete graph guarantees at least 7/16 of the optimal social welfare, and thus of the social welfare achievable under any discriminatory access design (since the optimal social welfare is monotonically increasing with the set of edges).

V. THE INEFFICIENCY OF CONTROLLED ALLOCATION PLATFORMS

A natural response to the potential inefficiency of the open access model is for platforms to optimize the matching of firms to markets directly. It is typically possible for the platform to have considerable knowledge on the firm costs and market demands, which makes the optimization of matching appealing. This leads to the design of controlled allocation platforms.

However, the control over allocations exerted by the platform may create unintended incentives for the firms. Though they cannot strategically choose prices or matches, they still have control over their participation in the platform. In the Cournot setting, this takes the form of strategic choices of production levels.

Our main result in this section highlights that these distorted incentives lead to inefficient market outcomes that can be far worse (in the worst case) than outcomes under open access designs. Further, these distorted incentives can create other challenges, such as the existence of multiple equilibria.

Theorem 8. *A Controlled Allocation Platform can have unbounded price of anarchy. In particular, there is a family of networks where the Nash equilibrium is unique and the price of anarchy is $\Omega(m)$.*

The contrast between Theorem 8 for controlled allocation platforms and Theorem 3 for open access platforms is stark.

Note that the networks exhibited to prove Theorem 8 are simple. They use a single firm with costless production. The construction begins with a single market and then, as markets are added one by one to the system, the parameters of each new market are such that the firm will have no incentive to increase production due to the reallocation under the controlled access, whereas the socially optimal (non-Nash) production level does increase, as does the optimal welfare.²

Before providing a proof of Theorem 8 we first state and prove the key structural lemma we use in the proof.

Lemma 9. *For a firm with costless production, a set of linear demand markets under controlled allocation is equivalent to a single market with a convex, piecewise linear demand curve. Conversely, any convex, decreasing, piecewise linear demand curve with finitely many linear segments can be realized by a set of linear demand markets under controlled access.*

Proof. The characterization of the socially optimal production in Lemma 1, with s fixed highlights that the platform will reallocate this amount to d_1, \dots, d_m such that $\sum_j d_j = s$, and for each market j where $d_j > 0$, p_j is equal to a same price p ; for each market j where $d_j = 0$, it must be that $\alpha_j \leq p$. This shows that, as s increases, the allocation will enter the markets one by one in the order in which α_j decreases.

We say a market becomes *active* when supply starts entering it. For a set of active markets, before the next market becomes active, the marginal increase in supply will be allocated in proportion to $1/\beta_j$ (in order to keep the prices the same). This fully describes the behavior of the platform.

Without loss of generality, assume the markets are ordered such that $\alpha_1 \geq \dots \geq \alpha_m$. From the firm’s point of view, the platform is equivalent to a single market with a piecewise linear demand curve: when the price is between α_1 and α_2 , the rate at which price drops when s increases is β_1 ; for $p \in [\alpha_2, \alpha_3]$, the rate is $1/(\frac{1}{\beta_1} + \frac{1}{\beta_2})$. In general, when the first k

²Note that, when a firm has no production cost, the socially optimal production is always to produce until the price in every market is driven to 0.

markets are active, prices drop at the rate of $(\sum_{j=1}^k \frac{1}{\beta_j})^{-1}$. We call this single demand curve the *aggregate demand curve*.

For a given production level, the area under the aggregate demand curve is equal to the welfare in the original markets. The aggregate demand curve fully characterizes the set of markets under controlled access. Note that whenever a new market joins, the rate at which price drops becomes slower, therefore the aggregate demand curve is always convex.

Conversely, we can show that any convex, decreasing, piecewise linear demand curve consisting of finitely many linear segments is equivalent to a set of linear demand markets under controlled access. We omit the details due to space constraints. \square

Proof of Theorem 8. By Lemma 9, we can focus on constructing an aggregate demand curve. Fix a constant $\lambda \in (0, \frac{1}{2})$. The aggregate demand curve we construct for m markets, $(m \geq 2)$, is

$$p(d) = \max_{0 \leq k < m} (\lambda^k - \lambda^{2k}d).$$

It is not hard to verify that this is the piecewise linear function that connects the following points: $(0, 1)$, $(\frac{1}{1+\lambda}, \frac{\lambda}{1+\lambda})$, $(\frac{1}{\lambda(1+\lambda)}, \frac{\lambda^2}{1+\lambda})$, \dots , $(\frac{1}{\lambda^{m-2}(1+\lambda)}, \frac{\lambda^{m-1}}{1+\lambda})$, $(\frac{1}{\lambda^{m-1}}, 0)$. Being the maximum of a family of decreasing linear functions, $p(d)$ is obviously a convex decreasing function.

We first calculate the optimal social welfare, the area under $p(d)$. The trapezoid whose vertices are $(\frac{1}{\lambda^{k-1}(1+\lambda)}, 0)$, $(\frac{1}{\lambda^k(1+\lambda)}, \frac{\lambda^k}{1+\lambda})$, $(\frac{1}{\lambda^{k+1}(1+\lambda)}, 0)$, $(\frac{1}{\lambda^k(1+\lambda)}, \frac{\lambda^{k+1}}{1+\lambda})$ has area

$$\frac{1}{2} \left(\frac{1}{\lambda^k(1+\lambda)} - \frac{1}{\lambda^{k+1}(1+\lambda)} \right) \left(\frac{\lambda^k}{1+\lambda} + \frac{\lambda^{k+1}}{1+\lambda} \right) = \frac{1-\lambda}{1+\lambda}.$$

There are $m-2$ such trapezoids under $p(d)$, and therefore the socially optimal welfare is $\Omega(m)$. On the other hand, the linear components of $p(d)$ are designed so that producing on any of the linear segment gives a maximal profit of $\frac{1}{4}$ (for the k -th segment, the profit maximizing production level is $\frac{1}{2\lambda^{k-1}}$). The firm is indifferent to best responding to any of the linear segments, and all the production levels $\frac{1}{2\lambda^k}$ for $k = 0, \dots, m-1$ are Nash equilibria. If we push the starting point of $p(d)$ from $(0, 1)$ to $(0, 1+\epsilon)$ for some $\epsilon > 0$, then producing $\frac{1+\epsilon}{2}$ (by responding only to the first market) will be the unique equilibrium, resulting in a social welfare of only $3(1+\epsilon)^2/8$. Therefore the price of anarchy is $\Omega(m)$. \square

The example establishing Theorem 8 highlights that a high production level can be more crucial to the welfare than an ideal distribution among the markets. Additionally, this also begs the question: Can the platform do better by optimizing a different quantity? Given our result in Section III, it is obvious that platforms can do better by doing nothing.

VI. CONCLUSION

This paper studies a trend in platform design today: the move from open access platform designs toward discriminatory access and controlled allocation platform designs. One reason for this shift is a belief that open access markets

are inefficient due to strategic behavior from firms and that limiting access or controlling allocations can lead to more efficient markets. In this paper, our results for networked Cournot markets counter this belief. We show that open access markets have a small constant factor efficiency loss (PoA 16/7), whereas controlled allocation platforms can have unbounded efficiency loss.

This paper represents a first step toward contrasting open access platforms with discriminatory access and controlled allocation platform designs. There are many interesting questions that remain. For example, is it possible to efficiently optimize the allowed connections in the discriminatory access model or is the task NP-hard? Further, our focus has been on social welfare due to the importance of network effects for platforms; however, what if the platform is more concerned with maximizing short term profit than long term growth?

More broadly, increased efficiency is only one motivation for moving to discriminatory access and controlled allocation platforms. Another is to provide a simple, easy-to-use interface for firms and consumers. Participants in open access platforms may suffer from information overload. How does this factor impact incentives and the efficiency of market outcomes?

Finally, this paper focused on matching and allowed pricing to emerge endogenously. Another trend in platform design today, exemplified by Uber, is fine-grained control on pricing. Understanding the interaction between controlled allocation and fine-grained pricing is a challenging open question.

ACKNOWLEDGMENT

This work is supported by grants from NSF AitF-1637598, CNS-1518941, CPS-154471 and the Ronald & Maxine Linde Institute of Economic and Management Science at Caltech.

REFERENCES

- [1] “Taming the beasts,” *The Economist*, 2016. [Online]. Available: <http://www.economist.com/news/business/21699465-european-governments-are-not-alone-wondering-how-deal-digital-giants-taming>
- [2] M. Abolhassani, M. H. Bateni, M. Hajiaghayi, H. Mahini, and A. Sawant, “Network cournot competition,” in *International Conference on Web and Internet Economics*. Springer, 2014, pp. 15–29.
- [3] D. Abreu and M. Manea, “Bargaining and efficiency in networks,” *Journal of Economic Theory*, vol. 147, no. 1, pp. 43–70, 2012.
- [4] D. Acemoglu, K. Bimpikis, and A. Ozdaglar, “Price and capacity competition,” *Games and Economic Behavior*, vol. 66, no. 1, pp. 1–26, 2009.
- [5] E. Adida and N. Ratisoontorn, “Consignment contracts with retail competition,” *European Journal of Operational Research*, vol. 215, no. 1, pp. 136–148, 2011.
- [6] E. Anshelevich and S. Sekar, “Price competition in networked markets: How do monopolies impact social welfare?” in *Web and Internet Economics*. Springer, 2015, pp. 16–30.

- [7] M. Armstrong, "Competition in two-sided markets," *The RAND Journal of Economics*, vol. 37, no. 3, pp. 668–691, 2006.
- [8] S. Banerjee, S. Gollapudi, K. Kollias, and K. Munagala, "Segmenting two-sided markets."
- [9] K. Bimpikis, S. Ehsani, and R. Ilklic, "Cournot competition in networked markets," in *EC*, 2014, p. 733.
- [10] S. Bose, D. W. Cai, S. Low, and A. Wierman, "The role of a market maker in networked cournot competition," in *53rd IEEE Conference on Decision and Control*. IEEE, 2014, pp. 4479–4484.
- [11] J. I. Bulow, J. D. Geanakoplos, and P. D. Klemperer, "Multimarket oligopoly: Strategic substitutes and complements," *Journal of Political economy*, vol. 93, no. 3, pp. 488–511, 1985.
- [12] D. Cai, S. Bose, and A. Wierman, "On the role of a market maker in networked cournot competition," 2016.
- [13] S. Cannon and L. H. Summers, "How uber and the sharing economy can win over regulators," *Harvard business review*, vol. 13, 2014.
- [14] S. Chawla and T. Roughgarden, "Bertrand competition in networks," in *Algorithmic Game Theory*. Springer, 2008, pp. 70–82.
- [15] L. Chen, A. Mislove, and C. Wilson, "Peeking beneath the hood of uber," in *Proceedings of the 2015 ACM Conference on Internet Measurement Conference*. ACM, 2015, pp. 495–508.
- [16] —, "An empirical analysis of algorithmic pricing on amazon marketplace," in *Proceedings of the 25th International Conference on World Wide Web*. International World Wide Web Conferences Steering Committee, 2016, pp. 1339–1349.
- [17] G. Christodoulou, A. Kovács, and M. Schapira, "Bayesian combinatorial auctions," *J. ACM*, vol. 63, no. 2, pp. 11:1–11:19, Apr. 2016.
- [18] D. Condorelli and A. Galeotti, "Bilateral trading in networks," 2012.
- [19] M. Dinerstein, L. Einav, J. Levin, and N. Sundaresan, "Consumer price search and platform design in internet commerce," National Bureau of Economic Research, Tech. Rep., 2014.
- [20] L. Einav, T. Kuchler, J. Levin, and N. Sundaresan, "Assessing sale strategies in online markets using matched listings," *American Economic Journal: Microeconomics*, vol. 7, no. 2, pp. 215–247, 2015.
- [21] M. Elliott, "Inefficiencies in networked markets," *American Economic Journal: Microeconomics*, vol. 7, no. 4, pp. 43–82, 2015.
- [22] A. Fradkin, "Search frictions and the design of online marketplaces," *NBER Working Paper*, 2014.
- [23] A. Gandini, "Working online: An exploration of social recruiting and digital marketplaces," in *The Reputation Economy*. Springer, 2016, pp. 69–83.
- [24] C. L. Guzmán, "Price competition on network," Tech. Rep., 2011.
- [25] K. Hasker and R. Sickles, "ebay in the economic literature: Analysis of an auction marketplace," *Review of Industrial Organization*, vol. 37, no. 1, pp. 3–42, 2010.
- [26] R. Ilklic, "Cournot competition on a network of markets and firms," 2009.
- [27] R. Johari and J. Tsitsiklis, "Efficiency loss in cournot games," 2005, publication 2639.
- [28] M. Kenney and J. Zysman, "Choosing a future in the platform economy: the implications and consequences of digital platforms," in *Kauffman Foundation New Entrepreneurial Growth Conference*, 2015.
- [29] —, "The rise of the platform economy," *Issues in Science and Technology*, vol. 32, no. 3, pp. 61–69, 2016.
- [30] E. Koutsoupias and C. H. Papadimitriou, "Worst-case equilibria," in *STACS 99, 16th Annual Symposium on Theoretical Aspects of Computer Science, Trier, Germany, March 4-6, 1999, Proceedings*, 1999, pp. 404–413.
- [31] M. Manea, "Bargaining in stationary networks," *The American Economic Review*, vol. 101, no. 5, pp. 2042–2080, 2011.
- [32] K. Miller, "Etsy: A site for artisans takes off," *Business Week June*, vol. 12, 2007.
- [33] F. Nava, "Efficiency in decentralized oligopolistic markets," *Journal of Economic Theory*, vol. 157, pp. 315–348, 2015.
- [34] T. Roughgarden, "Intrinsic robustness of the price of anarchy," in *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC, Bethesda, MD, USA, May 31 - June 2, 2009*, pp. 513–522.
- [35] —, "The price of anarchy in games of incomplete information," in *ACM Conference on Electronic Commerce, EC '12, Valencia, Spain, June 4-8, 2012*, pp. 862–879.
- [36] T. Roughgarden and É. Tardos, "How bad is selfish routing?" *J. ACM*, vol. 49, no. 2, pp. 236–259, 2002.
- [37] B. Stone, "Amid the gloom, an e-commerce war," *New York Times*, 2008.
- [38] A. Sundararajan, "What airbnb gets about culture that uber doesn't," *Harvard Business Review*, 2014.
- [39] V. Syrgkanis and E. Tardos, "Composable and efficient mechanisms," in *STOC*, 2013.
- [40] A. Vetta, "Nash equilibria in competitive societies, with applications to facility location, traffic routing and auctions," in *43rd Symposium on Foundations of Computer Science (FOCS 2002), 16-19 November 2002, Vancouver, BC, Canada, Proceedings*, 2002, p. 416.
- [41] Y. Wang, L. Jiang, and Z.-J. Shen, "Channel performance under consignment contract with revenue sharing," *Management science*, vol. 50, no. 1, pp. 34–47, 2004.
- [42] E. G. Weyl, "A price theory of multi-sided platforms," *The American Economic Review*, vol. 100, no. 4, pp. 1642–1672, 2010.
- [43] A. White and E. G. Weyl, "Insulated platform competition," *Available at SSRN 1694317*, 2016.