

CAMBRIDGE WORKING PAPERS IN ECONOMICS
CAMBRIDGE-INET WORKING PAPERS

Searching for Results: Optimal Platform Design in a Network Setting

George
Charlson
University of
Cambridge

Abstract

Online platforms that link buyers and sellers are able to shape the pattern of observation between market participants. We model buyer-seller interactions as a series of bipartite graphs, which are each realised with a probability chosen by the platform owner. Prominent sellers disproportionately increase competition, which decreases prices. To maximise profit, the platform owner ensures that the size of the neighbourhood of each buyer is the same in each state of the world and randomises buyer observation across all sellers on the platform, a result that still holds even when buyers are heterogeneous. When products are vertically differentiated, the platform owner faces a trade-off between biasing observation towards high-quality products and increasing competition. The extent to which platforms highlight high-quality products and the number of sellers depends on the characteristics of the market(s) in which they operate.

Reference Details

20118 Cambridge Working Papers in Economics
2020/52 Cambridge-INET Working Paper Series

Published 2 December 2020
Revised 13 April 2021

Key Words Networks, strategic interaction, network games, interventions, industrial organisation, platforms

JEL Codes D20, L20

Websites www.econ.cam.ac.uk/cwpe
www.inet.econ.cam.ac.uk/working-papers

Searching for Results: Optimal Platform Design in a Network Setting

George Charlson

Cambridge University

gc556@cam.ac.uk*

April 13, 2021

Abstract

Online platforms that link buyers and sellers are able to shape the pattern of observation between market participants. We model buyer-seller interactions as a series of bipartite graphs, which are each realised with a probability chosen by the platform owner. Prominent sellers disproportionately increase competition, which decreases prices. To maximise profit, the platform owner ensures that the size of the neighbourhood of each buyer is the same in each state of the world and randomises buyer observation across all sellers on the platform, a result that still holds even when buyers are heterogeneous. When products are vertically differentiated, the platform owner faces a trade-off between biasing observation

*I would like to thank Alex Teytelboym and Mark Armstrong for their invaluable contributions to this paper. Thank you also to Sanjeev Goyal, Mihai Manea, Matthew Elliott, Francis Bloch, Daniel Quigley, Meg Meyer and participants at conferences at Stony Brook Game Theory Festival, the Paris School of Economics and the Northwestern-Paris-Oxford theory seminar for their comments on this paper.

towards high-quality products and increasing competition. The extent to which platforms highlight high-quality products and the number of sellers depends on the characteristics of the market(s) in which they operate.

KEYWORDS: Networks, strategic interaction, network games, interventions, industrial organisation, platforms.

1 Introduction

Owners of large online marketplaces, like Airbnb or Amazon Marketplace, choose the ordering of sellers and number of results, both per search and in total, shown to buyers on the platform, and therefore determine the pattern of buyer-seller observations. On these platforms, buyers only observe a subset of sellers (Kim, Albuquerque and Bronnenberg, 2010 and Ringel and Skiera, 2016). Consumers engage far less with results on the second search page of online search engines than the first, and so on (Smith and Brynjolfsson, 2001, Baye et al, 2009 and Baye et al 2016), to the extent that the bottom item on the first page of results drew 140% more clicks than the first item on the second page (Chitika, 2013).

In choosing the search environment, platforms choose which buyers observe which sellers and how many sellers effectively compete with one another. A natural question is: what is the profit-maximising pattern of buyer-seller interactions? More specifically, how many sellers should compete for a given group of buyers and which sellers should those buyers observe?

We find that sellers being relatively prominent on a platform is costly in terms of aggregate platform profit, as prominence leads to competition being more concentrated among the prominent sellers, which has disproportionately large effects on prices across the network, even for those sellers who are less likely to be observed.¹

¹ Prominence here refers to the probability that buyers observe a given seller, and is therefore modeled differently to e.g. Armstrong, Vickers and Zhou (2009) and Armstrong and Zhou (2011), where prominent sellers are observed first and either accepted or rejected in a search model.

As a result, the platform has an incentive to display a fixed number of sellers to each group of buyers. Under the assumptions that there is some probability that buyers are also able to find sellers through means other than the platform’s internal search process and products are not vertically differentiated, this incentive results in it being optimal for the platform to randomise which set of sellers buyer groups observe, such that each buyer observes each seller with equal probability. This result continues to hold, even when buyers are heterogeneous with respect to how substitutable they view the products on offer.

Furthermore, when products are vertically differentiated, our model suggests that while the platform has an incentive to bias observation towards higher quality sellers, even the lowest quality sellers are observed with some probability if quality dispersion is not too high. If platforms were to limit search results for some behavioural reason² then high quality sellers would always be shown to each buyer in order to maximise demand, which is not necessarily optimal here.

Markets in which quality dispersion is high tend to result in higher profits, as the platform is able to increase the probability that high-quality sellers are observed. Whether the number of sellers observed by buyers is higher or lower in markets with lower quality dispersion depends on product substitutability. If products are not that substitutable, then the competition effects associated with a high-quality seller being observed dominate, and so fewer sellers are observed on average when quality dispersion is high.

However, when products are highly substitutable, there are few competitors in the market to start with, and the cost of missing out on sellers near the top of the quality distribution is in expectation higher when quality is highly dispersed, so the platform shows more sellers to the buyers in this setting compared to the case where seller quality is more similar.

²For example, rational inattention, see Hefti and Heinke (2015), or choice overload, see Chernev et al, (2015) for a meta-analysis and e.g. Nager and Gandotra (2016) and Moser et al. (2017) for applications to an online setting.

In terms of our approach, we develop a model of groups of buyers and sellers interacting on a network, with sellers competing on price.³ We eschew network bargaining models (Kranton and Minehart, 2001, Corominas-Bosch, 2004 and Polanski, 2007) in favour of a model in which each seller sets a single price for all buyers. This approach makes it possible to model an arbitrary set of observation patterns, characterise optimal prices and platform structure, even when there a large number of buyers and sellers that are potentially heterogeneous.

Our set-up allows us to characterise the equilibrium price of a seller as a decreasing function of the seller’s Bonacich centrality in the network, a result consistent with the broader games on networks literature (Ballester, Calvò-Armengol and Zenou, 2006, and Bramoullé, Kranton and D’Amours, 2014), and the application of this literature to an IO context (Bimpikis, Ehsani and Ilkiliç, 2018 and Elliott and Galeotti, 2019). In this setting, a seller’s centrality captures the amount of competition they face on the platform: more central players face more competition from other more central players, which in turn leads to lower prices.

From a theory perspective, our contribution is that the platform owner chooses the distribution of possible bipartite networks in order to maximise total profits given the equilibrium behaviour of the sellers. This methodology nests approaches in which a central planner can (costlessly) choose the weighted or unweighted edges of a deterministic network (see Sun, Zhaou and Zhou, 2021 and Li, 2020), but allows for further flexibility by allowing for the realisation of particular link structures to be correlated probabilistically, which, as we shall see, is generally profit increasing in this particular setting.⁴

³Our set-up thus bears a relation to work on captive buyers; see, Ireland (1993), McAfee (1994), Salvadori (2013) and Armstrong and Vickers (2019). This literature does not model buyer-seller interactions as a graph, but it does consider a pattern of seller competition that can be summarised in this way.

⁴Our methodology thus differs from other approach to the interventions in networks, for example the intervention by targeting specific nodes (of which there is a large literature, including, e.g. Ballester, Calvò-Armengol and Zenou, 2006; Galeotti and Goyal, 2009, Bloch and Querou, 2013;, Leduc, Jackson, and Johari, 2017; Akbarpour, Malladi, and Saberi, 2020, Belhaj, Deroïan and Safi, 2020) or changes to private returns on investment (see Galeotti, Golub and Goyal, 2020).

Given that the platform designer chooses an optimal platform in order to elicit a particular equilibrium action profile, our analysis shares some similarities from approaches that apply persuasion to the context of platforms, examples of which include where platforms segment markets by withholding information (Armstrong and Zhou, 2020 and Elliott, Galeotti and Koh, 2020) or use seller ratings as a means of biasing search results (Charlson, 2021). The latter approach shares some similarities with the analysis here, but our analysis gives the platform full control over the pattern of buyer observations and does not model seller ratings explicitly.

Our analysis suggests that competition authorities should examine intra-platform competition, in addition to inter-platform competition, which has been emphasised in the traditional platforms literature (Tirole and Rochet, 2003, Armstrong, 2006 and Tan and Zhou, 2019). Regulating the internal structure of the networks that underpin large online platforms may reduce the extent to which consumers are harmed by the formation of monopolistic platforms, by increasing choice and lowering prices.

2 Motivating example

Large, online platform owners must design platforms in which many sellers compete for groups of buyers. If sellers can only set one price, then network design has implications for the nature of competition between sellers for different buyers. If a platform owner can affect which buyers observe which sellers and sellers only set one price, then the platform owner faces a trade-off between more sellers being observed, increasing demand, and the resultant increase in competition.

As an example, suppose the probability that two buyers observe two sellers is strictly between 0 and 1. Then all the possible ex-post market structures (ignoring the case where no buyer observes any sellers) are shown in Figure 1.




Figure 1: The black nodes represent sellers, the white nodes represent buyers. Assuming sellers and buyers are identical, these configurations represent all possible market types when there are two sellers and two buyers.

Consider the nature of total profits and competition if prices were set **after** the realisation of a network structure. Market structures on the bottom and to the right of the diagram exhibit less competition, but profits are lost because the buyers will be assumed throughout to purchase at least some goods from any seller they observe. This issue is reduced in market structures above and to the left of the diagram; however, there, the two sellers are in more direct competition with one another, which reduces profit due to both sellers setting a lower price.

3 Model

Sellers, buyers and the platform owner

Suppose there is a finite set, \mathbf{B} , whose elements are “consumer segments”, in the sense that they are a finite mass of consumers who are assumed to share some trait, such as geographical location, age demographic, occupation, etc. We use n to denote the

number of consumer segments.

Similarly, let \mathcal{S} be a finite set of sellers, where $|\mathcal{S}| = m$. Sellers each sell a single type of completely divisible good, and each seller's good is an imperfect substitute for each of the goods.

Sellers and consumer segment interact on a platform, with each consumer segment observing a subset of S . These observations generate a network $G_\tau = (B \cup S)$, with an observation between a consumer segment i and seller j represented by an edge, $E_{ij} \in G_\tau$. We assume that the graph generating process is stochastic in the sense that there is a probability $\theta_\tau \in [0, 1]$ that a graph G_τ is generated for every possible $m - n$ bipartite graph, and hence $\sum_\tau \theta_\tau = 1$. Let $\boldsymbol{\theta}$ denote a vector whose τ th entry is θ_τ .

Consider a simple example of the above set-up, with three sellers (X, Y and Z) and three consumer segments (1, 2 and 3). Suppose the two graphs that can be realised with some positive probability are: (1) the complete graph, in which each buyer observes each seller and; (2) the graph depicted on the right of Figure 2, in which sellers Y and Z compete for consumer segment 2 but the other two segments are captive. Assume both of these graphs are realised with equal probability.




Figure 2: A case where two graphs can be realised.

Let \mathbf{p} denote a $m \times 1$ vector whose j th component is $p_j \in \mathbb{R}_+$, the price of j 's good. We assume that sellers set prices prior to the realisation of the links in the network,

but with full knowledge of the vector $\boldsymbol{\theta}$. Let a, γ_j, c_i be strictly positive scalars and μ_{ij}^τ be such that if there exists an edge between i, j , $E_{ij} \in G_\tau$, then $\mu_{ij}^\tau = 1$ and 0 otherwise. Consumer segment i 's expected demand function, $x_{ij}(\cdot) : \mathbb{R}^m \rightarrow \mathbb{R}_+$, for product j can be expressed as follows:⁵

$$\mathbb{E}[x_{ij}(\mathbf{p}; \boldsymbol{\theta})] = \sum_{\tau} \theta_{\tau} \mu_{ij}^{\tau} (a\gamma_j - ap_j + c_i \sum_{k \neq j} \mu_{ik}^{\tau} (p_k - \gamma_k)),$$

if the expression on the right is positive, and otherwise $\mathbb{E}[x_{ij}(\mathbf{p}; \boldsymbol{\theta})] = 0$. The parameter γ_j is a measure of the quality of the seller j . As $c_i > 0$, each product is a gross substitute for every other product.

Each seller, j , is assumed not to be able to price discriminate across buyers, and hence sets a single price $p_j \in R_+$. Sellers compete with one another on price, and set prices simultaneously. Seller j 's expected aggregate demand function is then defined $\mathbb{E}[x_j(\mathbf{p}; \boldsymbol{\theta})] = \sum_{i=1}^n \mathbb{E}[x_{ij}]$. Profit is similarly defined $\mathbb{E}[\pi_j(\mathbf{p}; \boldsymbol{\theta})] = \mathbb{E}[p_j x_j(\mathbf{p}; \boldsymbol{\theta})]$. Each seller's maximisation problem can be expressed:

$$\max_{p_j} \mathbb{E}[\pi_j(p_j, \mathbf{p}_{-j}; \boldsymbol{\theta})].$$

Let $\Gamma(\boldsymbol{\theta})$ represent the simultaneous move m -player game played on a network G with payoffs as specified above and strategy space R_+ . The platform owner has the following profit function:

$$\mathbb{E}[\pi_P(\mathbf{p}; \boldsymbol{\theta})] = \chi \sum_{j=1}^m \mathbb{E}[\pi_j(\mathbf{p}; \boldsymbol{\theta})],$$

where $0 < \chi < 1$.

Given the implicit assumption here that sellers have no marginal cost, the platform

⁵We show that this assumption is an approximation of the linear demand curve that is generated from a quadratic, quasi-linear demand curve in the Appendix.

owner’s profit function is equivalent to taking a proportion of seller revenues. A number of large platforms (e.g. Airbnb, Amazon Marketplace, Etsy and eBay) take a proportion of total revenue, and hence the dynamics on these platforms are captured by this assumption. If sellers were to have equal constant marginal costs, then the qualitative results of the model would remain unchanged.

An extension not considered here would be the case where sellers had different marginal costs, as the two models of platform revenue would potentially result in different optimal network structures.⁶ We also do not consider a fixed-fee based model of platform in our analysis: such a model seems less commonly used in the kind of large-scale buyer-seller platforms that we wish to analyse here.

The search environment

We consider a search environment in which the platform owner chooses a baseline distribution of networks, θ_b , optimising their above profit function. If a seller, j , is not observed by a buyer, i , in a baseline network G generated with positive probability by θ_b , then there is a probability, $v \in [0, 1) \forall i, j$, of “external observation”, where i observes j even though $E_{ij} \notin G$. Together, the baseline distribution of networks and these external observation probabilities induce a probability vector, θ .

The baseline distribution of networks, θ_b , reflects the groups of sellers presented to consumer segments by the platform owner, when, for example, they enter a search term on the platform’s internal search engine.

The external observation probability v captures the notion that buyers may find the seller through some other means, like an external search engine or word of mouth, even if they are not shown the seller by the platform intentionally through the plat-

⁶The difference between profit and revenue on these platforms is complicated by the fact that some of the costs are borne by consumers (i.e. postage and packaging), and platforms do not necessarily take a proportion of these elements of seller revenue. The recent changes to the Amazon algorithm, which reportedly takes into account a product’s profitability to the platform reflects this nuance (see, WSJ, 2019), and supports the possibility of seller profit being a component of the platform’s objective function.

form’s search process.⁷ If a graph G is realised with probability θ for the distribution $\boldsymbol{\theta}_b$, then its realisation probability for the distribution $\boldsymbol{\theta}$ is $\theta(1 - v)^{(mn-\phi)}$, where ϕ is the number of edges in G . The graph $G + E_{ij}$ where $E_{ij} \notin G$ is realised with probability $\theta v(1 - v)^{(mn-\phi)-1}$ and so on.

Of course, the existence of such leakages from the internal search process may generate an incentive to limit the number of sellers allowed to enter the platform. We leave this decision unmodelled, though note that the attractiveness such a strategy would depend on the platform’s owner’s ability to screen sellers for e.g. quality before they enter. If the platform is fully in control of the observation process, then $v = 0$.

Given the platform owner’s profit function above, they are assumed to choose $\boldsymbol{\theta}_b$ in order to solve the following maximisation problem:

$$\max_{\boldsymbol{\theta}_b} \mathbb{E}[\pi_P(\mathbf{p}; \boldsymbol{\theta}_b, v)].$$

We assume throughout that $\boldsymbol{\theta}$ is common knowledge, and prices are hence set after the realisation of $\boldsymbol{\theta}$, but, as stated previously, prior to the realisation of the actual observation network.

4 Equilibrium characterisation

We characterise the equilibrium price setting behaviour for a given vector of graph probabilities $\boldsymbol{\theta}$. Define $\beta_j(\boldsymbol{\theta}) := \sum_i \sum_\tau \theta_\tau \mu_{ij}^\tau$, as the expected number of consumer segments that observe j and $\hat{c}_{jk}(\boldsymbol{\theta}) = \frac{1}{2a\beta_j} \sum_\tau \theta_\tau \sum_i \sum_j \sum_k c_i \mu_{ij}^\tau \mu_{ik}^\tau$, which is a measure of the strength of the connection between j and k because it measures the weighted link between the sellers and shared buyers.

Define the *competition network*, $G_S(\boldsymbol{\theta})$ as a projection of $G(\boldsymbol{\theta})$, where the edge

⁷To get some perspective on the proportion of users who receive information using external sources, more than 10% of users find products on Amazon via means other than Amazon’s internal search process (Nasdaq, 2018).

between sellers j and k has the weight $\hat{c}_{jk}(\boldsymbol{\theta})$. The competition network of the probability vector $\boldsymbol{\theta}$ that generates the two graphs in Figure 2 with equal probability is shown in Figure 3 below.




Figure 3: Transforming the network G into the competition network G_S . Here it is assumed that $a = 1$ and $c = 0.25$ for each buyer and $\gamma_j = 1$ for all j .

Let γ_i denote the smallest component of the vector $\boldsymbol{\gamma}$. Throughout, we make the following assumption:

$$(A1) : a > \frac{\sum_i c_i \sum_{j \neq i} \gamma_j}{n \gamma_i} \quad \forall i$$

Define $R_S(\boldsymbol{\theta})$ as a zero-diagonal matrix of a network $G_S(\boldsymbol{\theta})$ with elements $\hat{c}_{jk}(\boldsymbol{\theta})$ and $C(\boldsymbol{\theta})\mathbf{1}_m = \sum_{k=0}^{\infty} R_S^k(\boldsymbol{\theta})\mathbf{1}_m$, which is the Bonacich centrality measure of the network $G_S(\boldsymbol{\theta})$. The following Proposition, which characterises the equilibrium price vector, then holds:

Proposition 1. *If (A1) holds, then the game $\Gamma(\boldsymbol{\theta})$ has a unique Nash equilibrium in pure strategies, which is the equilibrium price vector:*

$$\mathbf{p}^*(\boldsymbol{\theta}) = \boldsymbol{\gamma} - \frac{1}{2}C(\boldsymbol{\theta})\boldsymbol{\gamma}.$$

A seller's equilibrium price is decreasing in their centrality in $G_S(\boldsymbol{\theta})$. Sellers who are connected to more isolated consumer segments face relatively less competition than sellers who are largely connected to segments with different goods to choose from,

and therefore are able to set a higher price in equilibrium than other sellers. The above expression implies seller j 's price is increasing in γ_j but is decreasing in every other element of the vector $\boldsymbol{\gamma}$.

The assumption (A1) provides a restriction on each $\hat{c}_{jk}(\boldsymbol{\theta})$, which measures the substitutability of products relative to the effect own price has on demand. Specifically, (A1) guarantees both that (a): $x_{ij}(\boldsymbol{\theta}) > 0$ for all i, j pairs at equilibrium and (b) that $\mathbf{I} - \lambda R_S(\boldsymbol{\theta})$ is strictly diagonally dominant for all $\boldsymbol{\theta}$, and thus positive definite. Jointly, these two facts guarantee that the Nash equilibrium of the game both exists and is unique for any graph structure.

Recall that sellers target segments of consumers in the model, meaning that for $x_{ij}(\boldsymbol{\theta}) > 0$ to hold, it only need be that a small subset of consumers within the potentially large group, i , has positive demand for j 's product at equilibrium.

5 The role of prominence

To better understand some important properties of our price equilibrium, we assume in this section that there is a probability that i observes j of $w_{ij} = w_j$ for all i and these probabilities are determined by nature (rather than the platform) in the following sense: each w_j is the realisation of a random variable \tilde{w}_j according to the symmetric probability distribution, Λ , which is bounded such that $w_j \in [0, 1]$. The random variables \tilde{w}_j are independently and identically distributed, and generate the distribution over bipartite networks, $\boldsymbol{\theta}$. $\boldsymbol{\theta}$ is assumed to be common knowledge.

The observation probabilities of the consumer segments can be thought of as a measure of “prominence” in the sense that they capture the likelihood that the seller is observed by a given buyer. Seller j 's prominence in the network potentially increases profits as a result of increasing the probability of sales, but at the same time it imposes a cost on the rest of the network by increasing competition, reducing prices of every seller, including for the more prominent seller.

Recall that the centrality vector in G_S can be expressed $C(\boldsymbol{\theta})\mathbf{1} = \sum_{k=0}^{\infty} R_S^k(\boldsymbol{\theta})\mathbf{1}$. It follows that $\frac{\partial^2 C_j(\boldsymbol{\theta})}{\partial^2 w_j} \geq 0$. Recall that prices are falling in the centrality of the sellers in this setting. Hence, increasing w_j imposes a cost upon the platform owner because the centrality measure has a feedback effect such that increasing an observation probability w_j reduces j 's price, which reduces every other seller's price, which then reduces j 's price and so on. This feedback effect, which is a feature of the Bonacich centrality measure, is increasing as the centralities of the sellers in G_S become larger.

To model this discussion formally, suppose $\tilde{w}_j \sim \Lambda_1$ and let Λ_2 be a mean-preserving spread of Λ_1 such that when $\tilde{w}'_j \sim \Lambda_2$ constructed in the following way such that $\tilde{w}'_j = \tilde{w}_j + \epsilon_j$, where ϵ_j is symmetrically distributed and has mean 0, and is bounded such that $\tilde{w}'_j \in [0, 1]$. Let $\tilde{\boldsymbol{\theta}}_k$ denote the random probability vector generated by the distribution Λ_k . Then the following result holds:

Theorem 1. *Suppose $\gamma_j = \gamma \forall j$. Then, $E[p_j(\tilde{\boldsymbol{\theta}}_1)] > E[p_j(\tilde{\boldsymbol{\theta}}_2)]$ and $E[\pi_P(\tilde{\boldsymbol{\theta}}_1)] > E[\pi_P(\tilde{\boldsymbol{\theta}}_2)]$.*

The expected number of matches (i.e. $E[\sum_j \tilde{w}_j]$) is the same for both probability distributions. Hence, any differences in expected profit between the two are the result of differences in the expected prices.

As the quality vector $\boldsymbol{\gamma}$ is independent of centrality in this case, the expected price vector can be denoted: $E[\mathbf{p}^*(\tilde{\boldsymbol{\theta}})] = \boldsymbol{\gamma} - \frac{1}{2}E[C(\tilde{\boldsymbol{\theta}})]\boldsymbol{\gamma}$. Recalling that $\frac{\partial^2 C_k}{\partial^2 w_j} \geq 0$ and that the \tilde{w}_j s are independent of one another, we show in the Appendix that $E[C_j(\tilde{\boldsymbol{\theta}}_2)] > E[C_j(\tilde{\boldsymbol{\theta}}_1)] \forall j$. This implies that expected prices are lower in the case where each \tilde{w}_j is distributed according to the mean preserving contraction $\tilde{\boldsymbol{\theta}}_2$. Intuitively, this result is driven by the fact that high realisations of an observation probability \tilde{w}_j result in a disproportionately low price compared to low realisations of \tilde{w}_j .

Furthermore, as observation probabilities are independent of each other, if they have a distribution of $\tilde{\boldsymbol{\theta}}_2$ it results in a higher probability that two or more sellers are

prominent. We refer to the case where there is relatively intense competition between a subset of the sellers on the network as one in which competition is concentrated. There being a higher-than-average probability that two sellers compete with one another drives their own prices down, which propagates across the network.

6 Optimal networks

In this section, we use the equilibrium characterisation in Section 4 and the analysis in Section 5 to show which networks are optimal for both consumers and the platform.

Consumer surplus and profits

We first characterise the networks that maximise consumer surplus. Define the expected consumer surplus of consumer segment i for a given equilibrium price vector \mathbf{p}^* and demand vector \mathbf{x}_i^* is as follows:

$$E[CS_i(\mathbf{x}_i^*; \mathbf{p}^*)] = \frac{1}{2} \sum_{\tau} \theta_{\tau} \left[\sum_{j=1}^m \mu_{ij}^{\tau} x_{ij}^* (\gamma_j + \sum_{k=1}^m \mu_{ik}^{\tau} \frac{c_i}{a} (p_k^* - \gamma_k) - p_j^*) \right].$$

Define $CS(\mathbf{x}^*; \mathbf{p}^*) := \sum_i CS_i(\mathbf{x}_i^*; \mathbf{p}^*)$, where \mathbf{x}^* is an $m \times n$ matrix whose ij th component is $x_{ij}^*(\mathbf{p}^*)$. As $I - \lambda R_S$ is diagonally dominant by (A1), it is clear from the above expression that the expected value of each $CS_i(\mathbf{x}_i^*; \mathbf{p}^*)$ is falling in \mathbf{p}^* . It is also straightforward to show that expected consumer surplus, ceteris paribus, is increasing in the expected number of connections in G a buyer has. Define G_c as the complete graph, in which each consumer segment observes each seller. Let $\boldsymbol{\theta}_c$ denote the probability vector in which the complete graph G_c is yielded with probability 1. The following Proposition holds:

Proposition 2. *For any probability vector $\boldsymbol{\theta}_j \neq \boldsymbol{\theta}_c$, $E[CS(\mathbf{x}_i^*; \mathbf{p}^*) | \boldsymbol{\theta} = \boldsymbol{\theta}_c] > E[CS(\boldsymbol{\theta}) | \boldsymbol{\theta} = \boldsymbol{\theta}_j]$.*

The centrality of each agent is at its maximum for a given number of buyers and sellers when the network is complete. Intuitively, when the network is complete, each

buyer is competed for by each seller. A complete network maximises competition, which reduces the equilibrium price level of each seller.

While consumer surplus is maximised by the complete network, it does not necessarily maximise the platform’s profit. To maximise profit, the platform owner must choose a network structure that maximises the expected number of sellers each consumer segment observes while accounting for the constraint that increasing observability decreases prices. Proposition 3 formalises the above intuition:

Proposition 3. *For all γ , there exists a $\bar{c} \in \mathbb{R}_+$ such that if $c_i > \bar{c}$ for at least one i , then in any solution to the platform owner’s maximisation problem, θ^* , $\theta_c^* < 1$.*

As products on the platform become increasingly substitutable to a consumer segment, i , the owner’s incentive to reduce the number of sellers observed by i increases. If c_i becomes too large, the platform owner will show i fewer products, reducing expected sales, but increasing prices and profit.

Taken literally, Proposition 3 indicates that if products are sufficiently substitutable, platforms will or should intentionally display only a subset of sellers, hiding some from consumers entirely. It is unclear how prevalent the hiding of products on the platforms of interest are, but the Proposition captures the idea that there is a theoretical incentive for platforms to engage in such behaviour.

At the very least, the Proposition highlights the incentive from a platform perspective to reduce competition between sellers, and such an incentive could shape decisions such as the number of results displayed per search page. Chitika (2013) suggests that buyer observation is discontinuous, in the sense that the last search result on the first page of results is far more likely to be observed than the first result on the second page, and so displaying fewer search results per page would reduce the effective number of results observed by most consumers.

Profit-maximizing graphs with homogeneous consumers and sellers

The above analysis indicates that the platform owner has an incentive to reduce both the number of products observed and the extent to which sellers are prominent. We can use these observations to understand which networks are profit maximising for the platform when they are able to choose the baseline distribution of networks. Throughout this section, we assume that $c_i = c \forall i$, and $\gamma_j = \gamma \forall j$.

The platform owner's desire to reduce seller prominence implies the following result:

Proposition 4. *Any solution, θ_b^* , to the platform owner's maximisation problem, induces a θ^* and $G_S(\theta^*)$ such that $C_j(\theta^*) = C_k(\theta^*)$ for all j, k pairs where $\beta_j(\theta^*), \beta_k(\theta^*) > 0$.*

To understand why Proposition 4 holds, we define the concept of a neighbourhood switch. Take a graph G_τ where the τ th component of θ_b is $\theta_\tau > 0$. Define G_{jk} as a graph which is the result of performing a neighbourhood switch between two sellers j and k in G , such that for any i where $E_{ij} \in G$ and $E_{ik} \notin G$, $E_{ij} \notin G_{jk}$ and vice versa. Such a switch is depicted in Figure 4.




Figure 4: A neighbourhood switch between sellers X and Y .

Define a vector θ_{jk} where the probability that the graph G_{jk} is realised is equal to

θ_τ for each graph G_τ which has a positive component in the vector $\boldsymbol{\theta}$. Now define another probability vector, $\hat{\boldsymbol{\theta}}_j$, as follows: $\hat{\boldsymbol{\theta}}_j := (1 - (m - 1)\varepsilon)\boldsymbol{\theta} + \sum_k \varepsilon\boldsymbol{\theta}_{jk}$.

Let $C_j(\boldsymbol{\theta})$ be (jointly one of) the smallest component(s) of the vector $\mathbf{C}(\boldsymbol{\theta})\mathbf{1}_m$. Then $\hat{\boldsymbol{\theta}}_j$ involves there being some positive probability that each seller k (who by definition are weakly more central than j) will face the competition faced by j in every graph realised with positive probability. Due to seller prominence being costly to the platform, such a switch increases k 's price more than it decreases j 's price. As such a switch is performed on every seller, prices rise overall, and each price rise also has further second-order price effects, as prices are strategic complements. It follows that $\pi_P(\hat{\boldsymbol{\theta}}_j) > \pi_P(\boldsymbol{\theta})$.

Hence, the optimal competition graph structure is one in which each active seller is as central as every other active seller. If this is not the case, then the platform owner can always find a marginal re-allocation that increases the expected number of consumer segments observing the higher priced seller and increases prices across the network.

Proposition 4 does not fully characterise the optimal solution to the platform owner's problem. Instead, it provides a condition under which a graph $G_S(\boldsymbol{\theta})$ is the result of the platform owner's maximisation problem. It is possible to use this result to map the optimal set of competition graphs onto corresponding bipartite observation graphs.

Let φ_i^τ and $\hat{\varphi}_i(\boldsymbol{\theta}^*)$ denote the number of sellers consumer segment i observes in the graph G_τ and the expected number of sellers i observes across all graphs for the vector $\boldsymbol{\theta}^*$. We show in that each consumer segment's number of observations should be centered closely around this average in any optimal solution:

Theorem 2. *Any solution to the platform owner's problem, $\boldsymbol{\theta}_b^*$ and its induced probability vector, $\boldsymbol{\theta}^*$, are such that: (a) $\hat{\varphi}_i^*(\boldsymbol{\theta}^*) = \hat{\varphi}_i(\boldsymbol{\theta}^*) \forall i$; and (b) any component, $\theta_\tau^* > 0$, of $\boldsymbol{\theta}_b^*$ corresponds to a graph G_τ such that $\varphi_i^\tau = \lfloor \hat{\varphi}_i(\boldsymbol{\theta}_b^*) \rfloor$ or $\varphi_i^\tau = \lceil \varphi_i(\boldsymbol{\theta}_b^*) \rceil$ for some $\varphi \in (0, m]$ and all i . If $v = 0$, then this implies that $\varphi_i^\tau = \lfloor \hat{\varphi}_i(\boldsymbol{\theta}^*) \rfloor$ or*

$\varphi_i^T = \lceil \hat{\varphi}(\boldsymbol{\theta}^*) \rceil$ for all i .

A network G in which a consumer segment i observes $\hat{\varphi}(\boldsymbol{\theta}^*) + k$ (where $k > 1$) sellers has a disproportionately negative effect on profits compared with the otherwise identical network G' in which i observes $\hat{\varphi}(\boldsymbol{\theta}^*) - k$ sellers. The reason for this is that in G each of the $\hat{\varphi}(\boldsymbol{\theta}^*) + k$ sellers competes with $\hat{\varphi}(\boldsymbol{\theta}^*) + k - 1$ other sellers. Hence, the sum of links generated by i 's observation in any network is convex in the number of sellers observed.

In the Appendix, we show that in the case where there is probability of graphs such as G and G' being generated, it is always possible to find a reallocation of probabilities such that: (a) consumer segment i (and every other consumer segment) observes the same number of sellers in expectation and (b) prices increase.

The platform owner maximises aggregate profit by ensuring that the number of sellers observed by each buyer is as tightly focused around some mean, which is determined by the innate demand for the sellers' goods and how substitutable those goods are, as possible. For any given v , such an outcome minimises the probability of there being states in which there are a large number of sellers competing and some sellers being relatively more prominent than others, both of which are disproportionately costly to the platform owner.

When $v = 0$, the platform owner is able to fully control which sellers consumer segments observe, and hence show consumers either $\lfloor \hat{\varphi}(\boldsymbol{\theta}^*) \rfloor$ or $\lceil \hat{\varphi}(\boldsymbol{\theta}^*) \rceil$ sellers. When $v > 0$, it is still optimal to intentionally show each consumer segment $\lfloor \hat{\varphi}(\boldsymbol{\theta}_b^*) \rfloor$ or $\lceil \hat{\varphi}(\boldsymbol{\theta}_b^*) \rceil$ sellers, with $\lceil \hat{\varphi}(\boldsymbol{\theta}_b^*) \rceil < \lceil \hat{\varphi}(\boldsymbol{\theta}^*) \rceil$, as when $v > 0$ the average number of sellers observed will be greater than the average number of sellers intentionally shown to consumer segments by the platform owner.

When $\hat{\varphi}(\boldsymbol{\theta}^*)$ is not an integer, either any optimal probability vector $\boldsymbol{\theta}^*$ is non-degenerate, to ensure that the correct number of sellers are observed in expectation or m and n are such that some consumer segments can be shown $\lfloor \hat{\varphi}(\boldsymbol{\theta}^*) \rfloor$ sellers and some $\lceil \hat{\varphi}(\boldsymbol{\theta}^*) \rceil$ sellers so that the correct average number of sellers can be displayed

using a single, deterministic graph. Either way, when $v = 0$, there are solutions to the platform's optimisation problem in which there are only $\lceil \hat{\varphi}(\boldsymbol{\theta}^*) \rceil$ sellers for whom $\beta_j(\boldsymbol{\theta}) > 0$, and are thus active in the network. If $v > 0$, this is not the case, as by definition there is some probability every seller will be observed by every consumer segment. This fact has implications for the optimal network structure, as we outline below.

Comparative statics and the optimality of randomisation

Suppose that $v > 0$. When this holds, it is possible to further characterise the optimal graph structure, even in the case where v is arbitrarily small. Let $w_{ij}(\boldsymbol{\theta}^*)$ denote the total probability that i observes j for a probability vector $\boldsymbol{\theta}^*$. The following result holds:

Proposition 5. *If $v > 0$, then any solution to the platform owner's maximisation problem, induces a $\boldsymbol{\theta}^*$ such that $\beta_j(\boldsymbol{\theta}^*) = \beta \in \mathbb{R}_+ \forall j$. Hence for any $v \in [0, 1)$, m and n there exists a solution to the platform owner's maximisation problem that induces a $\boldsymbol{\theta}^*$ in which $w_{ij}(\boldsymbol{\theta}^*) = w$ for all i, j pairs and each consumer i observes $\hat{\varphi}(\boldsymbol{\theta}^*)$ sellers in expectation.*

The first result in Proposition 5 implies that the total probability that each seller is observed across all consumer segments must be the same for all sellers. If $\beta_j(\boldsymbol{\theta}^*) > \beta_k(\boldsymbol{\theta}^*)$ for some seller pair (j, k) , then it must be that $\beta_j(\boldsymbol{\theta}_b^*) > \beta_k(\boldsymbol{\theta}_b^*)$. Given at the optimum it must be that $\varphi_i^\tau = \lfloor \hat{\varphi}(\boldsymbol{\theta}^*) \rfloor$ or $\varphi_i^\tau = \lceil \hat{\varphi}(\boldsymbol{\theta}^*) \rceil$ for all i , we show in the appendix that this inequality in turn implies $C_j(\boldsymbol{\theta}^*) < C_k(\boldsymbol{\theta}^*)$.

Intuitively, conditional on being active, sellers who are more likely to be observed in the baseline graph are also more likely to be active in a state of the world in which there are fewer sellers, i.e. states where few or none of the sellers who are not observed in a given baseline graph are observed by some other means by the consumer segments. This reduces the centrality of these sellers compared with those

sellers who are less likely to be observed in a baseline graph, who are more likely to be externally observed and thus competing with sellers shown to consumers directly by the platform owner.

Our analysis implies then that one solution to the platform owner’s problem induces a vector $\boldsymbol{\theta}^*$ that sellers observe completely at random, such that each consumer segment observing each seller with the same probability and the expected number of sellers observed overall is $\hat{\varphi}(\boldsymbol{\theta}^*)$. This is not a unique solution in all cases, as different sellers could be shown to each consumer segment with different probabilities depending on n and m as long as each $\beta_j(\boldsymbol{\theta}^*) = \beta$ overall. However, randomising across sellers and consumer segments in this way is optimal for for all values of $v \in [0, 1)$, m and n .

We can also characterise the effect of changing parameters in the model on the optimal average number of sellers observed, $\hat{\varphi}(\boldsymbol{\theta}^*)$. To do so, we write $\hat{\varphi}(\boldsymbol{\theta}^*) = \hat{\varphi}^*(\boldsymbol{\theta}^*(c, \gamma))$, and state the following result:

Proposition 6. *Suppose $c' > c$ and $\gamma' > \gamma$. Then $\hat{\varphi}^*(\boldsymbol{\theta}^*(c, \gamma)) \geq \hat{\varphi}^*(\boldsymbol{\theta}^*(c', \gamma))$ and $\hat{\varphi}^*(\boldsymbol{\theta}^*(c, \gamma')) \geq \hat{\varphi}^*(\boldsymbol{\theta}^*(c, \gamma))$, with the inequalities strict if $\hat{\varphi}^*(\boldsymbol{\theta}^*(c, \gamma)) \in (1, m)$.*

As substitutability increases, the platform has an incentive to reduce the average number of sellers observed by each consumer segment, as this reduce the level of expected competition in the network. As seller quality increases, so too does the demand of every consumer for a given price vector, \boldsymbol{p} , which increases the platform owner’s incentive to show consumers additional products, increasing competition, but also expected sales.

7 Heterogeneous consumer segments and sellers

Thus far, we have considered the case in which both consumer segments and sellers are homogeneous. We now consider the case where sellers differ in terms of product quality, and consumer segments differ in terms of how substitutable they see each

seller's product as being. In doing so, we highlight the flexibility of the network approach to modelling buyer-seller interactions in this setting.

Vertical differentiation

In the previous section, we considered the case in which goods are horizontally differentiated, but are of the same quality. We now consider the case where products may differ in quality, which in the model corresponds to the case where $\gamma_j > \gamma_k$ for at least one pair of sellers j and k . In this subsection, we assume that $c_i = c \forall i$ to focus on the effect of vertical differentiation.

When one product is of higher quality than another, the platform owner has an immediate incentive to always show buyers the highest quality sellers. However, the analysis in Section 6 suggests that doing so imposes a cost on the platform owner, as it means that the likelihood that high quality sellers will be in the relatively-competition states (i.e. where $\varphi_i^\tau = \lceil \hat{\varphi}(\boldsymbol{\theta}^*) \rceil$ more often) will be larger and, when $v > 0$, it increases the expected level of competition lower quality sellers face, reducing prices.

To understand the platform owner's trade-off in more detail, let $\hat{\gamma} := \gamma_h - \gamma_l$, where γ_h and γ_l represent the quality of the best and worst quality sellers respectively. The following result holds:

Proposition 7. *Suppose $v > 0$ and $\hat{\gamma} > 0$. (i) If $\gamma_j \geq \gamma_k$ then $\beta_j \geq \beta_k$ and (ii) $\exists \delta \in \mathbb{R}_+$ such that when $\hat{\gamma} < \delta$, for any solution to the platform owner's maximisation problem, $\boldsymbol{\theta}_b^*$, there is at least one graph, G_τ , where $\theta_\tau > 0$ is a component of $\boldsymbol{\theta}_b^*$ and $E_{ij} \in G_\tau$ for all j .*

The platform has an incentive to bias observations towards high quality products. However, doing so increases the probability that low quality sellers compete in high competition states, as it is more likely that they are active as a result of being indirectly observed and thus compete with those sellers shown to the buyers directly.

These sellers then set lower prices than they would if observation probabilities were equal, which has a negative effect on the prices set by the high quality sellers.

If the difference between the highest and lowest qualities is not too high, then the platform owner is willing to forgo some of the profit associated with only ensuring that the highest quality products are observed in order to reduce the level of competition in the network. To do so, they display sellers of lower quality to consumers with some positive probability.

We show in the Appendix that Proposition 4 holds when products are vertically differentiated. Proposition 7 then implies that for any optimal probability vector θ^* each consumer segment is still optimally presented with either $[\hat{\varphi}(\theta^*)]$ or $[\hat{\varphi}(\theta^*)]$ sellers, but is shown a mix of high quality and low quality sellers in expectation when $\hat{\gamma} < \delta$. If the platform were motivated to display fewer products as a result of, for example, consumers facing choice overload or search costs, then it would always display the highest quality products available, which differs from the outcome of optimisation in our model. The selection of sellers presented to a given consumer segment is more likely to contain the best quality seller than the worst quality seller, but at least some consumers are directed to even the lowest quality sellers by the platform with positive probability.

Increasing the probability that high quality products are observed increases the level of effective competition in the network, which affects the optimal pattern of observations. Suppose that each $\tilde{\gamma}_j \sim \Phi$, where Φ is a symmetric and bounded probability distribution, such that the realisation of $\tilde{\gamma}_j$'s value, $\gamma_j > 0$ and $E[\tilde{\gamma}_j] = \bar{\gamma}$ for all j . Suppose that the platform owner sets the vector θ after the realisation of γ .

Let $\tilde{\gamma}$ denote the random quality vector associated with the case where each $\tilde{\gamma}_j \sim \Phi_i$. Suppose that if $\tilde{\gamma}_j \sim \Phi_1$ it is bounded such that $\tilde{\gamma}_j \in [\gamma_l, \gamma_h]$. Now define Φ_2 such that when $\tilde{\gamma}_j \sim \Phi_2$, $\tilde{\gamma}_j$ can be decomposed such that $\tilde{\gamma}_j = \tilde{\gamma}'_j + \varepsilon_j$, where $\tilde{\gamma}'_j \sim \Phi_1$ and ε_j is distributed symmetrically with mean 0 and is bounded such that

$\varepsilon_j \in [\varepsilon_l, \varepsilon_h]$.

We examine the ex-ante (i.e prior to the realisation of $\tilde{\gamma}$) profits and expected number of sellers observed by each segment when qualities are have a distribution of Φ_i in the following theorem:

Theorem 3. *i) For any value of c and $v \in [0, 1)$, $E[\pi_P(\theta^*)|\tilde{\gamma} \sim \Phi_2] \geq E[\pi_P(\theta^*)|\tilde{\gamma} \sim \Phi_1]$ and ii) for any value of $v \in [0, 1)$, $\exists c_T \in \mathbb{R}_+$ such that if $c \leq c_T$, $E[\hat{\varphi}(\theta^*)|\tilde{\gamma} \sim \Phi_1] \geq E[\hat{\varphi}(\theta^*)|\tilde{\gamma} \sim \Phi_2]$ and if $c > c_T$, $E[\hat{\varphi}(\theta^*)|\tilde{\gamma} \sim \Phi_2] > E[\hat{\varphi}(\theta^*)|\tilde{\gamma} \sim \Phi_1]$.*

To illustrate the results in Theorem 3, we consider a more limited case where under Φ_1 each $\tilde{\gamma}_j = \bar{\gamma}$ with probability 1. Consider first the claim relating to profit. If c is sufficiently small (e.g. equal to zero), then in expectation the optimal probability vector for either distribution will be such that $\theta_c^* = 1$. In this case, expected profit is the same under both distributions.

However, in the case where the platform owner restricts the number of sellers consumers observe, they are able to bias consumer observation towards high-quality products. In the case where Φ_1 results in each seller having the same quality with probability 1, this is clearly not possible, whereas the mean-preserving spread Φ_2 generates some high-quality and low-quality players in expectation. Thus, when c is sufficiently high, $E[\pi_P(\theta^*)|\tilde{\gamma} \sim \Phi_2] > E[\pi_P(\theta^*)|\tilde{\gamma} \sim \Phi_1]$ due to consumers being more likely to observe high-quality sellers.

Now consider the second result in Theorem 3. As c increases, the number of sellers observed by consumers reduces for either distribution of qualities. However, the expected loss of a segment observing fewer sellers to platform profit is increasing more slowly in the case where is no vertical differentiation. The reason for this is that as c becomes large, the expected quality of a seller that the platform owner is marginally willing to exclude in the case where quality is dispersed becomes greater than the mean quality level, $\bar{\gamma}$.

The platform owner is less willing to exclude such high-quality sellers from being observed. Hence, when c is sufficiently large, the optimal number of sellers a segment

observes is, in expectation, greater for the vertically differentiated case compared to the case where product quality is the same. This is shown in Figure 5.




Figure 5: Dark blue line denotes the case where $\tilde{\gamma} \sim \Phi_2$, light blue where $\tilde{\gamma} \sim \Phi_1$.

Heterogeneous consumers

Now we consider the case where each consumer segment's substitutability parameter, c_i , can differ from one another, but assume products are not vertically differentiated. A natural question in this setting is whether it would be more profitable for the platform to have some seller(s) serve consumer segments who see the products as highly substitutable, while others are shown to those who see the products as being less substitutable.

Our model suggests that, in fact, that the distribution of networks where each consumer segment observes every seller with equal probability is still optimal in this case, as Proposition 8 shows:

Proposition 8. *Suppose $c_i > c_k$ for at least one i, k consumer segment pair. Then, for any optimal vector θ^* , $\hat{\varphi}_i(\theta^*) \leq \hat{\varphi}_k(\theta^*)$. Furthermore, there exists a solution to the platform owner's maximisation problem, θ_b^* , that induces a θ^* such that $w_{ij}(\theta^*) = w_i \forall j$.*

The first result in Proposition 8 reflects the fact that when products are more substitutable for i than k and i observes the same number of sellers as k , prices and profits from i will be lower than from k . The platform has an incentive to reduce the number of products observed by i , and does so in equilibrium, assuming that it is not optimal to always show both consumer segments every seller.

The second result indicates that, even in the case where consumers are heterogeneous with respect to how substitutable they see the goods as being, it is still optimal to randomise across all sellers equally for each consumer segment. Competing for a consumer segment composed of consumers who view products as highly substitutable results in there being an incentive for sellers to reduce their prices. To increase prices, then, the platform has an incentive to ensure that there is some probability that sellers are observed by those consumer segments with lower levels of substitutability.

By ensuring that every seller is as central as every other seller for any optimal probability vector θ^* , the platform maximises profits; sellers set the same price and are induced in doing by there being some probability of them being observed by segments with high and low values of the substitutability parameter, c_i .

One way to implement an outcome in which every seller is as central as another in this case is by choosing a baseline probability vector that induces a θ^* such that each consumer segment observes every seller with equal probability. Of course, by the first result in the theorem, that probability will differ between consumer segments with different levels of c_i , and if $c_i > c_k$, then it will be lower for i than k .

8 Conclusion

We analyse the case where consumers observe a subset of sellers on a platform, which can be thought of as a bipartite observation network. The probability that an observation network is realised is determined either by nature or by the owner of the platform. Prices are set prior to the realisation of the network, but the realisation probabilities are common knowledge.

Using the characterisation of equilibrium, we observe how changes in network

structure affect prices. By examining a case in which observation probabilities are randomly determined, we show that prominent sellers impose a disproportionate cost on the platform owner's profits, due to the feedback effect inherent to equilibrium price setting behaviour. Prominent sellers are more likely in the case where observation probabilities are more dispersed, and such sellers disproportionately increase competition, decreasing prices.

Where there is no vertical differentiation, the optimal graph structure is one in which the expected number of sellers observed by consumers is as close to the average number of sellers observed across all possible networks. This minimises competition for a given number of expected buyer-seller links, increasing profits. When there is even a small probability that consumers observe sellers not explicitly presented to them by the platform, then it is always strictly better for the platform to randomise such that buyers observe each seller with equal probability, rather than choose any deterministic pattern of buyer-seller interaction.

When products are vertically differentiated, the platform owner has an incentive to increase the probability that sellers of higher quality are observed. This increases the effective competition faced by other sellers in the network, which reduces prices. To reduce the significance of this effect, the platform owner reduces the total number of sellers observed by consumers and, if products are not too vertically differentiated randomise so that each seller is deliberately presented to buyers some of the time.

As platforms have an incentive to reduce competition in order to increase prices, our analysis suggests that competition authorities would be well-advised to take seriously attempts by platforms to control intra-platform competition. Regulation, insofar as it has been directed at online platforms, has tended to focus on competition between platforms. As particular online platforms become increasingly dominant, this kind of competition becomes less relevant, and the incentives to increase prices by tweaking search algorithms or the use of private information will become increasingly important.

More broadly, the framework here could be used to examine the effect of entry, exit and mergers on networked markets. For example, it is possible to use our characterisation of the price equilibrium to identify firms who impose the most competitive pressure on the network. Identifying such firms has clear implications for merger control and which firms to support during economic crises.

Appendix

The demand function

We show that the linear demand curve for the game $\Gamma(\boldsymbol{\theta})$ is a form of the one generated by the following demand system. Let y_i denote i 's demand for a numeraire good. Suppose that i has the following quasi-linear, quadratic utility function:

$$u_i(\mathbf{x}_i) = \sum_{j=1}^m \gamma_j x_{ij} - \sum_{j=1}^m \kappa x_{ij}^2 - \sum_{j=1}^m x_{ij} \left(\sum_{k=1}^m \rho x_{ik} \right) + y_i$$

where $\kappa, \rho \in \mathbb{R}_+$. Suppose that each buyer has an income of ϖ . Assuming ϖ is sufficiently large and (A1) holds, the demand for each product is positive. Define the $m \times m$ matrix with 1 on its diagonals and ρ on its off-diagonals, $\boldsymbol{\kappa}$. Then, as discussed in Singh and Vives (1984) and Amir, Erikson and Jin (2015), the demand vector \mathbf{x}_i can be written $\mathbf{x}_i = \boldsymbol{\kappa}^{-1}(\boldsymbol{\gamma} - \mathbf{p})$.

Hence, for any consumer segment i and any seller j , the intercept term of the i 's demand for j 's product is some constant, a , multiplied by γ and their own price sensitivity term is also equal to a .

Proof of Proposition 1

Define $\alpha_j(\boldsymbol{\theta})$ as follows:

$$\alpha_j(\boldsymbol{\theta}) := \frac{\sum_i \sum_\tau \theta_\tau [\mu_{ij}^\tau (a\gamma_j - \sum_i \sum_k c_i \mu_{ik}^\tau \gamma_k)]}{2a\beta_j(\boldsymbol{\theta})}.$$

Note that assumption (A1) guarantees that: (a) $\alpha_j(\boldsymbol{\theta}) > 0 \forall j$ where $\beta_j(\boldsymbol{\theta}) > 0$ and (b) $[\mathbf{I} - \lambda R_S(\boldsymbol{\theta})]$ is positive definite for all $\boldsymbol{\theta}$. To see (a), as $c\gamma_k > 0$, then for a given vector $\boldsymbol{\gamma}$, and given values for the parameters c and a , $\alpha_j(\boldsymbol{\theta})$ is lowest when $\boldsymbol{\theta} = \boldsymbol{\theta}_c$, where $\boldsymbol{\theta}_c$ denotes the probability vector where $\theta_c = 1$, where θ_c is the probability that the complete bipartite graph G_c is realised. It is also clear that when $\boldsymbol{\theta} = \boldsymbol{\theta}_c$, $\alpha_l(\boldsymbol{\theta}_c) < \alpha_j(\boldsymbol{\theta}) \forall i \neq l$. When (A1) holds, $\alpha_l(\boldsymbol{\theta}_c) > 0$, which implies $\alpha_l(\boldsymbol{\theta}) > 0$ for all $\boldsymbol{\theta}$.

Now note that when (A1) holds, it must be the case that $na > (m-1) \sum_i c_i$, as $\frac{\sum_{j \neq i} \gamma_j}{\gamma_i} \geq m-1$. This immediately implies that the $\mathbf{I} - \lambda R_S(\boldsymbol{\theta}_c)$ is positive definite, as it is strictly diagonally dominant. Let $\varsigma_{ij}(\boldsymbol{\theta})$ denote the ij th component of $\mathbf{I} - \lambda R_S(\boldsymbol{\theta})$. The following trivially holds:

$$\left| \sum_{j \neq i} \varsigma_{ij}(\boldsymbol{\theta}_c) \right| \geq \left| \sum_{j \neq i} \varsigma_{ij}(\boldsymbol{\theta}) \right| \quad \forall i, \boldsymbol{\theta} \neq \boldsymbol{\theta}_c.$$

Hence, if the matrix $\mathbf{I} - \lambda R_S(\boldsymbol{\theta}_c)$ is diagonally dominant, then for any $\boldsymbol{\theta}$, the matrix $\mathbf{I} - \lambda R_S(\boldsymbol{\theta})$ is also diagonally dominant. The first result that $\alpha_j(\boldsymbol{\theta}) > 0 \forall j$ for all $\boldsymbol{\theta}$ guarantees that there exists a price, p'_l , such that for any vector of prices other than p'_l 's, $x_{il}(p'_l, \mathbf{p}_{-l}) > 0$. This holds for all i , and hence at any optimal solution it must be the case that: (a) $p_j^*(\boldsymbol{\theta}) > 0 \forall j$ and (b) $x_{il}(\mathbf{p}^*) > 0$.

Define: $E[\tilde{\pi}_j(\mathbf{p}; \boldsymbol{\theta})] = \frac{1}{2a\beta_j(\boldsymbol{\theta})} E[\pi_j(\mathbf{p}; \boldsymbol{\theta})]$. The maximisation problem $\max_{\mathbf{p}_j} E[\tilde{\pi}_j(\mathbf{p}; \boldsymbol{\theta})]$ has the same set of first-order conditions as the one that involves maximising j 's original profit function. It can be readily shown that the first-order condition (and therefore the resulting optimisation problem) for the payoff vector associated with the payoff described in (1) is equivalent to the first-order condition of the payoff vector associated with the original payoff function. The first-order condition of the payoff vector with individual components described in (1) is as follows $\boldsymbol{\alpha} = [\mathbf{I} - \lambda R_S(\boldsymbol{\theta})] \mathbf{p}(\boldsymbol{\theta})$, where $\boldsymbol{\alpha}$ is a $m \times 1$ vector whose j th component is $\frac{1}{2a\beta_j(\boldsymbol{\theta})} \alpha_j(\boldsymbol{\theta})$

As the matrix $\mathbf{I} - \lambda R_S(\boldsymbol{\theta})$ is positive definite, it is non-singular and the above first-order condition has a solution, which is denoted $\mathbf{p}^*(\boldsymbol{\theta})$. Rearranging this first-order condition leads to the expression: $\mathbf{p}^*(\boldsymbol{\theta}) = \boldsymbol{\alpha} [\mathbf{I} - \lambda R_S(\boldsymbol{\theta})]^{-1}$.

The first-order condition above yields a unique, interior solution. As shown above, (A1) guarantees that $p_j = 0$ cannot be a solution for any seller j 's maximisation problem, as there exists a p_j such that $x_{ij}(\mathbf{p}; \boldsymbol{\theta}) > 0$. Hence, there exists a $\varepsilon > 0$ such that $p_j = \varepsilon$ generates a strictly positive level of demand $x_{ij} > 0$, yielding a strictly positive profit.

Noting that $\frac{\sum_i \sum_\tau \theta_\tau [\mu_{ij}^\tau(a\gamma_j)]}{2a\beta_j(\boldsymbol{\theta})} = \frac{1}{2}\gamma_j$ in this setting and define $\boldsymbol{\gamma}$ as a $m \times 1$ vector whose j th element is equal to γ_j . It follows that $\mathbf{p}^*(\boldsymbol{\theta}) = \frac{1}{2} \sum_{k=0}^{\infty} R_S^k(\boldsymbol{\theta})\boldsymbol{\gamma} - \sum_{k=1}^{\infty} R_S^k(\boldsymbol{\theta})\boldsymbol{\gamma}$, which directly implies the result.

Proof of Theorem 1

First, note that as centrality is independent of $\boldsymbol{\gamma}$ the expected price of each seller can be written as $E[p(\boldsymbol{\theta}, \boldsymbol{\gamma})] = \boldsymbol{\gamma} - \frac{1}{2}E[C(\boldsymbol{\theta})]\boldsymbol{\gamma}$. Recall that $\tilde{w}'_{ij} = \tilde{w}_{ij} + \epsilon_{ij}$, where ϵ_{ij} is a symmetric random variable and $C(\boldsymbol{\theta})\mathbf{1} = \sum_{k=0}^{\infty} R_S^k(\boldsymbol{\theta})\mathbf{1}$.

This implies that $E[C_j(\boldsymbol{\theta})] = 1 + \sum_{\xi=1}^{\infty} P_{j\xi}(\boldsymbol{\theta})$, where $P_{j\xi}(\boldsymbol{\theta})$ is the expected sum of the weighted paths of length ξ that begin at j given the vector $\boldsymbol{\theta}$. As an example, a path of length 2 that goes from j to k and back to j is equal to $(\frac{\sum_i c_i}{2na} \frac{nw_j w_k}{nw_j}) (\frac{\sum_i c_i}{2na} \frac{nw_j w_k}{nw_k}) = (\frac{\sum_i c_i}{2na})^2 w_j w_k$ in this setting. The weighted sum of all paths of length 2 starting at j can hence be written $(\frac{\sum_i c_i}{2na})^2 \sum_{k \neq j} \sum_{s \neq k} w_k w_s$, and $(\frac{\sum_i c_i}{2na})^3 \sum_{k \neq j} \sum_{s \neq k} \sum_{t \neq s} w_k w_s w_t$ gives the same sum for paths of length 3, and so on.

Consider paths of length 1. Given that each $E[\epsilon_k] = 0$, $P_{j1}(\tilde{\boldsymbol{\theta}}_1) = P_{j1}(\tilde{\boldsymbol{\theta}}_2)$. Similarly, as $E[\epsilon_j \epsilon_k] = 0$, it must be that $P_{j2}(\tilde{\boldsymbol{\theta}}_1) = P_{j2}(\tilde{\boldsymbol{\theta}}_2)$. However, as $E[\epsilon_k^l] > 0$ for all even values of $l > 0$ and all k , it must be that, $E[(\tilde{w}'_k)^2 w_j] = E[(\tilde{w}'_k + \epsilon_k)^2 w_j] > E[(\tilde{w}_k)^2 w_j]$ for all $k \neq j$. As $E[(\tilde{w}'_k)^2 w_j]$ for $k \neq j$ is a component of $P_{j3}(\tilde{\boldsymbol{\theta}}_2)$, it follows that $P_{j3}(\tilde{\boldsymbol{\theta}}_2) > P_{j3}(\tilde{\boldsymbol{\theta}}_1)$. A similar argument holds for all $P_{j\xi}(\boldsymbol{\theta})$ where $\xi > 3$: $P_{j\xi}(\boldsymbol{\theta})$ is a function of at least $E[\epsilon_k^{\xi-1}]$ and $E[\epsilon_k^{\xi-2}]$ for all $k \neq j$ and, for any k , one of these is by definition is greater than zero, and hence, $P_{j\xi}(\tilde{\boldsymbol{\theta}}_2) > P_{j\xi}(\tilde{\boldsymbol{\theta}}_1)$ for $\xi \geq 3$.

It follows from the above that $E[C_j(\tilde{\boldsymbol{\theta}}_2)] > E[C_j(\tilde{\boldsymbol{\theta}}_1)] \quad \forall j$ and thus $E[p_j(\tilde{\boldsymbol{\theta}}_1)] > E[p_j(\tilde{\boldsymbol{\theta}}_2)]$. Now consider the ex-ante profit function of a seller j :

$$\mathbb{E}[\pi_j(\tilde{\boldsymbol{\theta}}_2)] = \mathbb{E}[p_j(\tilde{\boldsymbol{\theta}}_2)\alpha_j] - b\mathbb{E}[\beta_j p_j^2(\tilde{\boldsymbol{\theta}}_2)] + \mathbb{E}\left[\sum_{j \neq k} \hat{c}_{jk} p_j(\tilde{\boldsymbol{\theta}}_2) p_k(\tilde{\boldsymbol{\theta}}_2)\right].$$

Just for the sake of argument, we first assume that the parameters α_j , β_j , and each \hat{c}_{jk} are independent of the price vector \mathbf{p} . As $\mathbb{E}[\tilde{w}'_j] = \mathbb{E}[\tilde{w}_{ij}]$ and each element of set of observation probabilities is independent of every other element of that set, it follows that the expected profit generated by observation probabilities with distribution Λ_2 would be lower than Λ_1 . The reason for this is that: (a) $\mathbb{E}[p_j(\tilde{\boldsymbol{\theta}}_2)] < \mathbb{E}[p_j(\tilde{\boldsymbol{\theta}}_1)]$ and (b) (A1) implies that profit is concave in p_j . Hence, even if $\mathbb{E}[p_j(\tilde{\boldsymbol{\theta}}_2)] = \mathbb{E}[p_j(\tilde{\boldsymbol{\theta}}_1)]$, the following expression:

$$\mathbb{E}[p_j(\tilde{\boldsymbol{\theta}}_2)]\mathbb{E}[\alpha_j] - b\mathbb{E}[\beta_j]\mathbb{E}[p_j^2(\tilde{\boldsymbol{\theta}}_2)] + \mathbb{E}[C(\tilde{\boldsymbol{\theta}}_2)]$$

is less than the equivalent expression for $\tilde{\boldsymbol{\theta}}_1$, because $\text{var}(\epsilon_{ij}) > 0$.

However, the parameters α_j , β_j and each \hat{c}_{jk} are not independent of the realisation of each random observation probability as they are a function of \tilde{w}_{ij} . Expected demand in this environment can be written:

$$\mathbb{E}[\tilde{x}_{ij}(\mathbf{p}^*)] = \tilde{w}_j(a\gamma - ap_j + c_i \sum_{j \neq k} w_k(p_k - \gamma)).$$

Hence, it follows that $\text{cov}(\tilde{x}_{ij}(\mathbf{p}^*), \mathbf{p}^*) < 0$, which holds both because (A1) implies that demand conditional on i observing j is falling in price and because $\text{cov}(\tilde{w}_j'', \mathbf{p}^*) < 0$. Furthermore, $|\text{cov}(\tilde{w}_j'', \mathbf{p}^*)| > |\text{cov}(\tilde{w}_j', \mathbf{p}^*)|$. Hence, $\mathbb{E}[\pi_j(\tilde{\boldsymbol{\theta}}_2)] > \mathbb{E}[\pi_j(\tilde{\boldsymbol{\theta}}_1)]$.

Proof of Proposition 2

Given that $C(\boldsymbol{\theta})\mathbf{1} = \sum_{k=0}^{\infty} R_S^k(\boldsymbol{\theta})\mathbf{1}$ and the expression for the equilibrium price vector, it is clear that the complete network maximises the centrality of each node in G , which then minimises the price vector \mathbf{p} for a given m . At the same time, as (A1) holds, an edge E_{ij} between consumer segment i and seller j increases total demand, holding prices constant. Hence, a complete network, which implies that, holding price

constant, i 's expected consumer surplus is maximised where $\theta_c = 1$. It follows from these two facts that consumer surplus is maximised when $\boldsymbol{\theta} = \boldsymbol{\theta}_c$.

Proof of Proposition 3

Consider the case where $\theta_c = 1$ and let G denote a graph which is defined as $G_c - E_{ij} = G$ for some buyer i and the seller for whom γ_j is the smallest component in the vector $\boldsymbol{\gamma}$. Let $\theta_1 = \theta_G - \theta_c$ and $p_k^*(\boldsymbol{\theta}_c) = p_k^*$. By the envelope theorem (Milgrom and Segal, 2002):

$$\frac{\partial \pi_P}{\partial \theta_1} \Big|_{\theta_c=1} = -ap_j^*(\gamma - p_j^* + 2c_i \sum_{k \neq j} (p_k^* - \gamma)) + \sum_l \sum_{k \neq l} \hat{c}_{kl} p_l^* \frac{\partial p_k^*}{\partial \theta_1}.$$

As $(p_k^* - \gamma) < 0$, it follows that $\frac{\partial^2 \pi_P}{\partial \theta_1 \partial c} \Big|_{\theta_c=1} > 0$. Hence, there exists a \bar{c} such that if $c_i > \bar{c}$ for all i , then $\frac{\partial \pi_P}{\partial \theta_1} \Big|_{\theta_c=1} > 0$ and thus $\boldsymbol{\theta}_c$ is suboptimal.

Proof of Proposition 4

Let $C_j(\boldsymbol{\theta})$ be (jointly one of) the smallest component(s) of the vector $\mathbf{C}(\boldsymbol{\theta})\mathbf{1}_m$. We first examine the effect of the change from $\boldsymbol{\theta}$ to $\hat{\boldsymbol{\theta}}_{jk}$, which we define as $\hat{\boldsymbol{\theta}}_{jk} := (1 - \varepsilon)\boldsymbol{\theta} + \varepsilon\boldsymbol{\theta}_{jk}$, on the sum of the profits of j and k , holding p_i $i \neq j, k$ fixed. Note that in this case, the sum of these profits can be written:

$$\mathbb{E}[\pi_j(\boldsymbol{\theta}) + \pi_k(\boldsymbol{\theta})] = \sum_{j,k} [p_i(\boldsymbol{\theta})(\alpha_i - a\beta_i p_i(\boldsymbol{\theta}) + \sum_{l \neq i} \hat{c}_{il} p_l(\boldsymbol{\theta}))].$$

As a result of the fact that profits are increasing and concave in prices below the monopoly price (which is implied by (A1)), it follows that the proposed switch will result in an increase in the total profits the platform receives from j and k . The same logic applies for any seller $C_l(\boldsymbol{\theta}) > C_j(\boldsymbol{\theta})$ and if $C_l(\boldsymbol{\theta}) = C_j(\boldsymbol{\theta})$, then the proposed reallocation has no direct effect on prices. It follows that:

$$p_j(\boldsymbol{\theta}) + p_k(\boldsymbol{\theta}) < p_j(\hat{\boldsymbol{\theta}}_{jk}) + p_k(\hat{\boldsymbol{\theta}}_{jk}).$$

Hence:

$$\hat{c}_{jk}p_k(\boldsymbol{\theta}) + \hat{c}_{kj}p_j(\boldsymbol{\theta}) < \hat{c}_{jk}p_k(\hat{\boldsymbol{\theta}}_{jk}) + \hat{c}_{kj}p_j(\hat{\boldsymbol{\theta}}_{jk}).$$

Furthermore, given that $p_j(\boldsymbol{\theta}) < p_k(\boldsymbol{\theta})$, and that $\sum_{j,k} p_i(\boldsymbol{\theta})(\alpha_i - a\beta_i p_i(\boldsymbol{\theta}))$ is increasing and concave in $p_i \in [0, \frac{1}{2}\gamma]$, it follows that the effect of the change from $\boldsymbol{\theta}$ to $\hat{\boldsymbol{\theta}}_{jk}$ on the sum of the profits of j and k , holding p_i $i \neq j, k$ fixed is an increase in platform profits. This then implies that the direct effect of a change in probability vector from change from $\boldsymbol{\theta}$ to $\hat{\boldsymbol{\theta}}_j$ (i.e. where the proposed set of neighbourhood switches takes place between j and every other seller in the network).

We now turn to the second order effects of the proposed reallocation when $m > 2$. By second-order effects, we refer to the effect of the proposed set of neighbourhood switches between j and every $k \neq i$ has on i 's profits. Formally, we compare $\sum_i \pi_i((1 - (m - 2)\varepsilon)\boldsymbol{\theta} + \sum_{k \neq j} \varepsilon\boldsymbol{\theta}_{jk} - \varepsilon\boldsymbol{\theta}_{ji})$ with the sum of profits generated by $\boldsymbol{\theta}$, $\sum_i \pi_i(\boldsymbol{\theta})$. Define:

$$\Delta C_i := C_i((1 - (m - 2)\varepsilon)\boldsymbol{\theta} + \sum_{k \neq j} \varepsilon\boldsymbol{\theta}_{jk} - \varepsilon\boldsymbol{\theta}_{ji}) - C_i(\boldsymbol{\theta})$$

As per Bonacich (1972), $C_j(\boldsymbol{\theta}) = 1 + \sum_{k \neq j} \hat{c}_{jk}C_k(\boldsymbol{\theta})$, from which two observations follow. First, as the direct effect of each neighbourhood switch between j and k increases j 's centrality less than it decreases k 's, it follows that $\sum_i \Delta C_i(\boldsymbol{\theta}) > 0$. Furthermore, it must also be the case that if $C_i(\boldsymbol{\theta}) \geq C_l(\boldsymbol{\theta})$ then $|\Delta C_i(\boldsymbol{\theta})| \geq |\Delta C_l(\boldsymbol{\theta})|$, $i, l \neq j$.

The two above facts imply that the sum of second-order prices changes is positive and that the prices of more central sellers increase more than the prices of less central players. Again, as $\pi_P(\boldsymbol{\theta})$ is concave and increasing in $p_j \in [0, \frac{1}{2}\gamma]$ for all j , it follows that the sum of the second-order effects of a switch from $\boldsymbol{\theta}$ to $\hat{\boldsymbol{\theta}}_j$ are profit increasing.

The above analysis then jointly implies that $\pi_P(\hat{\boldsymbol{\theta}}_j) > \pi_P(\boldsymbol{\theta})$. This implies the result: for any vector in which there exists a pair of sellers j and k such that $C_j(\boldsymbol{\theta}) > C_k(\boldsymbol{\theta})$, there is always a series of neighbourhood switches that increases profits.

Hence, any solution to the platform owner's maximisation problem must be such that $C_j(\boldsymbol{\theta}) = C_k(\boldsymbol{\theta})$ for all j, k pairs

Proof of Theorem 2

Noting that $C_j(\boldsymbol{\theta}) = 1 + \sum_{k \neq j} \hat{c}_{jk} C_k(\boldsymbol{\theta})$, Theorem 2 implies that any solution, $\boldsymbol{\theta}_b^*$, to the platform owner's maximisation problem, induces a $\boldsymbol{\theta}^*$ and $G_S(\boldsymbol{\theta}^*)$ such that $\sum_{k \neq j} \hat{c}_{jk}(\boldsymbol{\theta}^*) = \hat{c}(\boldsymbol{\theta}^*) \in \mathbb{R}_+$. Define σ_{jy}^τ as the number of buyers for which seller j faces competition from exactly $y \in \{0, 1, \dots, m-1\}$ sellers in the graph G_τ , with $y = 0$ if G_τ is such that $E_{ij} = 0$ for all j . Then we can write $\sum_{k \neq j} \hat{c}_{jk}(\boldsymbol{\theta})$ as follows:

$$\sum_{k \neq j} \hat{c}_{jk}(\boldsymbol{\theta}) = \frac{c \sum_y \sum_\tau y \theta_\tau \sigma_{jy}^\tau}{a \beta_j(\boldsymbol{\theta})}, \quad (1)$$

which implies that $\frac{a}{c} \hat{c}(\boldsymbol{\theta}^*) + 1 = \hat{\varphi}_i^*(\boldsymbol{\theta}^*) = \hat{\varphi}(\boldsymbol{\theta}^*)$ for all i when $c_i = c$. For now, assume $v = 0$ and hence $\boldsymbol{\theta} = \boldsymbol{\theta}_b$. Consider a proposed profit-maximising vector $\boldsymbol{\theta}$ in which (a) $C_j(\boldsymbol{\theta}) = C_k(\boldsymbol{\theta})$ for all j, k pairs and (b) it is true for at least one segment i that $\varphi_i^\tau \leq [\hat{\varphi}^*]$ or $\varphi_i^\tau \geq [\hat{\varphi}^*]$, with at least one of the inequalities strict, in at least one graph τ realised with probability $\theta_\tau > 0$. We rule out that $\varphi_i^\tau = 0$ for any consumer segment i in any graph τ realised with probability $\theta_\tau > 0$, as this clearly suboptimal.

As each consumer segment is identical in such a graph it is possible to construct a probability vector, $\bar{\boldsymbol{\theta}}$, which yields the same profits as $\boldsymbol{\theta}$, but makes it easier to compare profits across graphs.

Take a graph τ generated with probability $\theta_\tau > 0$. As before, let τ_{ij} denote the graph generated by a neighbourhood switch between two consumer segments, i and j being performed on the graph τ . Let X denote total number of possible switches between $i, k \in B$.

Let $\bar{\boldsymbol{\theta}}$ denote the following probability vector. Suppose $\theta_\tau > 0$. Then the probability that τ is realised in $\bar{\boldsymbol{\theta}}$ is $\frac{\theta_\tau}{X+1}$, which is also equal to the realisation probability of each τ_{ij} for $i, j \in B$ where $i \neq j$. As $C_j(\boldsymbol{\theta}) = C_k(\boldsymbol{\theta})$ and consumers have identical preferences, $\pi_P(\bar{\boldsymbol{\theta}}) = \pi_P(\boldsymbol{\theta})$. The transformation makes it possible to show that $\boldsymbol{\theta}$ is

not a solution to the platform owner's profit maximisation problem.

We consider first the case where in θ , there exists a graph τ where $\varphi_i^\tau < \lfloor \hat{\varphi}(\theta^*) \rfloor$ and a graph τ' in which $\varphi_j^{\tau'} > \lceil \hat{\varphi}(\theta^*) \rceil$, where i and j may not be the same segment, and then consider afterwards the case where one of these inequalities is not strict.

Under this assumption, when the probability vector is $\bar{\theta}$, there is a strictly positive probability that a graph τ_H will be realised, where $\varphi_i^{\tau_H} > \lceil \hat{\varphi}(\theta^*) \rceil$ and i observes a set of sellers S_H . There is also a strictly positive probability that a graph, $\tau_{H,L}$, is realised, where $\tau_{H,L}$ is "paired" with τ_H in the sense that i observes a set of sellers $S_{H,L} \subset S_H$ and $\varphi_i^{\tau_{H,L}} < \lfloor \hat{\varphi}(\theta^*) \rfloor$.

Suppose τ'_H is a graph identical to τ_H except that i observes a set of sellers $S'_H \subset S_H$, such that $\varphi_i^{\tau'_H} = \lceil \hat{\varphi}(\theta^*) \rceil$. This is ensured by deleting the edge between i and j , E_{ij} , for at least one seller j where $E_{ij} \in \tau_H$. Let τ''_H be a graph identical to τ'_H except that i observes a set of sellers $S''_H \subseteq S'_H \subset S_H$, such that $\varphi_i(\tau''_H) = \lfloor \hat{\varphi}(\theta^*) \rfloor$, by deleting at most one edge between i and k where $E_{ik} \in \tau'_H$. Note that if $\bar{\varphi}$ is an integer, then τ'_H and τ''_H are identical, otherwise $S''_H \subset S'_H \subset S_H$.

Similarly, define $\tau'_{H,L}$ as a graph identical to $\tau_{H,L}$ except that i observes the set of sellers S'_H , such that $\varphi_i(\tau'_{H,L}) = \lceil \hat{\varphi}(\theta^*) \rceil$. This is ensured by adding an edge between i and j , E_{ij} , for at least one seller j where $E_{ij} \notin \tau$. $\tau''_{H,L}$ is constructed in an analogous way to τ''_H , and hence i observes a set of sellers $S''_H \subseteq S'_H \subset S_H$.

Let the constant $\eta > 0$ be such that it solves the expression $\eta \lceil \hat{\varphi}(\theta^*) \rceil + (1 - \eta) \lfloor \hat{\varphi}(\theta^*) \rfloor = \hat{\varphi}(\theta^*)$. We define the probability vector $\bar{\theta}'$ in the following way. $\bar{\theta}'_\tau = \bar{\theta}_\tau$ for all graphs except the probability that τ_H and its pair $\tau_{H,L}$, $\bar{\theta}'_H$ and $\bar{\theta}'_L$, are realised is zero. Instead, $\eta \bar{\theta}'_{\tau'_H} + (1 - \eta) \bar{\theta}'_{\tau''_H} = \bar{\theta}_{\tau_H}$ and $\eta \bar{\theta}'_{\tau'_L} + (1 - \eta) \bar{\theta}'_{\tau''_L} = \bar{\theta}_{\tau_L}$.

Each segment observes, in expectation, the same number of sellers in both $\bar{\theta}'$ and $\bar{\theta}$. If there is a difference in profit between the two, it is driven by differences in prices. Suppose that $j \in S''_H$. If it is also the case that $k \in S''_H$ then $\hat{c}_{jk}(\bar{\theta}') = \hat{c}_{jk}(\bar{\theta})$. However, by construction, $S''_H \subseteq S'_H \subset S_H$ and there exists a $l \in S_H$ but $l \notin S''_H$. It follows that $\sum_l \hat{c}_{jl}(\bar{\theta}') < \sum_l \hat{c}_{jl}(\bar{\theta})$ for at least one j, l pair. It follows that the

centrality of j and l are lower in $\bar{\theta}'$ than in $\bar{\theta}$. This implies that prices are higher across the network in $\bar{\theta}'$ than in $\bar{\theta}$, which in turn implies that $\pi_P(\bar{\theta}') > \pi_P(\bar{\theta})$, so θ cannot be a solution to the platform owner's maximisation problem.

Now consider the case in which there is a positive probability that a graph τ' where $\varphi_i^{\tau'} > \lceil \hat{\varphi}(\theta^*) \rceil$ is realised when the probability vector is θ , but there is no graph with positive realisation probability where $\varphi_i^\tau < \lfloor \hat{\varphi}(\theta^*) \rfloor$. Given that $\hat{\varphi}(\theta^*)$ is the number of sellers observed in expectation, it must be the case that $\varphi_i^\tau = \lfloor \hat{\varphi}(\theta^*) \rfloor$ for at least one i and τ pair. It follows that if the graph τ' is paired with a graph in which i observes exactly $\lfloor \hat{\varphi}(\theta^*) \rfloor$ sellers in the way described above, the vector $\bar{\theta}'$ will still be more profitable for the platform owner than θ .

Suppose there exists a graph τ where $\varphi_i^\tau < \lfloor \hat{\varphi}(\theta^*) \rfloor$, but for no graph generated with positive probability by the vector θ is it the case that $\varphi_j^{\tau'} > \lceil \hat{\varphi}(\theta^*) \rceil$ for any j, τ' pair. For this to hold it must be the case that $\hat{\varphi}^*$ is not an integer and that $\varphi_j^{\tau'} = \lceil \hat{\varphi}(\theta^*) \rceil$ and $\theta_{\tau'} > 0$ for at least one j, τ' pair.

If $\varphi_i^{\tau_H} = \lceil \hat{\varphi}(\theta^*) \rceil$ and $\varphi_i^{\tau_{H,L}} < \lfloor \hat{\varphi}(\theta^*) \rfloor$, then the preceding analysis implies that $S_H'' \subset S_H' \subset S_H$. Hence, by the same logic as the case where both original inequalities were strict, it must be true that $\bar{\theta}'$ will still be more profitable for the platform owner than θ , even in the case where no segment observes more than $\lceil \hat{\varphi}(\theta^*) \rceil$ sellers.

The same proof applies for $v > 0$, where $\theta_b \neq \theta$. Suppose θ_b is such that at least one segment i that $\varphi_i^\tau \leq \lfloor \hat{\varphi} \rfloor$ or $\varphi_i^\tau \geq \lceil \hat{\varphi} \rceil$, with at least one of the inequalities strict. Define $\bar{\theta}'_b$ and $\bar{\theta}_b$ in an analogous way to $\bar{\theta}'$ and $\bar{\theta}$, and let θ' and θ be the probability vectors induced by the baseline probability vectors $\bar{\theta}'_b$ and $\bar{\theta}_b$ respectively. Noting that v is common across all consumer segments and sellers, it must be that $\sum_l \tilde{c}_{jl}(\theta') < \sum_l \tilde{c}_{jl}(\theta)$, and so $\pi_P(\bar{\theta}') > \pi_P(\bar{\theta})$, which implies the result.

Proof of Proposition 5

Suppose that $\beta_s(\theta^*) > \beta_t(\theta^*)$ for some s, t pair. By Theorem 2 it must be the case that for any baseline graph G_τ that is realised with positive probability when $\theta = \theta_b$, $\varphi_i^\tau = \lfloor \hat{\varphi}(\theta^*) \rfloor$ or $\varphi_i^\tau = \lceil \hat{\varphi}(\theta^*) \rceil$ for all i .

Recall from equation (1) above that $\sum_{k \neq j} \hat{c}_{jk}(\boldsymbol{\theta}) = \frac{c \sum_y \sum_\tau y \theta_\tau \sigma_{jy}^\tau}{2a\beta_j(\boldsymbol{\theta})}$. As $\beta_s(\boldsymbol{\theta}^*) > \beta_t(\boldsymbol{\theta}^*)$, it more likely that if t is active then it is unintentionally observed they are unintentionally observed than is the case with s , and hence it must be that $\frac{\sigma_{sy}^\tau}{\beta_s(\boldsymbol{\theta})} < \frac{\sigma_{ty}^\tau}{\beta_t(\boldsymbol{\theta})}$ for all $y > \lceil \hat{\varphi}(\boldsymbol{\theta}^*) \rceil - 1$, and $\frac{\sigma_{sy}^\tau}{\beta_s(\boldsymbol{\theta})} > \frac{\sigma_{ty}^\tau}{\beta_t(\boldsymbol{\theta})}$ for $y = \lceil \hat{\varphi}(\boldsymbol{\theta}^*) \rceil - 1$ and $y = \lfloor \hat{\varphi}(\boldsymbol{\theta}^*) \rfloor - 1$. Hence, $\sum_{k \neq s} \tilde{c}_{sk}(\boldsymbol{\theta}) < \sum_{k \neq t} \tilde{c}_{tk}(\boldsymbol{\theta})$ when $\beta_s(\boldsymbol{\theta}^*) > \beta_t(\boldsymbol{\theta}^*)$. It follows that $C_s(\boldsymbol{\theta}) < C_t(\boldsymbol{\theta})$, violating the condition for optimality in Proposition 4. Thus, $\beta_j(\boldsymbol{\theta}^*) = \beta \forall j$.

Proof of Proposition 6

Let G_τ denote a graph where θ_τ for the optimal vector $\boldsymbol{\theta}_b^* = \boldsymbol{\theta}_b^*(c, \gamma)$ where $\varphi_i^\tau = \lceil \hat{\varphi} \rceil \in (1, m)$ for all i and $E_{ij} \notin G_\tau$ for some buyer i and seller j . Now consider a graph G'_τ , defined as $G_\tau + E_{ij} = G'_\tau$, realised with probability θ'_τ . Let $\theta'' = \theta'_\tau - \theta_\tau$ and $p_j^*(\boldsymbol{\theta}^*) = p_j^*$:

$$\frac{\partial \pi_P(\boldsymbol{\theta}^*(c, \gamma))}{\partial \theta''} = -ap_j^*(\gamma - p_j^* + 2 \sum_{k \neq j} c(p_k^* - \gamma)) + \sum_k \sum_{s \neq k} 2a\beta_k(\boldsymbol{\theta}^*) \hat{c}_{ks} p_k^* \frac{\partial p_s^*}{\partial \theta''}.$$

For an optimum probability vector, $\boldsymbol{\theta}^*(c, \gamma)$, $\frac{\partial \pi_P(\boldsymbol{\theta}^*(c, \gamma))}{\partial \theta''} = 0$. Note that $(p_k^* - \gamma) < 0$ and $|\frac{\partial p_s^*}{\partial \theta''}|$ is increasing in c , it follows that $\frac{\partial \pi_P(\boldsymbol{\theta})}{\partial \theta''}$ is falling in c for all $\hat{\varphi} \in (1, m)$. Hence $\frac{\partial \pi_P(\boldsymbol{\theta}^*(c', \gamma))}{\partial \theta''} < 0$ and thus $\hat{\varphi}^*(\boldsymbol{\theta}^*(c, \gamma)) > \hat{\varphi}^*(\boldsymbol{\theta}^*(c', \gamma))$.

Noting that $\frac{\partial \pi_P(\boldsymbol{\theta})}{\partial \theta''}$ is increasing in γ by (A1), it follows that $\frac{\partial \pi_P(\boldsymbol{\theta}^*(c, \gamma'))}{\partial \theta''} > 0$ and hence $\hat{\varphi}^*(\boldsymbol{\theta}^*(c, \gamma')) > \hat{\varphi}^*(\boldsymbol{\theta}^*(c, \gamma))$. If $\hat{\varphi} = 1$, then we have a corner solution, and it is possible that $\hat{\varphi}(\boldsymbol{\theta}^*(c, \gamma')) = \hat{\varphi}(\boldsymbol{\theta}^*(c, \gamma))$ and $\hat{\varphi}(\boldsymbol{\theta}^*(c, \gamma)) = \hat{\varphi}(\boldsymbol{\theta}^*(c', \gamma))$. Similarly, if $\hat{\varphi} = m$, then it is also possible that those equalities hold. Hence, for $\hat{\varphi}(\boldsymbol{\theta}^*(c, \gamma)) \geq \hat{\varphi}(\boldsymbol{\theta}^*(c', \gamma))$ and $\hat{\varphi}(\boldsymbol{\theta}^*(c, \gamma')) \geq \hat{\varphi}(\boldsymbol{\theta}^*(c, \gamma))$.

Proof of Proposition 7

Suppose $\gamma_k = \gamma_H$ and $\gamma_j = \gamma_L$. Suppose the converse of the statement in the Proposition holds, so that there is no graph G realised with positive probability under a baseline probability vector $\boldsymbol{\theta}_b$ such that $E_{ij} \in G$. If $v > 0$ then $\beta_j(\boldsymbol{\theta}) > 0$ even under this assumption. Let $\boldsymbol{\theta}$ be the probability vector induced by $\boldsymbol{\theta}_b$.

First, note that the proof of Proposition 4 does not rely on the assumption that $\gamma_j = \gamma_k$ for all j, k pairs: it is still the case that any graph in which $\varphi_j^\tau < [\hat{\varphi}(\boldsymbol{\theta}^*)]$ or $\varphi_j^\tau > [\hat{\varphi}(\boldsymbol{\theta}^*)]$ is suboptimal. By the proof of Proposition 5, it follows that $C_k(\boldsymbol{\theta}) < C_j(\boldsymbol{\theta})$.

Consider the vector: $\hat{\boldsymbol{\theta}}' := (1 - \varepsilon)\boldsymbol{\theta} + \varepsilon\boldsymbol{\theta}_{jk}$. The proof of Proposition 4 shows that there is a priori incentive for the platform to switch from inducing the vector $\boldsymbol{\theta}$ to $\hat{\boldsymbol{\theta}}'$ by performing a j, k neighbourhood switch on each G_τ whose component is $\theta_\tau > 0$ in the vector $\boldsymbol{\theta}$.

However, in the case where $\hat{\gamma} > 0$, such a switch imposes a cost on the platform as $v < 1$ and thus $\beta_k(\hat{\boldsymbol{\theta}}') < \beta_k(\boldsymbol{\theta})$, $\beta_j(\hat{\boldsymbol{\theta}}') > \beta_j(\boldsymbol{\theta})$. Even if $p_j = p_k$ (which does not hold here), $E[x_{ij}(\mathbf{p})|\mu_{ij} = 1] < E[x_{ik}(\mathbf{p})|\mu_{ij} = 1]$ for all i : each consumer segment's innate demand is higher for k than j .

Define the function $D_{jk}(\boldsymbol{\theta}, \hat{\gamma})$ as follows:

$$D_{jk}(\boldsymbol{\theta}, \hat{\gamma}) = \sum_{\tau} \theta_{\tau} [\pi_k^{\tau}(\boldsymbol{\theta}') - \pi_j^{\tau}(\boldsymbol{\theta})]$$

Where $\pi_j^{\tau}(\boldsymbol{\theta})$ is the profit j receives in G_{τ} when the probability vector is $\boldsymbol{\theta}$. We define $B_{jk}(\boldsymbol{\theta}, \hat{\gamma})$:

$$B_{jk}(\boldsymbol{\theta}, \hat{\gamma}) = \sum_{\tau} \theta_{\tau} [\pi_j^{\tau}(\boldsymbol{\theta}') - \pi_k^{\tau}(\boldsymbol{\theta})] + \sum_{l \neq k, j} \Delta \pi_l$$

where $\Delta \pi_l = \pi_l(\boldsymbol{\theta}') - \pi_l(\boldsymbol{\theta})$. Clearly, $B_{jk}(\boldsymbol{\theta}, \hat{\gamma}) + D_{jk}(\boldsymbol{\theta}, \hat{\gamma}) = \sum_l \Delta \pi_l$. $D_{jk}(\boldsymbol{\theta}, \hat{\gamma})$ is continuous, decreasing in $\hat{\gamma}$ and negative for $\hat{\gamma} > 0$, while $B_{jk}(\boldsymbol{\theta}, \hat{\gamma})$ is continuous and increasing in $\hat{\gamma}$, which follows from the fact that it is a positive function of $\pi_j^{\tau}(\boldsymbol{\theta}')$ and because $\pi_l(\boldsymbol{\theta})$ is a function of $\mu_{ij}^{\tau} c \gamma_j$ and $\mu_{ik}^{\tau} c \gamma_k$.

When $\hat{\gamma} = 0$, $B_{jk}(\boldsymbol{\theta}, \hat{\gamma}) + D_{jk}(\boldsymbol{\theta}, \hat{\gamma}) > 0$ by Proposition 4. Hence, $\exists \delta \in \mathbb{R}_+$ such that for all $\boldsymbol{\gamma}$ where $\hat{\gamma} < \delta$ such that $B_{jk}(\boldsymbol{\theta}, \hat{\gamma}) + D_{jk}(\boldsymbol{\theta}, \hat{\gamma}) > 0$. As this holds for $\boldsymbol{\gamma}_j$, it also holds for any seller, s , with $\boldsymbol{\gamma}_s \leq \boldsymbol{\gamma}_k$.

For (i), suppose the converse is true and that $\boldsymbol{\gamma}_k > \boldsymbol{\gamma}_s$, but $\beta_k < \beta_s$. Let $\boldsymbol{\gamma}_k - \boldsymbol{\gamma}_s := \boldsymbol{\gamma}'$ and $\hat{\boldsymbol{\theta}}'' := (1 - \varepsilon)\boldsymbol{\theta} + \varepsilon\boldsymbol{\theta}_{ks}$. As per the argument above, $C_k(\boldsymbol{\theta}) > C_s(\boldsymbol{\theta})$ and hence

$B_{ks}(\boldsymbol{\theta}, \gamma') > 0$. Given that $\gamma' > 0$, it must also be the case that $D_{ks}(\boldsymbol{\theta}, \gamma') > 0$. It immediately follows that $\pi_P(\hat{\boldsymbol{\theta}}'') > \pi_P(\boldsymbol{\theta})$.

Proof of Theorem 3

By the envelope theorem, $\frac{\partial \mathbb{E}[\pi_P(\boldsymbol{\theta}) | \Phi = \Phi_i]}{\partial \gamma_j}$ is linear in γ_j . This follows from the fact that the following result holds:

$$\frac{\partial \mathbb{E}[\pi_P(\boldsymbol{\theta}) | \Phi = \Phi_i]}{\partial \gamma_j} = \mathbb{E}[p_j^* \left(\frac{\partial \alpha_j(\boldsymbol{\theta})}{\partial \gamma_j} \right)] + \sum_i \sum_k 2a\beta_i(\boldsymbol{\theta}^*) \hat{c}_{ik} p_i^* \frac{\partial p_k^*}{\partial \gamma_j} | \Phi = \Phi_i.$$

The expression $\frac{\partial \alpha_j(\boldsymbol{\theta})}{\partial \gamma_j}$ is linear and increasing in γ_j . $\frac{\partial p_k^*(\boldsymbol{\theta})}{\partial \gamma_j} = -\frac{1}{2} \{R_S\}_{j,k}$ for all $k \neq j$ and $\frac{\partial p_j^*(\boldsymbol{\theta})}{\partial \gamma_j} = 1 - \frac{1}{2} \{R_S\}_{j,j}$, where $\{R_S\}_{ij}$ denotes the ij th component of the matrix R_S . By (A1), the sum which constitutes the second term in the right hand side of the above expression is positive and linear in γ_j . For small changes in γ_j , then, profit is approximately increasing linearly in γ_j . If $\mathbb{E}[\pi_P(\boldsymbol{\theta})]$ were increasing and linear in γ_j , the following statement holds:

$$\sum_i \mathbb{E}[\beta_i \tilde{\gamma}_i | \Phi = \Phi_j] \geq \sum_i \mathbb{E}[\beta_i \tilde{\gamma}_i | \Phi = \Phi_k] \leftrightarrow \mathbb{E}[\pi_P(\boldsymbol{\theta}) | \Phi = \Phi_j] \geq \mathbb{E}[\pi_P(\boldsymbol{\theta}) | \Phi = \Phi_k].$$

Proposition 7 indicates that $\beta_i = \beta_i(\tilde{\gamma}_i)$ where $\beta_i'(\cdot) \geq 0$. The following result holds:

$$\sum_i \mathbb{E}[\beta_i(\tilde{\gamma}_i) \tilde{\gamma}_i | \Phi = \Phi_2] = \sum_i (\mathbb{E}[\beta_i(\tilde{\gamma}_i) \tilde{\gamma}_i | \Phi = \Phi_1] + \mathbb{E}[\beta_i(\tilde{\gamma}_i) \varepsilon_i]).$$

It is clear that $\mathbb{E}[\beta_i(\tilde{\gamma}_i) \varepsilon_i] \geq 0$ as $\text{cov}(\beta_i(\tilde{\gamma}_i), \varepsilon_i) \geq 0$ and $\mathbb{E}[\varepsilon_i] = 0$. Hence:

$$\sum_i \mathbb{E}[\beta_i(\tilde{\gamma}_i) \tilde{\gamma}_i | \Phi = \Phi_2] \geq \sum_i \mathbb{E}[\beta_i(\tilde{\gamma}_i) \tilde{\gamma}_i | \Phi = \Phi_1].$$

Of course, for large changes in γ_j , $\mathbb{E}[\pi_P(\boldsymbol{\theta})]$, is not linear in γ_j . This is because j 's price is linearly increasing in γ_j , and hence j 's profit is a function of γ_j^2 . Given that (A1) holds, $\frac{\partial^2 \mathbb{E}[\pi_P(\boldsymbol{\theta})]}{\partial^2 \gamma_j} > 0$ for all j , which in turn implies that if $\beta_i(\tilde{\gamma}_i) = \beta_j$ that profit under Φ_2 would be larger than Φ_1 . This implies that:

$$\sum_i E[\beta_i \tilde{\gamma}_i | \Phi = \Phi_2] \geq \sum_i E[\beta_i \tilde{\gamma}_i | \Phi = \Phi_1] \rightarrow E[\pi_P(\boldsymbol{\theta}) | \Phi = \Phi_j] > E[\pi_P(\boldsymbol{\theta}) | \Phi = \Phi_k],$$

which is the result in (i).

With regards to (ii), we consider the case where c increases from 0. When $c = 0$, $\theta_c = 1$ for either distribution, as per Proposition 3. Let θ_G denote the realisation probability of a graph, G , in which each segment observes every seller except that i does not observe a seller k and thus observes $m - 1$ sellers. Define $\theta_\tau := \theta_c - \theta_G$. We consider the ex post expression $\frac{\partial \pi_P(\boldsymbol{\theta}^*; \boldsymbol{\gamma})}{\partial \theta_\tau} = \pi_{\theta_\tau}$.

When c is sufficiently small, $\pi_{\theta_\tau} \geq 0$ when $\boldsymbol{\theta} = \boldsymbol{\theta}_c$, in which case the optimal solution is $\boldsymbol{\theta}^* = \boldsymbol{\theta}_c$. It is clear that: $\frac{\partial \pi_{\theta_\tau}}{\partial c} < 0$, $\frac{\partial \pi_{\theta_\tau}}{\partial \gamma_k} < 0$ and $\frac{\partial |\pi_{\theta_\tau c}|}{\partial \gamma_j} > 0$ for some $j \neq k$. Furthermore, $\frac{\partial |\pi_{\theta_\tau c}|}{\partial \gamma_j}$ is independent of γ_k , as $\mathbf{p}^* = \boldsymbol{\gamma} - \frac{1}{2}C(\boldsymbol{\theta})\boldsymbol{\gamma}$. Abusing notation slightly, we can then write π_{θ_τ} as a function of γ_k and $\boldsymbol{\gamma}_{-k}$, the $(m - 1) \times 1$ vector of quality parameters not including k , $\pi_{\theta_\tau}(\gamma_k, \boldsymbol{\gamma}_{-k})$.

For any $(\gamma_l, \boldsymbol{\gamma}_h) = (\gamma_k, \boldsymbol{\gamma}_{-k})$, there exists a threshold level of c , $c'(\gamma_l, \boldsymbol{\gamma}_h)$ such that if $c \geq c'(\gamma_l, \boldsymbol{\gamma}_h)$ then $\pi_{\theta_\tau}(\gamma_l, \boldsymbol{\gamma}_h) \leq 0$, but if $c < c'(\gamma_l, \boldsymbol{\gamma}_h)$ then $\pi_{\theta_\tau}(\gamma_k, \boldsymbol{\gamma}_{-k}) > 0$. Note that, when $c = c'(\gamma_l, \boldsymbol{\gamma}_h)$, $\pi_{\theta_\tau}(\gamma_k, \boldsymbol{\gamma}_{-k}) > 0$ if $\gamma_k > \gamma_l$ and $\boldsymbol{\gamma}_{-k} \leq \boldsymbol{\gamma}_h$ or $\gamma_k < \gamma_h$ and $\gamma_k \geq \gamma_l$. Let $\hat{\gamma}_H = \gamma_h + \varepsilon_H$ and $\hat{\gamma}_L = \gamma_l + \varepsilon_L$, the lowest and highest possible values the quality of a seller can take when $\tilde{\gamma}_i \sim \Phi_2$. As Φ_2 is symmetric, it must be the case that:

$$\Pr(\tilde{\gamma}_j = \hat{\gamma}_H | \Phi = \Phi_2) = \Pr(\tilde{\gamma}_k = \hat{\gamma}_L | \Phi = \Phi_2) > \Pr(\tilde{\gamma}_j = \hat{\gamma}_H | \Phi = \Phi_1) = 0.$$

Let $\hat{\boldsymbol{\gamma}}_H$ denote an $(m - 1) \times 1$ vector with components all equal to $\hat{\gamma}_H$. When $c = c'(\hat{\gamma}_L, \hat{\boldsymbol{\gamma}}_H)$, $\Pr(\pi_{\theta_\tau} < 0 | \Phi = \Phi_2) > \Pr(\pi_{\theta_\tau} < 0 | \Phi = \Phi_1) = 0$. It follows that:

$$E[\hat{\varphi}(\boldsymbol{\theta}^*) | c = c'(\hat{\gamma}_L, \hat{\boldsymbol{\gamma}}_H), \Phi = \Phi_2] < m = E[\hat{\varphi}(\boldsymbol{\theta}^*) | c = c'(\hat{\gamma}_L, \hat{\boldsymbol{\gamma}}_H), \Phi = \Phi_1].$$

Suppose the largest component of $\boldsymbol{\gamma}$ is γ_h . Let G_j denote the graph in which seller h

and only h is observed by every consumer segment. Let $\boldsymbol{\theta}_H$ denote the probability vector in which $\theta_h = 1$.

Let γ_s be the largest component of the vector $\boldsymbol{\gamma}_{-k}$. Define G_s as a graph which is identical to G_j but where i observes h and seller s , and let θ_s denote the probability that this graph is realised. Let $\theta_d = \theta_s - \theta_h$. For a given quality vector $\boldsymbol{\gamma}$, when c is sufficiently large, $\pi_{\theta_d} < 0$ when $\boldsymbol{\theta} = \boldsymbol{\theta}_H$ and in this case $\boldsymbol{\theta} = \boldsymbol{\theta}_H$ is an optimal solution to the platform owner's problem. As before, it is clear that: $\frac{\partial \pi_{\theta_d}}{\partial c} < 0$, $\frac{\partial \pi_{\theta_d}}{\partial \gamma_s} > 0$ and $\frac{\partial |\pi_{\theta_d} c|}{\partial \gamma_k} > 0$.

For a given γ_h , the incentive to increase θ_d is greatest when $\gamma_s = \gamma_h$. Consider the marginal effect of increasing θ_d from 0 when $\boldsymbol{\theta} = \boldsymbol{\theta}_H$, using the envelope theorem:

$$\frac{\partial \mathbb{E}[\pi_P(\boldsymbol{\theta})]}{\partial \theta_d} = ap_s^*(\gamma_s - p_s^*) + 2a\beta_s(\boldsymbol{\theta}^*)\hat{c}_{sh}p_s^* \frac{\partial p_h^*}{\partial \theta_d} - c\gamma_h + 2a\beta_h(\boldsymbol{\theta}^*)\hat{c}_{hs}p_h^* \frac{\partial p_s^*}{\partial \theta_d} - c\gamma_s.$$

While $\frac{\partial p_h^*(\boldsymbol{\theta})}{\partial \theta_d} = -\{R_S\}_{hs}\gamma_h$ and $\frac{\partial p_s^*(\boldsymbol{\theta})}{\partial \theta_d} = -\{R_S\}_{sh}\gamma_h$, by (A1) the above expression is increasing in γ_h .

We can see that $\frac{\partial \pi_{\theta_d}}{\partial c} < 0$ and we also know there exists a values of c such that $\boldsymbol{\theta} = \boldsymbol{\theta}_c$ and that $\boldsymbol{\theta} = \boldsymbol{\theta}_H$ are solutions to the platform owner's problem. For a given γ_h , and assuming $\gamma_s = \gamma_h$, then there exists a $c''(\gamma_h)$ such that if $\gamma_s = \gamma_h$ and $c \leq c''(\gamma_h)$, then $\pi_{\theta_d} \geq 0$ and $c > c''(\gamma_h)$ then $\pi_{\theta_d} < 0$. The analysis above relating to $\frac{\partial \mathbb{E}[\pi_P(\boldsymbol{\theta})]}{\partial \theta_d}$ directly implies that $c''(\gamma_h)$ is increasing in γ_h .

Suppose $c = c''(\hat{\gamma}_H)$. It follows that:

$$\Pr(\tilde{\gamma}_s = \hat{\gamma}_H | \Phi = \Phi_2) > \Pr(\tilde{\gamma}_s = \hat{\gamma}_H | \Phi = \Phi_1) = 0.$$

Hence $\Pr(\pi_{\theta_d} < 0 | \Phi = \Phi_2) > \Pr(\pi_{\theta_d} < 0 | \Phi = \Phi_1) = 0$, and thus:

$$\mathbb{E}[\hat{\varphi}(\boldsymbol{\theta}^*) | c = c''(\hat{\gamma}_H), \Phi = \Phi_2] > 1 = \mathbb{E}[\hat{\varphi}(\boldsymbol{\theta}^*) | c = c''(\hat{\gamma}_H), \Phi = \Phi_1].$$

Now consider the function $\hat{\varphi}'(c) = \mathbb{E}[\hat{\varphi}(\boldsymbol{\theta}^*) | c, \Phi = \Phi_2] - \mathbb{E}[\hat{\varphi}(\boldsymbol{\theta}^*) | c, \Phi = \Phi_1]$. The

above analysis shows that there exists a c' where $\hat{\varphi}'(c') < 0$ and a c'' where $\hat{\varphi}'(c'') > 0$. Furthermore, $E[\hat{\varphi}(\boldsymbol{\theta}^*)|c, \Phi = \Phi_i]$ is a continuous function, and hence $\hat{\varphi}'(c)$ is as well. Hence, by the intermediate value theorem, there exists a $c_T \in \mathbb{R}$ such that $\hat{\varphi}'(c_T) = 0$. It follows that if $c \leq c_T$, $\hat{\varphi}'(c) \leq 0$ and if $c > c_T$, $\hat{\varphi}'(c) > 0$.

Proof of Proposition 8

For the second claim in the Proposition, note that the proof of Proposition 4 does not rely on the assumption that $c_i = c$ for all i . In which case, for any optimal probability vector, it must be that $C_j(\boldsymbol{\theta}^*) = C_k(\boldsymbol{\theta}^*)$ for all j, k pairs. Consider a probability vector $\boldsymbol{\theta}$ in which this condition holds and $w_{ij}(\boldsymbol{\theta}) \neq w_{ik}(\boldsymbol{\theta})$ for some consumer segment i and seller pair j, k , and consider the following transformation:

Take a graph τ generated with probability $\theta_\tau > 0$ under $\boldsymbol{\theta}$. As before, let τ_{jk} denote the graph generated by a neighbourhood switch between two sellers, j and k being performed on the graph τ . Let Y denote the total number of possible switches between any pair of sellers $j, k \in S$.

Let $\tilde{\boldsymbol{\theta}}$ denote the following probability vector. Suppose $\theta_\tau > 0$. Then the probability that τ is realised in $\tilde{\boldsymbol{\theta}}$ is $\frac{\theta_\tau}{Y+1}$, which is also equal to the realisation probability of each τ_{ij} for $i, j \in B$ where $i \neq j$. As $C_j(\boldsymbol{\theta}) = C_k(\boldsymbol{\theta})$ and sellers are identical, $\pi_P(\tilde{\boldsymbol{\theta}}) = \pi_P(\boldsymbol{\theta})$. Furthermore, $w_{ij}(\tilde{\boldsymbol{\theta}}) = w_{ik}(\tilde{\boldsymbol{\theta}})$ for all j, k , and hence if $\boldsymbol{\theta}$ was an optimal probability vector, so too will be $\tilde{\boldsymbol{\theta}}$.

For the first claim, we can show that for all optimal baseline distribution of networks, $\boldsymbol{\theta}_b^*$, $\varphi_i^\tau = \lfloor \varphi_i \rfloor$ or $\varphi_i^\tau = \lceil \varphi_i \rceil$ for some $\varphi_i \in [1, m]$ for all i . To see this, we note that the proof of Theorem 2 can be applied to each buyer node individually by noticing that $\tilde{\boldsymbol{\theta}}$ is analogous to $\bar{\boldsymbol{\theta}}$. Applying the same proof method to $\tilde{\boldsymbol{\theta}}$ then yields the same conclusion as that Theorem.

Suppose contrary to the first statement in the proposition that $\hat{\varphi}_i(\boldsymbol{\theta}^*) = \hat{\varphi}_k(\boldsymbol{\theta}^*) \in (1, m)$ for some optimal $\boldsymbol{\theta}^*$. Let $\varphi_i^\tau = \lfloor \hat{\varphi}_i(\boldsymbol{\theta}^*) \rfloor$ where $E_{ij} \notin G_\tau$ for i and seller j . Now consider a graph G'_τ , defined as $G_\tau + E_{ij} = G'_\tau$, realised with probability θ'_τ .

Define θ'_ϕ in the same way for consumer segment k . Let $\theta''_\tau = \theta'_\tau - \theta_\tau$, $\theta''_\phi = \theta'_\phi - \theta_\phi$ and $p_j^*(\boldsymbol{\theta}^*) = p_j^*$. By the envelope theorem and the fact that $\varphi_i^\tau = \lfloor \varphi_i \rfloor$ or $\varphi_i^\tau = \lceil \varphi_i \rceil$ for all i :

$$\frac{\partial \pi_P(\mathbf{p}; \boldsymbol{\theta}_b^*, v)}{\partial \theta''_\tau} = -ap_j^*(\gamma - p_j^* + 2 \sum_{k \neq j} c_i(p_k^* - \gamma)) + \sum_k \sum_{s \neq k} 2a\beta_k(\boldsymbol{\theta}^*) \hat{c}_{ks} p_k^* \frac{\partial p_s^*}{\partial \theta''_\tau} = 0,$$

with the analogous result for θ_ϕ . Recall that $(p_k^* - \gamma) < 0$ and by the result in Proposition 1, it must be that $|\frac{\partial p_s^*}{\partial \theta''_\tau}| > |\frac{\partial p_s^*}{\partial \theta''_\phi}|$ for all s , with $\frac{\partial p_s^*}{\partial \theta''_\tau}, \frac{\partial p_s^*}{\partial \theta''_\phi} < 0$. Hence, if the above statement holds for $\frac{\partial \pi_P(\mathbf{p}; \boldsymbol{\theta}_b^*, v)}{\partial \theta''_\tau}$, it cannot hold for $\frac{\partial \pi_P(\mathbf{p}; \boldsymbol{\theta}_b^*, v)}{\partial \theta''_\phi}$ when $\hat{\varphi}_i(\boldsymbol{\theta}^*) = \hat{\varphi}_k(\boldsymbol{\theta}^*)$, yielding a contradiction. It follows that all optimal baseline distribution of networks, $\boldsymbol{\theta}_b^*$, $E[\varphi_i] < E[\varphi_k]$, and hence $\boldsymbol{\theta}_b^*$ induces a $\boldsymbol{\theta}^*$ such that $\hat{\varphi}_i(\boldsymbol{\theta}^*) < \hat{\varphi}_k(\boldsymbol{\theta}^*)$.

If $\lfloor \hat{\varphi}_i(\boldsymbol{\theta}^*) \rfloor = m$, then $\lfloor \hat{\varphi}_k(\boldsymbol{\theta}^*) \rfloor = 1$ and in which case $\hat{\varphi}_i(\boldsymbol{\theta}^*) = \hat{\varphi}_k(\boldsymbol{\theta}^*)$. Similarly, if $\lceil \hat{\varphi}_k(\boldsymbol{\theta}^*) \rceil = 1$ then $\lceil \hat{\varphi}_i(\boldsymbol{\theta}^*) \rceil = 1$. It follows that $\hat{\varphi}_i(\boldsymbol{\theta}^*) \leq \hat{\varphi}_k(\boldsymbol{\theta}^*)$ for any optimal probability vector $\boldsymbol{\theta}^*$.

References

- [1] AKBARPOUR, M., MALLADI, S., AND SABERI, A. Just a few seeds more: Value of network information for diffusion. Tech. rep., Graduate School of Business, Stanford University, 2020.
- [2] AMIR, R., ERICKSONZ, P., AND JIN, J. On the microeconomic foundations of linear demand for differentiated products. Discussion Paper Series, School of Economics and Finance 201507, School of Economics and Finance, University of St Andrews, July 2015.
- [3] ARMSTRONG, M. Competition in two-sided markets. *The RAND Journal of Economics* 37, 3 (2006), 668–691.
- [4] ARMSTRONG, M., AND ZHOU, J. Paying for prominence. *The Economic Journal* 121, 556 (2011), F368–F395.

- [5] ARMSTRONG, M., AND ZHOU, J. Consumer information and the limits to competition. Economics Series Working Papers 888, University of Oxford, Department of Economics, Nov. 2019.
- [6] BALLESTER, C., CALVÒ-ARMENGOL, A., AND ZENOU, Y. Who's who in networks. wanted: The key player. *Econometrica* 74, 5 (2006), 1403–1417.
- [7] BAYE, M., DE LOS SANTOS, B., AND WILDENBEEST, M. Search engine optimization: What drives organic traffic to retail sites? *Journal of Economics Management Strategy* (2016).
- [8] BAYE, M. R., AND MORGAN, J. Brand and price advertising in online markets. *Management Science* (2009).
- [9] BELHAJ, M., DERIAN, F., AND SAF, S. Targeting in networks under costly agreements. Tech. rep., Aix-Marseille University, 2020.
- [10] BIMPIKIS, K., EHSANI, S., AND ILKILIÇ, R. Cournot competition in networked markets. *Management Science* (2018).
- [11] BLOCH, F., AND QUEROU, N. Pricing in social networks. *Games and Economic Behavior* 80 (2013), 243 – 261.
- [12] BONACICH, P. Factoring and weighting approaches to status scores and clique identification. *The Journal of Mathematical Sociology* 2, 1 (1972), 113–120.
- [13] BRAMOULLÉ, Y., KRANTON, R., AND D'AMOURS, M. Strategic interaction and networks. *American Economic Review* 104, 3 (March 2014), 898–930.
- [14] CHARLSON, G. Rating the competition: Seller ratings and intra-platform competition. Tech. rep., Cambridge INET, 2021.
- [15] CHERNEV, A., BOCKENHOLT, U., AND GOODMAN, J. Choice overload: A conceptual review and meta-analysis. *Journal of Consumer Psychology* 25, 2 (2015), 333–358.
- [16] CHITIKA. The value of google result positioning. Available at perma.cc/7AGC-HTDH.

- [17] COROMINAS-BOSCH, M. Bargaining in a network of buyers and sellers. *Journal of Economic Theory* 115, 1 (2004), 35–77.
- [18] DE FRANCESCO, M., AND SALVADORI, N. Bertrand-edgeworth competition in an almost symmetric oligopoly. *Studies in Microeconomics* 1, 2 (2013), 213–219.
- [19] ELLIOTT, M., AND GALEOTTI, A. The role of networks in antitrust investigations. *Oxford Review of Economic Policy* 35, 4 (12 2019), 614–637.
- [20] GALEOTTI, A., ELLIOTT, M., AND KOH, A. Market segmentation through information. Tech. rep., Cambridge University, 2020.
- [21] GALEOTTI, A., AND GOYAL, S. Influencing the influencers: A theory of strategic diffusion. *The RAND Journal of Economics* (2009).
- [22] GALEOTTI, ANDREA, G. B., AND GOYAL, S. Targeting interventions in networks. *Econometrica* (2020).
- [23] HEFTI, A., AND HEINKE, S. On the economics of superabundant information and scarce attention. *Economia* (2015).
- [24] IRELAND, N. J. The provision of information in a bertrand oligopoly. *The Journal of Industrial Economics* 41, 1 (1993), 61–76.
- [25] KIM, J. B., ALBUQUERQUE, P., AND BRONNENBERG, B. J. Online demand under limited consumer search. *Marketing Science* 29, 6 (2010), 1001–1023.
- [26] KRANTON, R. E., AND MINEHART, D. F. A theory of buyer-seller networks. *American Economic Review* 91, 3 (June 2001), 485–508.
- [27] LEDUC, M. V., JACKSON, M. O., AND JOHARI, R. Pricing and referrals in diffusion on networks. *Games and Economic Behavior* 104 (2017), 568 – 594.
- [28] LI, X. Designing weighted and directed networks under complementarities. Tech. rep., Economics Experimental Lab, Nanjing Audit University, 2019.
- [29] MATTIOLI, D. Amazon changed search algorithm in ways that boost its own products. *Wall Street Journal* (2019).
- [30] MCAFEE, R. Endogenous availability, cartels, and merger in an equilibrium price dispersion. *Journal of Economic Theory* 62 (02 1994), 24–47.

- [31] MILGROM, P., AND SEGAL, I. Envelope theorems for arbitrary choice sets. *Econometrica* 70, 2 (2002), 583–601.
- [32] MOSER, C., PHELAN, C., RESNICK, P., SCHOENEBECK, S. Y., AND REINECKE, K. No such thing as too much chocolate: Evidence against choice overload in e-commerce. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (New York, NY, USA, 2017), CHI '17, Association for Computing Machinery, p. 4358–4369.
- [33] NAGAR, K., AND GANDOTRA, P. Exploring choice overload, internet shopping anxiety, variety seeking and online shopping adoption relationship: Evidence from online fashion stores. *Global Business Review* 17, 4 (2016), 851–869.
- [34] NASDAQ. Amazon tops in product search, but can it keep growing? *Nasdaq* (2018).
- [35] POLANSKI, A. Bilateral bargaining in networks. *Journal of Economic Theory* 134, 1 (2007), 557–565.
- [36] RINGEL, D. M., AND SKIERA, B. Visualizing asymmetric competition among more than 1,000 products using big search data. *Marketing Science* 35, 3 (2016), 511–534.
- [37] ROCHET, J.-C., AND TIROLE, J. Platform Competition in Two-Sided Markets. *Journal of the European Economic Association* 1, 4 (06 2003), 990–1029.
- [38] SINGH, N., AND VIVES, X. Price and quantity competition in a differentiated duopoly. *The RAND Journal of Economics* 15, 4 (1984), 546–554.
- [39] SMITH, M. D., AND BRYNJOLFSSON, E. Consumer decision-making at an internet shopbot: Brand still matters. *The Journal of Industrial Economics* 49, 4 (2001), 541–558.
- [40] SUN, Y., ZHAO, W., AND ZHOU, J. Structural interventions in networks, 2021.
- [41] TAN, G., AND ZHOU, J. The effects of competition and entry in multi-sided markets. Tech. rep., National University of Singapore, 2020.