Essays in Strategic Decision-Making in Complex Environments

George Charlson

Mansfield College

University of Oxford

A thesis submitted for the degree of Doctor of Philosophy

Contents

1 Price Competition and Platform Design in Two-Sided Networks **7**

1.1 Introduction 7

1.2 Literature review 11

1.3 Motivating example 14

1.4 Model 16

1.5 Equilibrium characterisation in the general case 19

1.6 Comparative statics 24

1.7 Endogenous observation probabilities 28

1.8 Network structure and prominence 33

1.9 Discussion 42

1.10 Conclusion 45

2 Rating the Competition: Seller Ratings and Intra-Platform Competition **64**

2.1 Introduction 64

2.2 Literature review 68

2.3 The model 70

2.4 Determining observation probabilities 74

2.5 The price setting equilibrium . . . 76
2.6 Weighting ratings . . . 80
2.7 The profitability of matching environments . . . 88
2.8 Consumer welfare . . . 94
2.9 Discussion . . . 96
2.10 Conclusion . . . 100

3 To Lead or To Follow? The Incentives of Party Leadership and Legislators in a Multi-Principal Environment 117

3.1 Introduction . . . 117
3.2 Literature review . . . 121
3.3 The model . . . 123
3.4 Analysis . . . 133
3.5 Commitment . . . 137
3.6 Discussion . . . 147
3.7 Conclusion . . . 149

Bibliography 160

Acknowledgments

Thank you to Alex Teytelboym and Mark Armstrong, both of whom provided supervisory and advisory support on this thesis. I would also like to thank Sanjeev Goyal, Mihai Manea, Matthew Elliott, 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 Chapter 1. Finally, I thank Claire Jones, who helped copy-edit the document.

Abstract

Chapter 1: The owners of large online platforms, like Airbnb or Amazon Marketplace, are able to affect the probability that a given buyer observes a given seller on their platform. Buyers and sellers therefore interact on a bipartite graph with links between them weighted by the probability of observation. Platform owners are then able to design the network by choosing the observation probabilities. A seller's price is decreasing in their Bonacich centrality. In order to maximise profits, the platform owner connects an entrant to buyers according to a measure of their own centrality in the graph. While a network where buyers observe and trade with all sellers with probability one maximises consumer surplus, such a network does not necessarily maximise the platform's profit. The platform owner faces a trade-off between increasing observability and reducing competition, which explains why buyers tend not to be able to observe all sellers on online platforms. If sellers are identical, increasing seller prominence for any consumer segment increases competition, and profit is maximised when it is randomised across all sellers.

Chapter 2: Product ratings and reviews are commonplace on large online platforms, like Airbnb and Amazon Marketplace. One use for these ratings is to order search results. Platform owners are able to choose the extent to which ratings can be used to determine the probability a given seller is observed. Since demand is higher for high-quality products, there is an incentive to increase the probability that highly rated sellers are observed by biasing search results towards them. Biasing search results in this way results in competition being more concentrated, reducing prices. In markets where demand is sensitive to quality, a search process that is biased towards highly rated sellers

is more profitable than an environment where matches are random. However, when the variance in idiosyncratic preference is relatively large, the loss in profits associated with missing beneficial matches renders a random match environment more profitable. Whether it is profitable to use ratings as a means of ordering search results therefore depends on the properties of the market(s) in which the platform operates.

Chapter 3: Politicians serve two masters: voters and their own party leadership, and allocate effort between tasks that benefit one or the other of these principals. Depending on the extent to which they value the incentives offered by the party leader, a politician may choose to build a reputation either as an effective legislator, increasing their probability of receiving those incentives, or as a representative of the people, increasing their probability of re-election. This framework allows us to analyse the party leader's incentives in the case where they lack commitment power: the legislative benefit affect the probability that politicians are re-elected, and the party leader cannot to take this into account when they are unable to contract with the legislator ex-ante. The party leader therefore suffers from a lack of ability to commit. They may also face second-mover disadvantage because they can only reward the politicians who are re-elected, implying that the effectiveness of their incentive is reducing as the first-best level of legislative effort becomes large.

Introduction

We examine two disparate environments across three chapters. The first two chapters analyse competition and platform design in networked markets hosted on a platform. The third chapter characterises the incentives of a politician whose time allocation is distorted by two principals: a representative voter and their party leader in the legislature.

The two set-ups are both complex in the sense that one decision maker is constrained in some way in designing an environment that incentivises another decision maker (or decision makers) to maximise the payoff of the designer. In maximising aggregate profits, the platform owner in the first two chapters must account for the strategic interaction between a large number of sellers, and is constrained in the ability to affect which sellers interact with one another. For example, we examine cases where the platform owner is only able to affect the probability that certain types of sellers are observed, or can only choose the probability that an entrant is observed, holding other observation probabilities fixed.

In Chapter 3, the fact that the party leader can only reward politicians that are re-elected means that the leader faces incentives that are non-standard compared to much of the traditional literature on moral hazard. By incentivising a politician to allocate effort towards legislative tasks, the party leader is more likely to be able to achieve their short-term goals, but makes the politician's re-election less likely. Furthermore, the fact that the party leader moves second influences the effectiveness of the party leader's action: while spending a large amount of effort on legislative tasks may impress the party, doing so makes it less likely that the politician is re-elected, and hence they cannot receive any reward the party leader offers them.

Chapter 1 considers the case where the owner of a large, online platform is able to affect the probability that groups of consumers who are tied by some characteristic, like age or geographic location, observe horizontally differentiated sellers. Sellers compete with one another by setting a single price, which is chosen prior to knowing which buyers observe them.

The interactions between buyers and seller are captured by a bipartite graph with edges weighted according to the observation probabilities of the buyers. Using this framework, we characterise the equilibrium price vector and find that, under certain conditions, a seller's price is decreasing in their Bonacich centrality a seller-only graph which captures the competitive overlap between sellers. Sellers who are more central in this graph face more direct and/or indirect competition than other sellers, and therefore set a lower price.

Having characterised the price vector, we consider the owner's platform design incentives. Suppose a new seller enters the market - which consumer segments should observe them, and with what probability? Assuming the platform owner receives a proportion of the profits from each seller, segments who observe fewer sellers in expectation and are less central in the bipartite graph should observe the entrant with a higher probability than segments who have larger values on both of these metrics. If a consumer segment i observes few sellers in expectation and has low centrality, then i observing the entrant has a relatively small effect on competition and therefore prices across the network.

We show that while the optimal network in terms of consumer surplus is one in which every seller observes every buyer with probability 1, the platform owner has an incentive to reduce observation probabilities in order to reduce competition. This incentive is consistent with empirical work which suggests that buyers only observe a relatively small fraction of products available in any given market on large platforms like Amazon.

Our analysis also indicates that making a seller prominent imposes a cost on the owner. Assuming each seller is observed with some baseline probability, then making a seller more likely to be prominent increases the expected centrality of every seller in the network, reducing prices. Assuming sellers are of equal quality, we show that under certain conditions the platform owner maximises profits by setting the probability of prominence the same for each seller, as this minimises the effect on prices.

When sellers have different levels of quality, the platform owner has an incentive to ensure that high quality sellers are observed with a higher probability, but faces a trade-off in doing so in terms of more concentrated competition. Chapter 2 examines this trade-off by considering the case where part of the demand function of consumer segments is stochastic. The platform owner does not observe consumer demand directly but is able to use ratings to bias search such that buyers are more likely to observe high quality sellers.

The owner has an incentive to bias search towards highly-rated sellers, since doing so increases demand. However, it also induces more concentrated competition, as highly-rated sellers are more likely to be in competition with one another. This reduces prices, which in turn reduces profit. We find that this effect is more pronounced in cases where goods are relatively substituable.

If a consumer segment values quality more than other segments, then the platform owner has an incentive to bias their observation function towards highly-rated sellers more than less quality conscious buyers. In order to reduce the level of concentrated competition imposed by this bias, the probability that less quality conscious buyers observe high quality products is reduced.

We also compare the profits associated with an environment where sellers and buyers are randomly matched and a "ratings-based" environment in which the bias towards high quality products is fixed. As the extent to which consumers value quality increases, the difference in profits between the ratings-based and random matching environments increases. The cost associated with more concentrated competition becomes relatively less important as the cost of missing high quality matches increases.

Under certain conditions, biasing observation towards high quality sellers reduces the expected number of sellers that a given consumer segment observes compared with the random match environment. This implies that a consumer segment is more likely not to observe a product that is idiosyncratically valued highly by that segment in the ratings-based environment. As the variance of the idiosyncratic element of demand increases therefore, the random match environment becomes relatively more profitable.

Our analysis therefore suggests that the nature of the market or markets that the platform hosts is crucial for determining the extent to which ratings are used to rank products. In markets where buyers quality is a large driver in consumer demand, such as the hotel industry, ratings-based search increases profits. However, in markets driven largely by idiosyncratic demand, such as arts and crafts, random matching may be preferred.

Chapter 3 examines the effort allocation incentives of politicians in the United States Congress. While many politicians frontload their effort levels in their constituency, we note that the empirical literature also indicates that a number of politicians actually spend more time in the legislature early in their career rather than later. We explain this discrepancy with a model in which the politician is incentivised by both voters (as is usual in models of political reputation) and party leadership in Congress (which tends not to be).

Politicians are assumed to be of two types that innately value legislative and constituency effort differently. Politicians allocate their effort across legislative and constituency tasks according to their preferences and the incentives offered by the two principals. The representative voter and the party leader observe the politician's action imperfectly, and provide rewards to the politician according to a binary signal they receive that is related to the politician's effort on legislative and constituency tasks.

Voters reward politicians who generate a high signal by re-electing them, and the party leader provides legislative rewards in accordance with their belief about the politician's type. Politicians who value legislative rewards (e.g. positions on congressional committees) sufficiently have the incentive to frontload their legislative effort so that the party leader believes that they value legislative tasks, and are therefore provided with legislative rewards.

The fact that the party leader can only reward the politician if the politician is re-elected implies that the party leader is worse off if they lack commitment power. When the direct reward to being re-elected is relatively low, the legislative incentives offered by the party actually serves as an incentive to increase the effort the politician puts into constituency tasks. The legislative incentives in this case serve as indirect reward to being re-elected, and the party leader would in some cases be better off not offering legislative rewards at all than the amount they offer in equilibrium.

Furthermore, if the party leader lacks commitment power then they are unable to take into account the fact that the level of incentive they offer affects both the probability that a politician is re-elected and the type of politician who is re-elected. The fact that the party leader moves second implies that the marginal effect of their incentive on effort allocation is larger for politicians who innately value constituency effort more highly.

Our analysis also shows that the party leader may suffer from second-mover disadvantage: under certain conditions, the party leader's ability to affect the politician's optimal effort allocation is larger than if they were able to reward the politician regardless of whether they are re-elected.

Chapter 1

Price Competition and Platform Design in Two-Sided Networks

1.1 Introduction

Many economic interactions occur in settings where sellers are only able to sell to a subset of buyers. This stratification might occur, for example, because: sellers are unable to supply some consumers due to geographic constraints or due to consumer preferences (see, for example, Spiegler, 2006); consumers might be uncertain which firms are active in a market, as in Janssen and Rasmusen (2002); or the owner of a platform may deliberately not show all the products available. In such settings, a seller potentially faces different levels of competition for each individual consumer that they can supply - some buyers could be supplied by a large number of sellers, while others may only be able to buy goods from a single seller.

Online platform owners appear to choose to design search such that buyers only observe a subset of sellers. Ringel and Skiera (2016) find that 99.9% of potential buyers of LED TV sets only observe sixteen of a possible 1,124 products on a German price comparison website. Similarly, Kim, Albuquerque and Bronnenberg (2010) examine the camcorder market on Amazon and find that the median search set for consumers of these products was eleven out of a total of more than ninety camcorders available on the platform. Furthermore, which products consumers are able to observe via Amazon's search function differs depending on the search terms used, suggesting that platform design directly influences the number and nature of the sellers observed by buyers on a platform.

We model the online platform environment as a network that connects sellers with the buyers who could potentially demand goods from them. In general, the literature that models buyer-seller interactions in a network setting has focused on cases where buyers bargain with sellers they are connected to (see, e.g., Kranton and Minehart, 2001, Corominas-Bosch, 2004 and Polanski, 2007) over a single, indivisible good. This approach seems particularly relevant in relatively "thin" markets, populated by a small amount of buyers and sellers, and where the goods being sold are discrete. In thin networks with discrete goods, individual buyers and sellers can make bilateral agreements with one another easily.

However, many real-world cases of networks connecting buyers and sellers are not thin markets. For example, online platforms, such as Amazon and Airbnb have a very large number of users, with sellers interacting with a large number of buyers at any one time. This would make bilateral bargaining between users difficult, and in general prices on these platforms are not bilaterally negotiated or set by the owner of the platform itself. A natural assumption in a thick market is that sellers choose a single price which is the same for each potential buyer.

We characterise the equilibrium price setting behaviour of sellers competing for consumer segments (i.e. groups of consumers who share some characteristic like age or location) in a potentially large-scale stochastic network. The network is stochastic in the sense that we assume that there is a probability of a given buyer observing a given seller and that the actual network is only realised after price setting has taken place.

We find that a seller's price is linked to their Bonacich centrality in a network connecting sellers and consumer segments who observe them with some probability. This result is consistent with the literature on strategic interaction in networks (e.g. Ballester, Calvò-Armengol and Zenou, 2006), as well as more recent work applying price or quantity competition to network environment (see, for example, Elliott and Galleotti, 2019). When sellers are symmetric in terms of the substitutability of their goods and their own-price elasticity, we find that a seller's price is falling in their centrality.

The equilibrium pricing behaviour of sellers implies that the platform owner faces a trade-off when increasing the probability that sellers are observed by buyers: increasing the probability of observation increases the likelihood of sales, but at the same time it increases the level of competition between sellers on the network. This latter effect reduces seller prices, which in turn reduces the aggregate profits.

The same trade-offs are also present in the case where the platform owner chooses the probability that consumers observe an entrant. The platform owner has an incentive to make it more likely that the entrant is observed by consumer segments that are: (a) relatively uncontested in expectation (i.e. their probability of observing a given number of sellers is relatively low) and (b) are relatively uninfluential on seller prices. Hence, entrants should be more likely to be observed by buyers who are less likely to observe sellers on aggregate and/or observe sellers who experience relatively little competition across the whole network.

We also find that products being more likely to be prominent than others imposes a cost on the platform owner. If a product is likely to be prominent but buyers are still likely to observe at least some products, then the sellers of those products are placed in more concentrated competition with the prominent seller, and, indirectly with one another. Prominence therefore reduces the price of those sellers in direct competition with the prominent seller, which drives a fall in prices, which then produces a feedback effect: sellers in competition with sellers with lower prices must decrease their price in order to compete effectively, as prices are strategic complements (see, for example, Bulow, Geanakoplos, and Klemperer, 1985) .

The feedback effect implied by the equilibrium pricing behaviour implies that making a seller more likely to be prominent (i.e. more likely to be observed by any given consumer segment) results in lower expected prices than the case where sellers are equally likely to be prominent. When products have the same quality, the platform owner maximises their profit by completely randomising which seller is prominent for each consumer segment.

If products differ in quality, then the platform owner has an incentive to make higher quality products more prominent, as this will increase sales. However, doing so again concentrates expected competition between the more prominent sellers and the other sellers in the network. Hence the platform owner generally faces a trade-off between concentrated competition on the one hand, and increasing the sales of high-quality products on the other.

Our analysis more generally indicates that monopoly platforms have an incentive to lower competition on the platform by changing its structure. Given that the optimal structure from a consumer surplus perspective is one in which each buyer observes each seller with a probability of one, the fact that platforms may choose to only allow consumers to observe a fraction of the total number of goods on offer is harmful to welfare.

This suggests that competition authorities may wish to examine intra-platform competition, in addition to inter-platform competition. 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.

1.2 Literature review

In terms of equilibrium characterisation, our finding that a seller's price in equilibrium is decreasing in their Bonacich centrality in a seller-only network is consistent with Ballester, Calvò-Armengol and Zenou (2006), which finds that if direct effects are sufficiently small, the equilibrium action of each player in a network is proportional to their Bonacich centrality in a network. The price setting behaviour of sellers are strategic complements in this setting, which might suggest, as in Bramoullé, Kranton and D'Amours (2014), that the more central a seller in the network, the higher their price.

However, the model here does not conform to that result due to the fact that the underlying relationship between buyers and sellers is determined by the preferences of buyers, which are pinned down by a quadratic utility function. As the number of options a buyer considers increases, their demand for a given seller's good at any non-zero price decreases. This "competition effect" outweighs the strategic complementarity effect found in Bramoullé, Kranton and D'Amours (2014), and hence a seller's equilibrium price is decreasing in their Bonacich centrality.

The link we find between pricing and centrality is consistent with earlier work in industrial organisation, which examines the role "captive" buyers have on optimal pricing. Ireland (1993) and McAfee (1994) consider a framework in which sellers of a homogeneous good compete for consumers and have “independent reach” - a buyers observe each firm with some firm-specific probability, and the fact that a consumer observes one firm does not affect the probability that the consumer observes another firm. The unique mixed equilibrium of this game is one in which the lowest price within each seller’s strategy set is the same, but the strategy set of the seller with the largest proportion of captive consumers contains the highest maximum price. De Francesco and Salvadori (2013) and Armstrong and Vickers (2019) extend this analysis to cases where firms have different capacities and only the largest firm has captive customers respectively.

The relationship between the price of competing sellers in a network and their Bonacich centrality in that network found here is consistent with other studies which examine pricing in networks. For example, Candogan, Bimpikis and Ozdaglar (2012) and Chen, Zenou and Zhou (2018) analyse monopolistic pricing in cases where there are networked consumers.

Cominetti, Correa and Stier-Moses (2009) and Chawla and Roughgarden (2009) compare the efficiency of competitive equilibria of a game involving Bertrand competition between firms who act as intermediaries who control the flow of some good (for example) by controlling the edges of a network compared with the optimally efficient network flow. Similarly, Melo (2017) employs a discrete choice framework to show that a node’s Bonacich centrality determines the shape of its market demand.

Recently, a literature on competition networks has developed. Bimpikis, Ehsani and Ilkiliç (2018) examines networked quantity competition. Firms potentially compete for multiple markets, and each firm has a non-separable cost function, such that the quantity produced in one market affects the marginal cost of production in another.

In this setting, it can be shown that quantities in that model are proportional to their Bonacich centrality with a negative decay factor, which captures the idea that the more competitors a firm is connected to via a market, the lower their profit.[1]

Elliott and Galeotti (2019) show that in a Hotelling-type model in which sellers compete on price and are differentiated by location, a seller j's price is determined not only by the sellers with whom directly j competes, but also those sellers in other markets that compete with j's competitors. This result is a corollary of our more general characterisation of the price equilibrium of the sellers on the platform.

The way in which a platform can influence equilibrium prices and profits by choosing observation probabilities represents an intervention in the network that differs considerably to approaches deployed in other contexts. For example, Ballester, Calvò-Armengol and Zenou (2006) use their characterisation of equilibrium to identify the key player in the network, the removal of whom would allow a central planner to reduce total activity the most with the removal of a single player.

Birge, Candogan and Chen (2018) construct a model in which firms are connected to buyers on a network controlled by a platform and choose their price in order to attract buyers with different valuations of a single good. They characterise the effect of network structure on the profitability of commissions and subscriptions from a platform's perspective, finding that loss in revenues is potentially unbounded when all sellers are charged either the same commission or the same subscription fee regardless of their location in the network.

Galeotti, Golub and Goyal (2019) examines a case where a central planner can partly determine (at an exogenously imposed cost) the payoff of players in a network,

[1] Bonacich (1987) calls the centrality measure in the case where the scalar is negative a measure of "power".

which in turn determines the equilibrium actions of each player. While there is similarity between this approach and the one utilised here, our intervention in the network is on its design, rather than on the payoff functions of the agents. Furthermore, the cost of intervention in our model arises endogenously from the effect that changing observation probabilities has on increasing competition.

We examine the effect of prominence on price competition. Armstrong, Vickers and Zhou (2009) analyses the case where consumers engage in costly search to learn the price and match value of a series of products. If one product is more prominent than another, then it is observed first, with buyers choosing whether to buy it or search for other products. They find that the prominent firm sets a lower price than other sellers, who set a higher price than in the case where matching is random, increasing profits. In constrast, our analysis finds that prominence causes prices to fall, because prominence increases the expected intensity of competition.

1.3 Motivating example

Large, online platform owners must design platforms in which many sellers compete for many 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) can be represented by Figure 1.1.

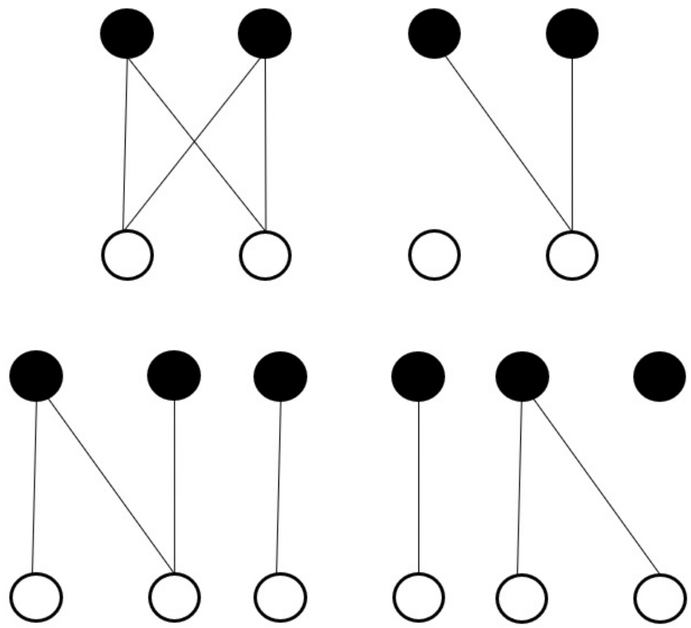

Figure 1.1: The black nodes represent sellers, the white nodes represent buyers. Assuming sellers and buyers are identical, these configurations represent all possible market types when there are two sellers and two buyers

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.

The trade-off highlighted by this simple example is one that the platform owner faces when designing the network. Increasing the probability that a buyer observes a seller increases the competition faced by every seller in the network, but this in turn increases sales.

1.4 Model

Suppose there is a finite set, 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.

Let S be a finite set of sellers, where $|S| = m$. 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 buyers interact on a platform. A consumer segment, i, observes a seller j with probability $0 \leq w_{ij} \leq 1$. Assume w_{ij} is independent of w_{ik} for all (i, j, k) and is determined exogenously. Such an observation process generates a weighted network $G = (B \cup S, E)$, where E is the set of weighted edges from B to S. G has an adjacency matrix, R, which is a zero diagonal matrix with components w_{ij}.

As an example, suppose the probability that three subsets of buyers, 1, 2 and 3, observe three sellers, X, Y and Z, can be summarised in the following table:

Observation probability	X	Y	Z
1	0.1	0.7	0.3
2	0	0.1	0.5
3	0	0	0.3

The probabilities above can be represented as a bipartite network, as shown in Figure 1.

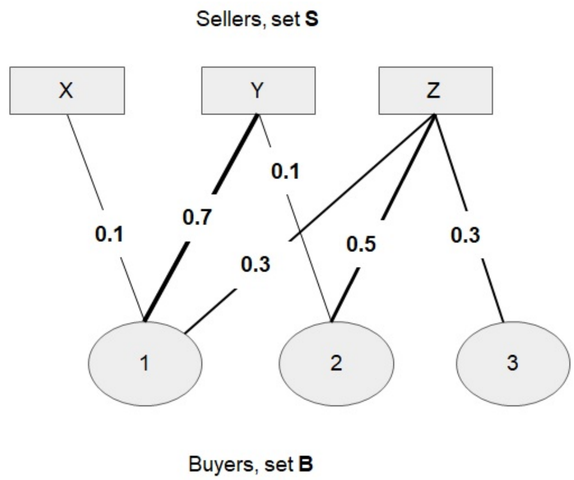

Figure 1.2: A weighted bipartite network - the number on each line indicates the probability that a given buyer observes the relevant seller

Let $\boldsymbol{p}$ denote a $m \times 1$ vector whose jth entry is $p_j \in \mathbb{R}_+$, the price of j's good. We assume that sellers set prices with full knowledge of the adjacency matrix R, but prior to the realisation of the links in the network. In other words, sellers know the buyers' observation probabilities, but do not know which buyers actually observe them.

We will assume that if a consumer segment i observes a seller j, then their demand function, $x_{ij}(.) : \mathbb{R}^m \to \mathbb{R}_+$, for product j can be expressed as follows:

$$x_{ij}(\boldsymbol{p}) = a\gamma_{ij} - b_{ij}p_j + \sum_{k=1}^{m} c_{ik}(p_k - \gamma_{ik}). \tag{1.1}$$

where a, b_{ij}, c_{ik}, γ_{ij} are all positive scalars, with $c_{ii} = 0$. We will assume throughout that a is large enough such that $x_{ij} > 0$ for each observed good: this restriction will be discussed in more detail below. As $c_{ik} > 0$ for all i, k, each product is a gross substitute for every other product.

Note that the demand function in (1.1) is an ex-post demand function, in that it is generated after the realisation of the network structure is realised, and hence after the sellers set prices. Define i's ex-ante demand function for a good j as follows:

$$\mathrm{E}[x_{ij}(\boldsymbol{p}; G)] = w_{ij}(a\gamma_{ij} - b_{ij}p_j + \sum_{k=1}^{m} w_{ik}c_{ik}(p_k - \gamma_{ik})). \tag{1.2}$$

Seller j's aggregate demand function is then defined:

$$\mathrm{E}[x_j(\boldsymbol{p}; G] = \sum_{i=1}^{n} \mathrm{E}[x_{ij}]. \tag{1.3}$$

Then, assuming the marginal cost of each seller is zero, the profit function for seller j is:

$$\mathrm{E}[\pi_j(\boldsymbol{p})] = p_j \mathrm{E}[x_j(\boldsymbol{p})]. \tag{1.4}$$

Sellers are assumed to observe the structure of graph G: they observe the matrix of observation probabilities. On the basis of this information, sellers choose their prices. A 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. Therefore, each seller's maximisation problem can be expressed:

$$\max_{p_j} \mathrm{E}[\pi_j(p_j,\ p_{\_j};\ \boldsymbol{G})]. \tag{1.5}$$

Let $\Gamma(G)$ represent the simultaneous move m-player game played on a network G with payoffs as specified in (1.5) and strategy spaces R_+.

We will also assume that there is a platform owner who has the following profit function:

$$\mathrm{E}[\pi^P(\boldsymbol{p})] = \chi \sum_{j=1}^{m} \mathrm{E}[\pi_j(\boldsymbol{p})]$$

where $0 < \chi < 1$.

1.5 Equilibrium characterisation in the general case

Buyers are not strategic agents in the game, and therefore it is possible to define a network that is strategically equivalent to the original network G, but only includes seller nodes. Such a network can be used to characterise the equilibrium of the game $\Gamma(G)$. Define α_j and β_j as follows:

$$\alpha_j(G) := \sum_{i=1}^{n} w_{ij}(a\gamma_{ij} - \sum_{k=1}^{m} w_{ik}c_{ik}\gamma_{ik})$$

and:

$$\beta_j(G) := \sum_{i=1}^{n} w_{ij}b_{ij}.$$

α_j represents the expected value of the intercept of the aggregate demand function of j, and β_j represents the expected aggregate price sensitivity. Transforming the profit function in (1.5) to put it in terms of a seller-only network yields the following payoff function for seller j:

$$\mathrm{E}[\pi_j(\boldsymbol{p})] = p_j(\alpha_j - \beta_j p_j + \sum_{k=1}^{m} \hat{c}_{jk}p_k), \tag{1.6}$$

where $\hat{c}_{jk} = \sum_{i=1}^{n} c_{ij}(w_{ij}w_{ik})$, which is therefore a measure of the strength of the connection between j and k because it measures the weighted link between the sellers and shared buyers. Rescaling (1.6) by $1/\beta_i$ and multiplying by $\frac{1}{2}$ generates the following:

$$\mathrm{E}[\tilde{\pi}_j(\boldsymbol{p})] = p_j\tilde{\alpha}_j - \frac{1}{2}p_j^2 + \sum_{k=1}^{m} \tilde{c}_{jk}p_jp_k \tag{1.7}$$

where $\tilde{\alpha}_j = \frac{\frac{1}{2}\alpha_j}{\beta_j}$ and $\tilde{c}_{jk} = \frac{\frac{1}{2}\hat{c}_{jk}}{\beta_j}$. The maximisation problem $\max_{p_i} \tilde{\boldsymbol{\pi}}(\boldsymbol{p_i})$ has the same set of first-order conditions as the one that involves maximising a vector containing the profit functions in (1.6). Define $\lambda = \max\{\widetilde{c}_{ij}|i \neq j\}$ and R_S as a symmetric zero diagonal matrix of a network G_S with components $\frac{\widetilde{c}_{ij}}{\lambda}$. This transformation yields a *competition network*, G_S, which is a projection of G. The competition network of the network in Figure 1.2 is shown in Figure 1.3 below.

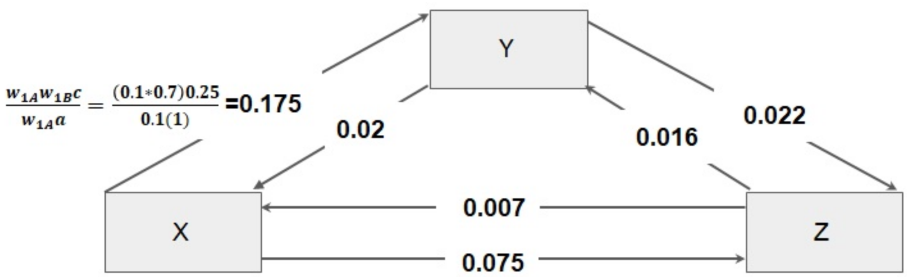

Figure 1.3: Transforming the network G into the competition network G_S. Here it is assumed that $b = 1$ and $c = 0.25$ for each buyer and $\gamma_j = 1$ for all j.

The links in G_S represent the relative importance of the connection between two sellers. In the graph in Figure 1, seller X is competing for their only (potential) buyer with seller Y. Hence, the directed link between X and Y has a much larger value than the weight between Y and X, as Y is also in competition with Z for consumer segment 2.

Let $\tilde{\boldsymbol{\alpha}}$ represent a $m \times 1$ vector with element j $\tilde{\alpha}_j$. To ensure that an equilibrium of the game $\Gamma(G)$ is unique, it is necessary to ensure that demand is positive for all sellers. Define the $m \times 1$ vector:

$$Z_{\tilde{\alpha}}(G) := [\boldsymbol{I} - \lambda R_C]^{-1}\tilde{\boldsymbol{\alpha}},$$

where R_C is the adjacency matrix of the complete seller-only graph, with $w_{ij} = 1$ for all i, j. Define $\boldsymbol{\Sigma}$ as a $m \times m$ matrix with a for each of its diagonal components and with $-c_{ij}$ for each of its ijth non-diagonal components. Given that $Z_{\tilde{\alpha}}(G) < \boldsymbol{\gamma}$, $\boldsymbol{\Sigma}(\boldsymbol{\gamma} - \boldsymbol{Z}_{\tilde{\boldsymbol{\alpha}}})$ is increasing in a. Hence, there exists an $\bar{a} \in \mathbb{R}_+$ such that if $a \geq \bar{a}$ then $\boldsymbol{\Sigma}(\boldsymbol{\gamma} - \boldsymbol{Z}_{\tilde{\boldsymbol{\alpha}}}) > 0$.

Define $C_{\tilde{\alpha}}(G_S, \lambda) = C_{\tilde{\alpha}} = [\boldsymbol{I} - \lambda R_S]^{-1}\tilde{\boldsymbol{\alpha}}$, which is the weighted transformed Bonacich centrality measure of the network G_S. The following theorem, which characterises the equilibrium price vector, then holds[2]:

Theorem 1.1. *Suppose $a > \bar{a}$. Then the game $\Gamma(G)$ has a unique Nash equilibrium in pure strategies, which is the equilibrium price vector:*

$$\boldsymbol{p}^*(\boldsymbol{G}) = \boldsymbol{C}_{\tilde{\boldsymbol{\alpha}}}(\boldsymbol{G}_{\boldsymbol{S}}, \boldsymbol{\lambda}) \tag{1.8}$$

iff $(\boldsymbol{I} - \lambda R_S)$ is positive definite.

There exists a unique Nash equilibrium price vector of $\Gamma(G)$ that is equal to the Bonacich centrality of the sellers in the network G_S multiplied by $\tilde{\boldsymbol{\alpha}}$, assuming that $(\boldsymbol{I} - \lambda R_S)$ is positive definite and at the equilibrium vector of prices, each consumer demand a non-zero amount of each good.

The vector $\tilde{\boldsymbol{\alpha}}$ is a function of the matrix G_S. The more direct connections a seller has in G, the larger the sum of the intercept terms of the demand functions of the buyers in the network, increasing the $\tilde{\alpha}$ term of each seller. This implies that a more central seller will have an incentive to set a higher price. At the same time, the more

[2]Proofs can be found in the Appendix throughout.

sellers that are connected to a given buyer, the more competition that each of those sellers faces, which leads to an incentive to lower prices. This implies that a more central seller will choose a lower price. Which of these effects dominates depends in general on the parameters of the model. The next section shows that, in the symmetric case (i.e. where the demand function of each buyer is identical), the equilibrium vector is decreasing in the network's centrality vector.

The condition on $(\boldsymbol{I} - \lambda R_S)$ in Theorem 1.1 provides a restriction on each c_{ij}, which is measures the substitutability of the model, relative to the effect own price has on demand, which is captured by a. For example, if $(\boldsymbol{I} - \lambda R_S)$ is diagonally dominant, which holds when $\sum_j \hat{c}_{ij} < \beta_i$ for all i, then it is positive definite.

Returning to the example in Figures 1.1 and 1.2, Theorem 1.1 suggests that the centrality of each seller and their price is as follows for the assumptions laid out in Figure 1.2:

	Centrality	Price
X	1.26	3.70
Y	1.20	3.98
Z	1.06	4.87

Seller Z has a (potentially) uniquely connected buyer, and therefore has lower centrality in the competition network than X or Y. As a result, Z's price is higher than either X or Y's. Z's price is the lowest as they are likely to face competition from Y, as the connection probability between Y and consumer segment 1 is relatively high (0.7).

Equilibrium in the symmetric case

Assume that $\gamma_{ij} = \gamma$, $b_{ij} = a$ for all (i, j) and $c_{ij} = c$ for all (i, j).[3] Then the demand function of the consumer segment, i, can be expressed as follows:

$$x_{ij}(\boldsymbol{p}) = a\gamma_j - ap_j + c\sum_{k=1}^{m}(p_k - \gamma_k). \tag{1.9}$$

Define $\Gamma_c(G)$ as an m-person game played on a connected network G by sellers with payoffs consistent with the demand vector implied by the equality in (1.9), and assume that the demand parameters are such that $(\boldsymbol{I} - \lambda R_S)$ is diagonally dominant (and hence positive definite) and that $a > \bar{a}$. Let $\boldsymbol{p}^*(G)$ represent the equilibrium price vector of this game and $C(G_S, \lambda)$ be vector of Bonacich centralities of the nodes in graph G_S. Define $\boldsymbol{\gamma}$ as an $m \times 1$ vector with jth element γ_j. Theorem 1.1 implies the following result:

Corollary 1.1. *The unique pure-strategy Nash equilibrium of* Γ_c *is:*

$$\boldsymbol{p}^*(G) = \frac{1}{2}(\boldsymbol{\gamma} - \boldsymbol{\gamma}^T C(G_S, \lambda)). \tag{1.10}$$

Hence, a seller's equilibrium price is decreasing in their centrality in G_S. Sellers that are connected to more isolated consumer segments (particularly those who are captive) 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 $\boldsymbol{\gamma}$.

[3] We show that this assumption is an approximation of the linear demand curve that is generated from a quadratic, quasi-linear demand curve in the first section of the Appendix.

1.6 Comparative statics

Let $\Gamma_c^{'}$ represent the symmetric version of the game described in Section 5 above but with the further restriction that $\gamma_j = \gamma$. Define the vector $\boldsymbol{p}^*(G, c, \gamma, a)$ as the equilibrium vector of this game.

Consider the effect of a change in the parameters in the consumers' utility function will have on the equilibrium price level in the game $\Gamma_c^{'}$. The following proposition concerning γ and c follows from the analysis above:

Proposition 1.1. *Suppose the game $\Gamma_c^{'}$ is being played and that $\gamma^{'} > \gamma$ and $c^{'} > c$. Then: (i)$\boldsymbol{p}^*(G, c, \gamma^{'}, a) > \boldsymbol{p}^*(G, c, \gamma, a)$ and (ii) $\boldsymbol{p}^*(G, c, \gamma, a) > \boldsymbol{p}^*(G, c^{'}, \gamma, a)$.*

The parameter γ determines the intercept of the consumers' demand function for each of the goods. If γ increases, then each seller has a direct incentive to increase their price, and, because the network is connected, the have a second-order incentive to raise their price because all other prices have also risen.

Increasing c has the effect of increasing the substitutability between the goods of each seller. An increase in c therefore increases competition between sellers, which reduces prices directly, which, again, induces further price reductions because prices are strategic complements.

Suppose a network $\boldsymbol{G'}$ is created by adding a node, i, to $\boldsymbol{G}$ and an edge, $w_{ij}^{'}$, between i and one seller $j \in \boldsymbol{G}$. Assume that $w_{ik}^{'} = 0 \; \forall k \neq j$. Then the following proposition holds:

Proposition 1.2. *Suppose the game Γ_c is being played. Then, $\boldsymbol{p}^*(\boldsymbol{G'}) > \boldsymbol{p}^*(\boldsymbol{G})$.*

A seller's price is increasing in the number of captive segments who observe it. As the number of captive buyers a seller has increases, as the relative importance of competition from other sellers decreases. Since the sellers' actions are strategic complements, an increase in the price of a seller results in an increase in the price of sellers who are in direct competition with that seller. If the network is connected, it follows that the price of all sellers will increase - the sellers connected indirectly to j increase their price, which leads to the sellers indirectly connected to them increasing their price and so on, until the entire set of sellers has increased their price.

Now consider the effect of seller entry. Suppose a network $\boldsymbol{G}''$ is created by adding both a seller node, j, an edge, $w_{ij}^{'} > 0$, between j and one consumer segment $i \in \boldsymbol{G}$. The new graph $G^{''}$ is generated as a result of entry by the seller j. The following proposition shows the effect of this entry on equilibrium prices:

Proposition 1.3. *Suppose the game Γ_c is being played. Then, $p^*(\boldsymbol{G}) > p^*(\boldsymbol{G}^{''})$ if G_S is connected.*

Seller entry unambigously reduces prices across a connected network. Every seller that consumer segment i observes experiences an increase in direct competition from j, which results in a fall in prices. This fall in price results in sellers who compete with sellers who are observed by i also lowering their prices, and so on. The result is lower prices across the network.

The above results make it possible to analyse optimal pricing and entry in the case where sellers have "independent reach", which occurs when the the fact that a consumer observes a firm does not affect the probability that the consumer observes another firm.

An illustrative example: the case of independent reach

To see the implications of the analysis above, consider the case where G is a connected network with m sellers and n consumer segments. Suppose that the proportion of consumer segments that are connected to a seller i is $\sigma_i > 0$, and the proportion of consumer segments connected to two sellers i and j is $\sigma_i \sigma_j$, and so on. In industrial organisation, this situation is called *independent reach*, and it generates a network such as the one in Figure 1.4.

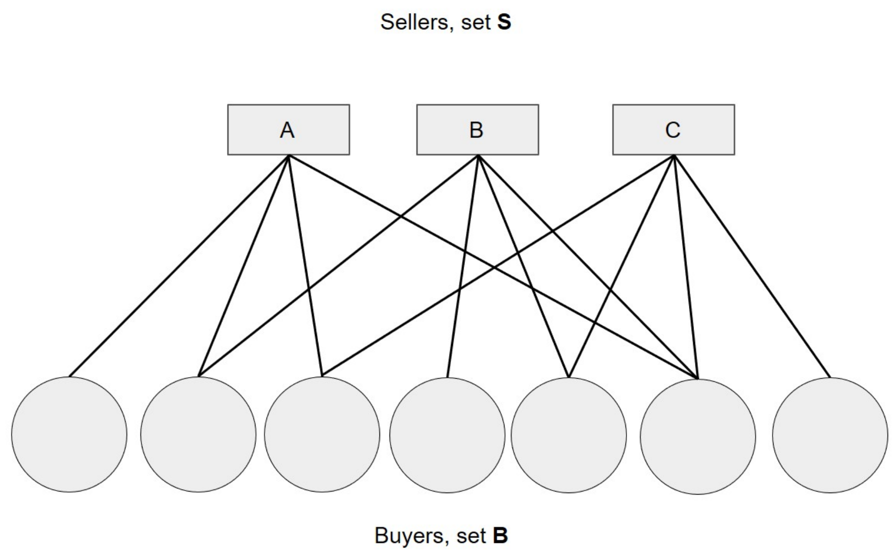

Figure 1.4: An example of a network generated by each buyer having "independent reach".

Suppose each of the m sellers have independent reach to n consumer segments. This implies that in expectation there is a positive proportion of buyers connected to a seller i uniquely and another proportion connected to i and some other seller j only and so on. Such a set up generates a network $G_\sigma(m, n, \boldsymbol{\sigma})$. Call the seller within G_σ with the weakly highest σ_i seller 1, such that $\sigma_1 \geq \sigma_2 \geq ...\sigma_m$. Theorem 1.1 implies the following about the relative prices of the sellers in the case of independent reach:

Corollary 1.2. *Suppose sellers play $\Gamma_c^{'}$ on G_σ. Then the equilibrium prices satisfy the following inequality: $p_1^c \geq p_2^c \geq ... \geq p_m^c$.*

Seller 1 has the lowest centrality in the network G_S because they have the largest proportion of captive buyers and near-captive buyers in expectation. Though the strength of their connection with a seller k, $\hat{c}_{1k}$ is larger than $\hat{c}_{ik}$ for all $i \neq 1$, 1's connections in G_S, which take into account the own-price sensitivity parameter β_1, are such that $\tilde{c}_{1k} < \tilde{c}_{ik}$ for all (i, k) pairings where $i \neq 1$. This is because the seller with the largest reach has the most captive buyers, and therefore is able to set the highest price. The same logic holds for each seller, which implies that the larger a seller's reach, the higher their price.

The result in Corollary 1.2 is consistent with Ireland (1993) and McAfee (1994) in an undifferentiated Bertrand setting, which find that a firm's expected profit is increasing in that firm's reach.

Independent reach can also be used to analyse seller entry. Suppose there is a set of m sellers with a $m \times 1$ reach parameter vector $\boldsymbol{\sigma}$, generating a network G_σ with equilibrium price vector $\boldsymbol{p}^*(G_\sigma)$ when the game $\Gamma_c^{'}$ is played. Also suppose that there is a set of $m+1$ sellers such that each of the components $1, ..., m$ of the vector $\boldsymbol{\sigma}$ are such that $\sigma_i = \sigma_i^{'}$ and let $\sigma_m^{'} > \sigma_{m+1}^{'} > 0$. Define $G_{\sigma'}$ as the $m+1$-seller, n-buyer network generated by the parameter vector $\sigma^{'}$. Proposition 1.3 has the following corollary:

Corollary 1.3. *Take any element $1, ..., m$ of the equilibrium price vector $p^*(G_\sigma)$, $p_i(G_\sigma)$. $p_j^*(G_\sigma)$ is strictly greater than $p_j^*(G_{\sigma'})$.*

In the case where the entrant has independent reach, each of the incumbent's centralities in the competition network increase as a result of entry, which in turn implies an increase in competition and therefore a decrease in equilibrium price for each seller.

1.7 Endogenous observation probabilities

Advertising

Thus far, the structure of the network and the probability that a buyer observes a given seller have been treated as exogenous. We consider the case where an entrant or the platform are able to shape the nature of competition by shaping the structure of the network. Throughout this section, assume that players play the game Γ_c once the structure of the network is realised.

First, consider the case where, through advertising (for example) a seller j, can choose the probability of being observed by a single buyer, i, holding the remainder of the network fixed. Define:

$$G^{'} \equiv G + E^{'}_{ij}$$

where $-1 \leq E^{'}_{ij} \leq 1$ and $0 \leq w_{ij}(G) + E^{'}_{ij} \leq 1$. $G^{'}$ is then the graph generated by seller j's choice of $E^{'}_{ij}$, which in turn generates an observation probability. Seller j's profit maximisation problem can be stated as follows:

$$\max_{E^{'}_{ij}} \pi^*_j(\boldsymbol{p}; G) \tag{1.11}$$

subject to the constraints above. The following lemma holds:

Lemma 1.1. *There is a unique solution to the seller's maximisation problem, $w^*_{ij} \leq 1$.*

The seller's maximisation problem is concave in $E^{'}_{ij}$, as each seller's centrality is convex in any observation probability w_{ik}. This implies that there is a unique solution to the seller's maximisation problem, which corresponds with a unique value for i's probability of observing j, w^*_{ij}. Note that there will be cases in which $w^*_{ij} = 1$.

The seller faces a trade-off: increasing w_{ij} increases revenue directly because there is a higher probability that they are observed by i. At the same time, increasing an observation probability increases the centrality of j, which results in more competition across the network.

The above trade-off is also present if we consider the case in which the platform owner maximises their profit by choosing E'_{ij} for a given network G. However, the platform owner also internalises the fact that a reduction in prices across the network also reduces the profits of sellers other than j. By the same logic as in the proof of Lemma 1, there is a single observation probability that corresponds to the solution of the platform owner's maximisation problem, which we will denote w^P_{ij}. The following proposition holds:

Proposition 1.4. *The solution to the platform owner's maximisation problem, $w^P_{ij} \leq w^*_{ij}$ with the inequality strict if $w^*_{ij} < 1$.*

The platform has an incentive to set w_{ij} below the level that is optimal for the individual seller because of the effect changing the weight has on the profits of other players in the network. Hence, even if $w^*_{ij} = 1$, if the effect of the weight change on the price of sellers other j is sufficiently large, then $w^P_{ij} < 1$.

When an individual seller chooses the probability that they are observed in the network, they do not take into account the competition effect on aggregate profits. Increasing any observation probability weakly increases competition on the network, reducing prices and platform profits. Individual sellers only take into account their own profit levels, and hence do not fully internalise this effect if they choose their own observation probabilities.

Profit-maximising seller entry

Let G denote a bipartite graph of n buyers and $m \geq 2$ sellers. Suppose that there exists a seller $j \in G$ for whom $w_{ij} = \tilde{w} > 0$ for all i, and suppose that the platform owner chooses the $n \times 1$ vector of observation probabilities, $\boldsymbol{w_j}$, whose ith element is the observation probability w_{ij}. We use this set up to represent the case where a new seller enters the market and the platform owner needs to choose the optimal set of observation probabilities to maximise profits, keeping all other observation probabilities fixed. The observation probability w^* is thus a default likelihood that buyers would observer the new seller if the platform owner did not intervene.

Define $\boldsymbol{w_j}$ as an $n \times 1$ vector whose ith entry is w_{ij}. The platform owner's maximisation problem can be stated as follows:

$$\max_{\boldsymbol{w_j}} \mathrm{E}[\hat{\pi}_P^*] = \chi \max_{\boldsymbol{w_j}} \sum_{j=1}^{m} \mathrm{E}[\hat{\pi}_j^*] \tag{1.12}$$

To solve this maximisation problem, we define the following variable, which captures how influential a given consumer segment is on the centrality of sellers in G_S, and which we will call "buyer influence":

$$C_i^B(G) := \sum_{j=1}^{m} w_{ij} C_j(G_S, \lambda). \tag{1.13}$$

Let w_{ij}^P represent an element of a vector $\boldsymbol{w_j^P}$, which is a solution to the platform's optimisation problem stated above. Denote the equilibrium price vector when $\boldsymbol{w_j} = \tilde{w}\boldsymbol{1_n}$ as the $m \times 1$ vector $\boldsymbol{\tilde{p}}$ with jth element $\tilde{p}_j$.

We characterise the first-order conditions of the platform owner's problem as follows. By the envelope theorem, while a marginal change in an observation probability, w_{ij}, changes j's price, it does not affect j's profits. Seller j's profit change with w_{ij} insofar as it affects profit directly (i.e. by generating an increase in demand from i) and through a change in price of the other sellers. Similarly, the profits of all other sellers who are not j change as a result of a change in the prices of the other sellers.

We denote the marginal benefit associated with increasing a weight w_{ij} above $\tilde{w}$ as follows:

$$\mathrm{MB}_{ij} = a\tilde{p}_j(\gamma - \tilde{p}_j) + \sum_{k=1}^{m} \frac{\partial \tilde{c}_{jk}}{\partial w_{ij}} \tilde{p}_j(\tilde{p}_k - \gamma) > 0.$$

Increasing w_{ij} increases the probability that i observes j, which, by assumption, results in positive demand and therefore increased expected profit. However, increasing w_{ij} increases the competition between the existent sellers and the new seller. Increasing the probability that a consumer segment observes a seller has a direct effect on the expected demand for every other good and the prices of those goods. The marginal cost associated with increasing w_{ij} can be written as follows:

$$\mathrm{MC}_{ij} = \sum_{k=1}^{m} \frac{\partial \tilde{c}_{jk}}{\partial w_{ij}} \tilde{p}_j(\tilde{p}_k - \gamma) + \sum_{l=1}^{m} \sum_{k=1}^{m} \hat{c}_{kl} \frac{\partial \tilde{p}_k}{\partial w_{ij}} \tilde{p}_l.$$

The solution to the maximisation problem in (1.12) equates the absolute value of these two expressions. Note that, as per Bonacich (1972), the centrality of k in G_S can be written:

$$C_k(G_S, \lambda) = \sum_{k=1}^{m} \hat{c}_{kl} C_l(G_S, \lambda).$$

The centrality of k is a linear function of $C_j(G_S, \lambda)$. As a result, it can be shown that $\frac{\partial C_k(G_S,\lambda)}{\partial w_{ij}}$ can be written as a weighted sum of both $C_j(G_S, \lambda)$ and $C_k(G_S, \lambda)$. Specifically, the weights in this expression are a function of $\sum_i w_{ik}$, the sum of observation probabilities of every consumer segment. As both of these expressions are weighted sums of the centrality of every other seller, it follows that $\frac{\partial C_k(G_S,\lambda)}{\partial w_{ij}}$ is increasing in $C_i^B(G)$.

The analysis above suggests that there is a link between the influence measure defined in (1.13) and the solution to the platform owner's maximisation problem. Theorem 1.2 makes this link explicit:

Theorem 1.2. *Suppose Γ_c is played after the platform owner chooses G_S by choosing the vector $\boldsymbol{w}_j$. If $C_i^B(G) < C_l^B(G)$ and $\sum_k w_{ik}(G) \leq \sum_k w_{lk}(G)$, then $w_{ij}^P \geq w_{lj}^P$ for any solution to the platform owner's optimisation problem.*

The marginal benefit of increasing an observation probability w_{ij} is related to the addition revenue generated by sales from j. Increasing the probability that j is observed straightforwardly results in an increase in the expected demand for j's product, increasing profit. The extent to which profit increases as a result of an increase in w_{ij} is tempered by the fact that the ex-post demand for j's good is reducing in the number of additional products buyers observe.

Figure 1.5 illustrates an example of where of the platform owner's optimal solution in the case where a new seller enters.

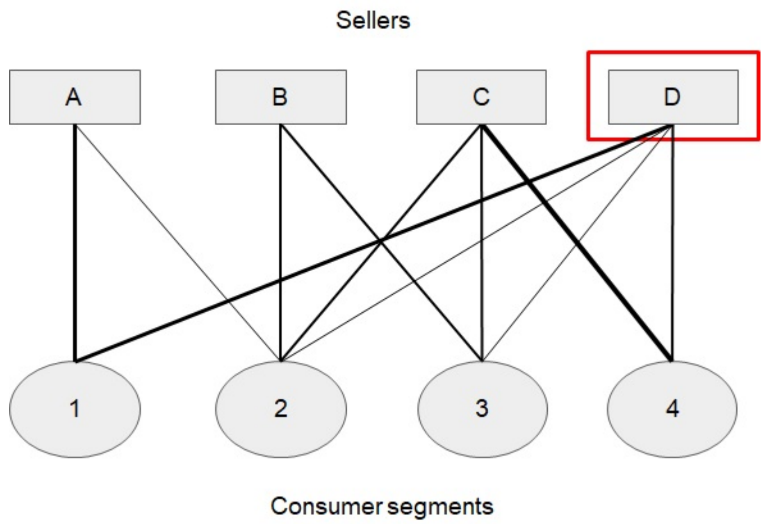

Figure 1.5: profit-maximising entry of seller D, keeping the remainder of the network fixed. The thicker the lines in the diagram, the larger the observation probability.

Consumer segments 1 and 4 are connected to only one seller, with segment 1 also having the lowest sum of observation probabilities. Hence, 1's probability of observing D in any solution to the platform owner's maximisation problem is larger than the probability that 4 observes D. The buyer influences of both 2 and 3 are both larger than the influence of 1 and 4, and hence they observe D with lower probability than 1 and 4.

1.8 Network structure and prominence

In Section 7, we saw that the platform owner faces a trade-off between increasing a seller's expected sales by increasing the probability that they there are observed by a consumer segment i and the increase in competition that this imposes on the rest of the sellers who are observed by i. In this section, we examine this trade-off more generally by characterising the network structure or structures that maximise profit and/or consumer welfare.

Consumer surplus

Assume the game Γ is being played, $a > \bar{a}$ and that the matrix $I - \lambda R_S$ is diagonally dominant. Define the expected consumer surplus of consumer segment i for a given equilibrium price vector $\boldsymbol{p^*}$ and demand vector $\boldsymbol{x_i^*}$ is as follows:

$$\mathrm{E}[CS_i(\boldsymbol{x_i^*};\boldsymbol{p}^*)] = \sum_{j=1}^{m} \frac{1}{2} w_{ij} \mathrm{E}[x_{ij}^* | w_{ij} = 1] \left(\frac{a\gamma_{ij}}{b_{ij}} + \sum_{k=1}^{m} w_{ik} \frac{c_{ik}}{b_{ij}} (p_k^* - \gamma_{ik}) - p_j^* \right) \tag{1.14}$$

Define $\mathrm{CS}(\boldsymbol{x}^*;\boldsymbol{p}^*) := \sum_i \mathrm{CS}_i(\boldsymbol{x_i^*};\boldsymbol{p}^*)$, where $\boldsymbol{x}^*$ is an $m \times n$ matrix whose ijth component is $x_{ij}^*(\boldsymbol{p^*})$. If $I - \lambda R_S$ is diagonally dominant, then it is clear from the above expression that the expected value of each $CS_i(\boldsymbol{x_i^*};\boldsymbol{p}^*)$ is falling in p^*. It is also straightfoward to show that expected consumer surplus, ceteris paribus, is increasing in the expected number of connections in G a buyer has. Clearly then, CS is a function of the network G and can be written $\mathrm{CS}(G)$. Call the set of all possible m, n bipartite networks $G(m, n)$. Define G_c as a complete network where $w_{ij} = 1$ for all (i, j) pairs. The following proposition holds:

Proposition 1.5. *Assume the game Γ is being played, $a > \bar{a}$ and that the matrix $I - \lambda R_S$ is diagonally dominant. For any $G \in G(m, n)$, $G \neq G_c$, $E[CS(G_c)] > E[CS(G)]$.*

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. This result is consistent with Bimpikis, Ehsani and Ilkiliç (2018), who find that a complete network maximises consumer welfare in Cournot setting if buyers' demand functions are homogeneous.[4]

[4] Strictly, the model of Bimpikis, Ehsani and Ilkiliç (2018) is one where firms compete for "markets"

Hence, despite the fact that in a Cournot model the sellers' actions are strategic substitutes, while in a Bertrand setting they are strategic complements, the driving logic in both cases is that competition reduces prices and a complete network maximises competition.

Profit-maximising networks

Consider first the case where $c = 0$. If c is zero, then that each seller's product is not a substitute for the other goods in the market, which implies that each interaction effect parameter linking the two sellers in G_S is equal to zero. Let $\pi(G) = \sum_{i=1}^{m} \pi_i(G)$. It follows that:

Proposition 1.6. *Suppose $c = 0$ and the game Γ_c is being played. For all $G \in G(m, n)$ where $G \neq G_c$, $\pi(G_c) > \pi(G)$.*

When goods are not substitutes, sellers are not in competition with one another. Prices are therefore set at the monopoly level. The platform owner then has an incentive to increase any probability $w_{ij} < 1$ as $x_{ij} > 0$ for all i, j. It follows that the complete network maximises seller profit, which in turn maximises the platform owner's profit.

The analysis in Section 7 implies that, despite demand from each consumer segment being strictly positive for any good they observe, the complete network does not necessarily maximise the platform's profit. The platform owner has an incentive to reduce the probability that sellers are observed in order to increase prices. In decreasing ob-

rather than consumers. They find that if markets are of the same size, which would be equivalent to homogeneous buyers in this model.

servation probabilities, the platform owner potentially (assuming a is sufficiently large) reduces demand. In order to maximise profit then, 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.

To illustrate the platform's trade-off, suppose first that $\upsilon \in [0, 1]$ and that:

$$w_{ij}(\upsilon) = \upsilon.$$

Assume that the platform owner then solves:

$$\max_{\upsilon} \hat{\pi}_P^* = \chi \max_{\upsilon} \sum_{j=1}^{m} \hat{\pi}_j^* \tag{1.15}$$

The following result holds:

Proposition 1.7. *Suppose that Γ_c is played after the platform owner's choice of υ. Then there exists a unique solution to the platform owner's profit maximisation problem, $\upsilon^* \leq 1$, with the inequality strict if c is sufficiently large. Furthermore, $\frac{\partial \upsilon^*}{\partial a} > 0$ and $\frac{\partial \upsilon^*}{\partial c} < 0$.*

Proposition 1.7 highlights the platform owner's trade-off with respect to network design. Increasing the probability that each of the seller's is observed increases profits as sales are increasing in the probability that each buyer observes each seller. At the same time, increasing sales increases competition, reducing prices. Each seller's price is a concave function of υ implying that a unique solution to the maximisation problem in (1.15) exists for $\upsilon \in [0, 1]$.

Increasing c increases the substitutability of the goods being sold on the platform. This increases the level of competition in the network for a given level of observability, v, reducing prices. The platform owner then has an incentive to reduces observability across the network to increase prices.

Increasing a, on the other hand, results in an increase in sales for a given price level. Hence, the platform owner has an incentive to increase the probability that sellers are observed, even if the implied increase in competition unambiguously reduces prices.

Proposition 1.7 implies that the platform owner potentially has the incentive to reduce the extent to which consumer segments observe sellers in order to increase prices. For a given observation probability, as the number of sellers, $m \to \infty$, the probability that any consumer segment will observe every seller on the network (which is equal to v^m) tends to zero. We discuss the policy implications of this finding in the discussion section below.

Prominence

The above analysis and the characterisation of the equilibrium price vector also indicate that the platform owner faces another trade-off when they are able to choose network structure more generally. The observation probabilities 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. Increasing a 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.

As the number of sellers and the centrality of those sellers in the in the network G_S

increases, $w_{ij} > w_{ik}$ has an increasingly large effect on prices. The centrality vector in G_S can be expressed:

$$C(G_S, \lambda)\mathbf{1} = \sum_{k=0}^{\infty} \lambda^k G_S^k \mathbf{1}.$$

It follows that $\frac{\partial^2 C_k}{\partial^2 w_{ij}} > 0$. Recall that seller prices are falling in the centrality of the sellers in this setting. Hence, setting $w_{ij} > w_{ik}$ imposes a cost upon the platform owner because the centrality measure has a feedback effect such that increasing an observation probability w_{ij} (weakly) reduces j's price, which reduces k'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.

In order to formally analyse the effects of prominence on the network, we consider the case where each seller has a baseline observation probability, $\upsilon > 0$, which we take as fixed, and that one seller is prominent for each consumer segment. Which seller is prominent for a consumer segment can differ for each segment. If a seller j is prominent for a segment i, they are observed with probability $w_{ij} = \upsilon + \xi < 1$ where $\xi > 0$. Otherwise, sellers are observed with probability υ.

Denote the probability that a seller j is prominent with any consumer segment as $\rho_j \in [0, 1]$, and assume that the probability that a seller being prominent for one consumer segment is independent of the probability that they are prominent for another seller group. Thus, $\sum_{j=1}^{m} \rho_j = 1$. Then, the platform owner chooses the vector of probabilities $\boldsymbol{\rho}$ to solve the following problem:

$$\max_{\boldsymbol{\rho}} \mathrm{E}[\hat{\pi}_P^*(\boldsymbol{\rho}; \upsilon, \xi)] = \chi \max_{\boldsymbol{\rho}} \sum_{j=1}^{m} \mathrm{E}[\hat{\pi}_j^*(\boldsymbol{\rho}; \upsilon, \xi)] \tag{1.16}$$

subject to the constraint that:

$$\sum_{j=1}^{m} \rho_j = 1.$$

We assume that the platform first chooses the prominence probability vector, the who seller becomes prominent is then realised, generating a network G and then the sellers play the game $\Gamma(G)$. This implies that the expected price vector can be expressed

$$\mathrm{E}[\boldsymbol{p}(G)] = \frac{1}{2}(\boldsymbol{\gamma} - \boldsymbol{\gamma}^T \mathrm{E}[C(G_S, \lambda)])$$

where:

$$\mathrm{E}[C(G_S, \lambda)] = \sum_{k=0}^{\infty} \lambda^k \mathrm{E}[G_S^k(\boldsymbol{\rho})]$$

Let $\boldsymbol{\rho}^*$ be a solution to the platform owner's maximisation problem. Theorem 1.3 characterises the results of the platform owner's problem as follows:

Theorem 1.3. *Suppose the platform owner chooses $\boldsymbol{\rho}$ to maximise profit and $\Gamma_c^{'}$ is played once $\boldsymbol{\rho}$ is chosen. Then, if $m, n \geq 2$, $\boldsymbol{\rho}^* = \frac{1}{m}\mathbf{1}_m$ is the unique solution to the platform owner's maximisation problem. This implies that, at the optimum, $E[w_{ij}] = \upsilon + \frac{1}{m}\xi$ for all i, j pairs.*

Making one seller more likely to be prominent is costly for the platform owner for two reasons. Firstly, prominence concentrates competition, which results in a fall in prices that in turn reduces platform profits. By randomising prominence, the platform owner ensures that the level of expected competition in the network is as low as possible given the prominence parameter ξ and υ. The feedback effect inherent within the price competition on the network results in outcomes where one seller is more prominent to a larger number of consumer segments than other sellers increases competition. The more consumer segments a seller is prominent to, the more direct and indirect competition the seller imposes upon other sellers in all consumer segments, which reduces prices.

Secondly, increasing the likelihood that one seller is more prominent than other sellers increases the probability of outcomes in which that seller is prominent in many markets and the other sellers are prominent in few or no markets. Sellers who are not prominent in many or any markets are less likely to have captive or near captive buyers. As a result, they put more relatively more weight on the competition they face from other sellers.

A seller prominent in many markets puts less weight on the competition they face; however, this effect is diminishing as the number of buyers for whom that seller is prominent increases. Mathematically, this effect arises from the fact that $\tilde{c}_{ij} = \frac{\hat{c}_{ij}}{\beta_i}$, and hence $\tilde{c}_{ij}$ is convex in β_i for all $j \neq i$. It follows that in outcomes where one seller is prominent for more consumer segments than other sellers, the sum of the weighted walks starting from each seller is greater than outcomes where seller prominence is more evenly distributed. This implies setting some prominence probability $\rho_i > \rho_j$ $j \neq i$ increases the probability of outcomes where prices and profits are lower than outcomes where prominence is more evenly distributed.

As a result of both the competition effect and the captive buyer effect described above, the unique solution to the platform owner's maximisation problem is that $\boldsymbol{\rho}^* = \frac{1}{m}\mathbf{1}_m$, assuming $m, n \geq 2$. If $n = 1$, then as one of the sellers must be prominent, any set of probabilities is a solution to the maximisation problem, and if $m = 1$, then the probability that the seller is prominent is 1 by definition. Note that the solution to the seller's problem implies that $\mathrm{E}[w_{ij}] = \upsilon + \frac{1}{m}\xi$ for all i, j pairs: the expected probability that any seller observed by any consumer segment is the same for all sellers and consumer segments.

Now suppose that there is a single seller whose product is of higher quality than the other products on the market. We capture this by assuming that there is a seller j for whom $\gamma_j > \gamma = \gamma_k \; k \neq j$, which implies that demand for j's product is higher than demand for all other sellers' products when prices are equal. Define $\boldsymbol{\rho}^{**}$ as a solution to the platform owner's maximisation problem. Proposition 1.8 examines the platform owner's choice of $\boldsymbol{\rho}$ in this case:

Proposition 1.8. *Suppose the platform owner chooses $\boldsymbol{\rho}$ to maximise profit, that $\gamma_j > \gamma = \gamma_k$, $k \neq j$ for all j and Γ_c is played once $\boldsymbol{\rho}$ is chosen. Then there exists an $\tilde{a} \in \mathbb{R}_+$such that if $a > \tilde{a}$, any $\boldsymbol{\rho}^{**}$, $\rho_j > \rho_k$, $k \neq j$, and hence $E[w_{ij}]$ >$E[w_{ik}]$ for all i and k.*

Suppose $\boldsymbol{\rho} = \boldsymbol{\rho}^{*}$. A marginal increase in the expected prominence of product j increases expected demand for j directly for all buyers, and, as $\gamma_j > \gamma$, this effect outweighs the implied decrease in demand from another seller, k, whose expected prominence must decrease because $\sum_{i=1}^{m} \rho_i = 1$.

However, increasing the expected prominence of j also decreases demand for other products for two other reasons. First, it causes a direct fall in demand for every other product because by equation (1.2) above, $\frac{\partial x_{ik}}{\partial \gamma_j} < 0$ for all $k \neq j$. Second, from $|\frac{\partial \mathrm{E}[p_i^*]}{\partial \rho_j \partial \gamma_j}|$ is strictly positive, and hence increasing the prominence of j at $\boldsymbol{\rho} = \boldsymbol{\rho}^{*}$. has the effect of decreasing the price of every seller other than j.

The latter two effects are not functions of a, but the direct increase in demand associated with an increase in j's expected prominence is linearly increasing a. Hence, if a is sufficiently large, then it is profitable to increase ρ_j above the ρ^*, so that $\rho_j > \rho_k$, $k \neq j$. When this holds, the benefit associated with increasing the probability that a buyer observes a seller with a popular product is higher than increasing any other observation probability, due to the fact that the popular product yields more revenue when it is observed. As a result, the equilibrium observation probability for the popular product is higher than the observation probability of the other products.

Proposition 1.8 highlights that platform owners have an incentive to increase the expected prominence of particular, high quality (or otherwise high demand) products, even if this results in additional direct and indirect competition on the network. This incentive may explain why platforms such as Amazon tend to highlight particularly popular or branded products to users: see, for example, Barach, Golden, and Horton (2019) and Kim, Albuquerque and Bronnenberg (2010).

An example of the expected weights of a network generated in the case analysed by Proposition 1.8 are shown in Figure 1.6.

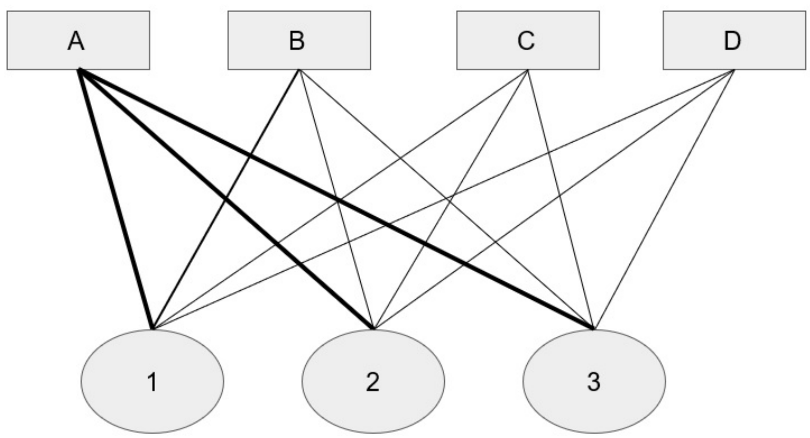

Figure 1.6: Network structure when $\gamma_A > \gamma_B = \gamma_C = \gamma_D$. The thicker the lines in the diagram, the larger the observation probability between the two connected players.

1.9 Discussion

When a platform owner's revenue is a proportion of the profits of sellers on their platform, they have an incentive to reduce the probability that buyers observe sellers. While ensuring that a consumer segment observes a seller increases sales, it produces a cost because it reduces prices through increased competition. If product substitutability is sufficiently high, the platform owner reduces observability to maintain high prices, reducing consumer welfare.

The result that platform owners have an incentive to reduce is consistent with the observed behaviour of online platforms. Several empirical studies (e.g. Kim, Albuquerque and Bronnenberg, 2010 and Ringel and Skiera, 2016) show that buyers tend only to be able to observe a fraction of the total sellers on the platform. Our analysis explain these findings by showing that it is consistent with the profit-maximising behaviour of a platform owner who takes a proportion of the profits of the sellers on the platform.

A similar logic applies to seller entry. If a seller enters the market, the platform owner has an incentive to ensure that they are more likely to be observed by consumer segments that are: (a) relatively uncontested in expectation (i.e. their probability of observing a given number of sellers is relatively low) and (b) are relatively uninfluential on seller prices. The expected demand an entrant receives from a relatively contested consumer segment is relatively less than the expected demand from an uncontested consumer segment for the same observation probability. Furthermore, the effect on the demand of other sellers of increasing the probability of consumer segment i observing the entrant is increasing the extent to which i is contested.

In Section 7, we defined the buyer influence of a consumer segment i as the sum of the centralities of the sellers observed by i weighted by the observation probabilities. The lower the buyer influence of a consumer segment, the smaller the effect on prices of other sellers associated with a seller being observed by that consumer segment. The platform owner therefore has a potential incentive to reduce the probability that some consumer segments observe entrants in order to maintain relatively high prices across the network.

Our analysis implies that seller prominence imposes a cost on the platform owner. As a seller's price is falling in their centrality in the network, a single seller being relatively more prominent to a consumer segment than other sellers results in a fall in the prices of those sellers, reducing profits. Due to the fact that the Bonacich centrality measure is convex in any one observation probability, this price effect increases as the probability that the seller is prominent increases. Hence, if one seller has to be more prominent than others, because, for example, the platform displays search results as a list of sellers, and product quality is the same, then the probability that any one seller is prominent should be equal.

However, if products have different levels of quality, then the platform owner benefits from making high quality products more observable, and low quality products less observable. Increasing the probability high-quality products are observed then imposes a trade-off: if high-quality sellers are more prominent, sales increase for a given price vector, but increasing that prominence can reduce prices. We examine this trade-off in a stochastic context in Chapter 2.

There are a few of potentially important issues left unaddressed by our analysis. We have assumed that marginal costs are zero. If sellers are heterogeneous with respect to their marginal cost, then the platform owner may have an incentive to increase the prominence of low cost sellers in the network at the expense of sellers with higher marginal costs. This would imply that some form of paid prominence could be profit-maximising from the platform owner's perspective, as Armstrong and Zhou (2011) point out.

Furthermore, we do not consider the case where the platform owner is able to segment the market by making one seller prominent with one consumer segment and less prominent with another. This is an extension to the analysis here, and would involve the same kind of trade-offs present in our model. Seller prominence is costly (assuming buyers observe sellers with relatively high probability), but is likely to be less costly than the case where they are prominent to all consumer segments, as the platform owner is able to partially segment the market.

1.10 Conclusion

We consider the case where buyers have consideration sets containing a subset of all the sellers in a market, who sell differentiated substitutes and compete on price. Whether a given buyer observes a given seller in the market is assumed to be determined after prices are set, but the probability of observation is assumed to be common knowledge. Buyers and sellers are therefore connected to one another in a weighted network, and the structure of this network affects the equilibrium price of each seller. Specifically, if the substitutability between goods is sufficiently small, there is a unique, interior pure strategy equilibrium where each seller's price is falling in their Bonacich centrality in a sellers-only network that is strategically equivalent to the original bipartite network. The more central a player in this competition network, the more competition they face, and the lower their price.

Using the characterisation of equilibrium, it is possible to see how changes in network structure affect prices. Increasing the number of captive buyers unambiguously increases the price of every buyer in the network as the direct effect of increasing the price of one player, which in turn propagates across the entire network if it is complete. In general, however, assessing the effect on the network price vector across the network is more complex. While seller entry or the addition of contested buyers increases the competition effects experienced by a seller, it also increases the relative importance of captive buyers. If the latter effect is sufficiently large, entry may result in an increase in price.

The framework can also be used to determine the optimal probability of each consumer observing a given seller from the perspective of the seller and of a platform that owns the entire the network. In general, the platform prefers the observation probabilities to be (weakly) lower than the seller, because the platform needs to take into account the fact the higher the level of competition, the lower the prices are across the network, which lowers profit. This effect is not fully internalized in the case where the individual seller choose the vector of observation probabilities.

When a new seller enters, the platform owner can increase expected sales by increasing the probability that the seller is observed. However, increasing the probability that the entrant is observed reduces prices through increased competition. The platform owner has an incentive to ensure that the entrant is more likely to be observed by consumer segments that are both relatively uncontested in expectation (i.e. their probability of observing a given number of sellers is relatively low) and relatively uninfluential on seller prices.

Consumer surplus is maximised by a complete network. For any number of buyers and sellers, a complete network maximises the level of competition between sellers, which reduces the equilibrium price level. However, platform profit is not necessarily maximised by the complete network, despite the fact that, by assumption, each consumer segment has strictly positive demand for any products they observe. Platform owners may therefore have an incentive to set the default observation probability below 1, in order to reduce the amount of competition on the platform and increase prices.

We also find that seller prominence potentially imposes a cost on the platform owner.

Due to the feedback effect inherent to actions determined by the sellers' Bonacich centralities in the network, a seller being more likely to be prominent results in an increase in competition that is larger than the corresponding decrease in competition associated with a reduction in prominence. When sellers are identical, the platform owner maximises their profit by making it equally likely that every seller is prominent.

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 inter-platform competition. As online platforms become more established and dominant, this kind of competition becomes less relevant, and the incentives to increase prices examined here become increasingly important.

Appendix

The symmetric demand function

We show that the linear demand curve for the game Γ_c is a simplified 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:

$$u_i(\boldsymbol{x_i}) = \sum_{j=1}^{m} \gamma x_{ij} - \sum_{j=1}^{m} \mu x_{ij}^2 - \sum_{j=1}^{m} x_{ij}(\sum_{k=1}^{m} \rho x_{ik}) + y_i$$

where $\mu, \rho \in \mathbb{R}_+$. Suppose that each buyer has an income of l. As $l \to \infty$ and $a > \bar{a}$, the demand for each product is positive. Define the $m \times m$ matrix $\boldsymbol{\mu}$ as follows:

$$\boldsymbol{\mu} = \begin{bmatrix} 1 & \rho & ... & \rho \\ \rho & \ddots & \rho & \vdots \\ \vdots & \rho & \ddots & \vdots \\ \rho & \dots & \rho & 1 \end{bmatrix}$$

Then, as discussed in Amir, Erikson and Jin (2015), the demand vector $\boldsymbol{x}_i$ can be written:

$$\boldsymbol{x_i} = \boldsymbol{\mu}^{-1}(\boldsymbol{\gamma} - \boldsymbol{p}).$$

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 Theorem 1.1

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 (1.9) 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 by (1.7) is as follows:

$$\tilde{\boldsymbol{\alpha}} = [\boldsymbol{I} - \lambda R_S]\boldsymbol{p}.$$

The matrix $\boldsymbol{I} - \lambda R_{\tilde{S}}$ is non-singular and the above first-order condition has a solution, which is denoted $\boldsymbol{p^*}$. Rearranging this first-order condition leads to (1.9).

The first-order condition above yields a unique interior solution. It is necessary to rule out potential corner solutions. Suppose that there exists an equilibrium in which at least one seller j sets their price, $p_j = 0$, yielding a profit of 0. By assumption, $a > \bar{a}$ and hence $x_{ij} > 0$ for all i, j pair. There then exists a $\varepsilon > 0$ such that $p_j^{'} = \varepsilon$ generates a strictly positive level of demand $x_{ij}^{'} > 0$ for at least one buyer j. This would yield a strictly positive level of profit, which implies that j would have an incentive to deviate - yielding a contradiction.

Proof of Corollary 1.1

Note that $\frac{\frac{1}{2}\alpha_j}{\beta_j} = \frac{1}{2}\gamma$ in this setting and define $\boldsymbol{\gamma}$ as a $m \times 1$ vector whose jth element is equal to γ_j. Then the right-hand side of (1.10) can be rewritten as follows:

$$C_{\tilde{\alpha}}(G_S, \lambda) = \frac{1}{2}\sum_{k=0}^{\infty} \lambda^k R_S^k \boldsymbol{\gamma} - \sum_{k=1}^{\infty} \lambda^k R_{\tilde{S}}^k \boldsymbol{\gamma}$$

thus:

$$C_{\tilde{\alpha}}(G_S, \lambda) = \frac{1}{2}\boldsymbol{\gamma} + \frac{1}{2}\boldsymbol{\gamma}^{\boldsymbol{T}} \sum_{k=0}^{\infty} \lambda^k R_S^k - \sum_{k=0}^{\infty} \lambda^k R_S^k \boldsymbol{\gamma}$$

and:

$$C_{\tilde{\alpha}}(G_S,\lambda) = \frac{1}{2}\boldsymbol{\gamma} - \frac{1}{2}\boldsymbol{\gamma}^T \sum_{k=0}^{\infty} \lambda^k R_S^k,$$

which implies:

$$C_{\tilde{\alpha}}(G_S,\lambda) = \frac{1}{2}\boldsymbol{\gamma} - \frac{1}{2}\boldsymbol{\gamma^T} C(G_{\bar{S}},\lambda).$$

Proof of Proposition 1.1

Given the expression for the price vector, and noting that $\boldsymbol{p}^*(G) = \frac{1}{2}(\boldsymbol{\gamma} - \boldsymbol{\gamma^T} C(G_S,\lambda)) > 0$, the result immediately follows.

For the result in (ii), first recall that $\tilde{c}_{jk} = \frac{\sum_i cw_{ij}w_{ik}}{\beta_j}$. Define G_S^τ as a graph with (m,n) nodes with edges such that $\tilde{c}_{jk}^\tau = \frac{\sum_i c^{'} w_{ij}w_{ik}}{\beta_j}$. It follows that, as G_S (and therefore G_S^τ) is connected, the ijth entry of the centrality matrix $C(G_{S,}\lambda)$ is less than the ijth entry of $C(G_{S,}^\tau\lambda)$, which in turn implies that:

$$C(G_{S,}^\tau\lambda)\mathbf{1}_m > C(G_{S,}\lambda)\mathbf{1}_m$$

increasing in c. As $\boldsymbol{p}^*(G) = \frac{1}{2}(\boldsymbol{\gamma} - \boldsymbol{\gamma}' C(G_S,\lambda))$ the result follows.

Proof of Proposition 1.2

Recall that for the graph G_S, $\tilde{c}_{jk} = \frac{\sum_i cw_{ij}w_{ik}}{\beta_j}$. Let $\tilde{c}_{jk}^{'}$ denote this value in the graph $G_S^{'}$. $\beta_j^{'} > \beta_j$, as $b[\sum_l^n w_{lj} + w_{ij}] > b[\sum_l^n w_{lj}]$. As i is uniquely connected to j, $cw_{lj}^{'}w_{lk}^{'} = cw_{lj}w_{lk}$ for all l,j,k. This implies in turn that $\tilde{c}_{jk}^{'} < \tilde{c}_{jk}$ for all k and $\tilde{c}_{lk}^{'} < \tilde{c}_{lk}$ for all $l,k \neq j$.

The entries of the centrality matrix $C(G_{S,}\lambda)$ are equal to $\sigma_{ij} = \sum_l \lambda^l r_{ij}^{[l]}(G_S)$,where $r_{ij}^{[l]}(G_S)$ measures the weighted walks of length l that start at node i and end at node

j. As G_S is connected, it follows that:

$$C(G'_{S,}\lambda)\mathbf{1}_m < C(G_{S,}\lambda)\mathbf{1}_m,$$

which implies the result.

Proof of Proposition 1.3

Note that the entries of the centrality matrix $C(G_{S,}\lambda)$ are equal to $\sigma_{ij} = \sum_l \lambda^l r_{ij}^{[l]}(G_S)$,where $r_{ij}^{[l]}(G_S)$ measures the weighted walks of length l that start at node i and end at node j. Consider walks of length 1 in G and G''. If the graph is complete, then there exists at least one walk from j and a seller $k \in \boldsymbol{G''}$ that is not present in $\boldsymbol{G}$. By similar reasoning, there are also walks of length $n \in [2, \infty]$ in $\boldsymbol{G''}$that do not exist in $\boldsymbol{G}$. Every walk in $\boldsymbol{G}$ exists in $\boldsymbol{G''}$ and the weight on those edges in both graphs is the same. To see this, note that $\tilde{c}_{ij} = \frac{\hat{c}_{ij}}{\beta_i}$, which do not depend on the weight on other edges in the graph.

Hence the weighted sum of walks from all nodes is greater in $\boldsymbol{G}''$ than $\boldsymbol{G}$, which implies that:

$$C(G_S, \lambda)\boldsymbol{\gamma} < C(G''_S, \lambda)\boldsymbol{\gamma}.$$

Given that $\boldsymbol{p}^* = (\boldsymbol{\gamma_j} - \frac{1}{2}\boldsymbol{\gamma}'_{\boldsymbol{j}} C(G_S, \lambda))$, the result is proved.

Proof of Corollary 1.2

Call the proportion of buyers who have n connections and are connected to seller i, $\sigma_{n,i}$. Then the strength of the connection between two seller, i and j, can be written as follows:

$$\tilde{c}_{ij} = \frac{c\sigma_i\sigma_j}{an\sigma_i} = \frac{c\sigma_j}{an}.$$

It follows that if $\sigma_i > \sigma_j$, then $\tilde{c}_{ik} < \tilde{c}_{jk}$ for all k, including $k = i$. Hence, the centrality of i in the network G_S is lower than the centrality of j, which implies that $p_i^* > p_j^*$ by the result of Corollary 1.1.

Proof of Corollary 1.3

This corollary follows immediately from Proposition 1.3.

Proof of Lemma 1.1

By the envelope theorem, taking the derivative of seller j's maximisation problem gives the following condition:

$$\frac{\partial \pi_j(\boldsymbol{p^*})}{\partial w_{ij}} = ap_j^*(\gamma - p_j^*) + (cp_j^* \sum_{k=1}^{m} r_k w_{ik}(p_k^* - \gamma) + \sum_{k=1}^{m} [\hat{c}_{jk} \frac{\partial p_k^*}{\partial w_{ij}}] p_j^*) = 0 \qquad (1.17)$$

where $r_k = 1$ when $k \neq j$ and $r_k = 0$ when $k = j$. To guarantee that a solution to this problem exists and is a maximum, we must analyse $\frac{\partial p_k^*}{\partial w_{ij}}$ and show that its derivative is negative.

Recall that $C(G_S, \lambda)\mathbf{1} = \sum_{k=0}^{\infty} \lambda^k G_S^k \mathbf{1}$. It follows that C_k can be written as some additive function of each of the powers w_{ij} from 1 to ∞:

$$C_k(w_{ij}) = g_k(w_{ij}, w_{ij}^2, ..., w_{ij}^{\infty}) = 1 + \sum_{n=1}^{\infty} \varphi_n(\boldsymbol{R}) w_{ij}^n$$

where $0 < \varphi_n < 1 \, \forall n$. It is possible to see that all the criteria of Abel's uniform convergence test are satisfied such that $g_k(w_{ij}, w_{ij}^2, ..., w_{ij}^{\infty})$ is uniformly convergent. Define f_n as the nth term in the infinite sum that chracterises $C_k(w_{ij})$,

$$J(w_{ij}) = \sum_n f_n'(w_{ij})$$

and:

$$K(w_{ij}) = \sum_n f_n''(w_{ij}).$$

It is clear that $J(w_{ij})$ and $K(w_{ij})$ converge uniformly, which, as Jeffreys and Jeffreys (1988) show, is sufficient to guarantee that $\frac{\partial C_k(G_{S,}\lambda)}{\partial w_{ij}} > 0$ as it guarantees that the derivative of $g_k(w_{ij}, w_{ij}^2, ..., w_{ij}^\infty)$ is the sum of the derivatives of each component of the sum which composes it. Furthermore, it is clear by the same logic that $\frac{\partial^2 C_k(G_{S,}\lambda)}{\partial^2 w_{ij}} > 0$. Noting that p_k is negative in $C_k(G_{S,}\lambda)$, the seller's maximisation problem is strictly concave in w_{ij}.

The fact that $\frac{\partial^2 C_j(G_{S,}\lambda)}{\partial^2 w_{ij}} < 0$ is sufficient to guarantee that there exists a $w^* \in [0, 1]$ that is a unique global maximum. Namely, w^* solves the first-order condition in (1.17).

Proof of Proposition 1.4

The first-order condition of the platform owner's maximisation problem can be written:

$$\frac{\partial \pi_P(\boldsymbol{p}^*)}{\partial w_{ij}} = \chi \sum_{k=1}^{m} \frac{\partial \pi_k(\boldsymbol{p^*})}{\partial w_{ij}} = 0.$$

The derivative of j's profit with respect to w_{ij} is expressed in (1.17) above. The derivative of another seller, k's profit function can be similarly expressed:

$$\frac{\partial \pi_k^*}{\partial w_{ij}} = p_k^*(\sum_{k=1}^{m} r_k c w_{ik} p_j^* + \sum_{l=1}^{m} \hat{c}_{kl} \frac{\partial p_l^*}{\partial w_{ij}}), \tag{1.18}$$

As shown in proof of Lemma 1.1, $\frac{\partial^2 p_l^*}{\partial^2 w_{ij}} < 0$, and hence the platform owner's maximisation problem has a unique solution w_{ij}^P that solves the above first-order condition:

$$\frac{\partial \pi_P^*}{\partial w_{ij}} = \chi \sum_{k=1}^{m} \frac{\partial \pi_k^*}{\partial w_{ij}} = 0. \tag{1.19}$$

The above implies that the solution to the two sets of first-order conditions are maxima. Now note that the platform's first-order condition can be as follows:

$$\frac{\partial \pi_j^*}{\partial w_{ij}} + cp_j^* \sum_{k=1}^{m} w_{ik}(p_k^* - \gamma) + \sum_{l=1}^{m} \sum_{k=1}^{m} [\hat{c}_{kl} \frac{\partial p_k^*}{\partial w_{ij}}] p_l^*) = 0. \tag{1.20}$$

The second term in this expression is negative, as $p_k^* < \gamma$. As the second derivative of π_j^* in w_{ij} is also negative, it follows that the unconstrained value of w_{ij} that solves the above must be smaller than the unconstrained value of w_{ij} that solves i's optimisation problem. Hence, $w_{ij}^P \leq w_{ij}^*$, with the inequality strict if $w_{ij}^* < 1$.

Proof of Theorem 1.2

By the envelope theorem, the marginal benefit of increasing w_{ij} can be expressed:

$$\text{MB}_{ij} = ap_j(\tilde{w})(\gamma - p_j(\tilde{w}) + c \sum_{k=1}^{m} p_j(\tilde{w}) w_{ik}(p_k - \gamma).$$

and the marginal cost of increasing w_{ij} can be expressed:

$$\text{MC}_{ij} = c \sum_{k=1}^{m} w_{ik}(p_j(\tilde{w})p_k - \gamma) + \sum_{l=1}^{m} \sum_{k=1}^{m} [\hat{c}_{kl} \frac{\partial p_k^*}{\partial w_{ij}}] p_l^*.$$

As $p_k = \gamma(1 - \frac{1}{2}C_k(G_S, \lambda))$ and hence $p_k < \gamma$, it follows that if $\sum_k w_{ik} > \sum_k w_{lk}$ and $\sum_k w_{ik} C_k(G_S, \lambda) > \sum_k w_{lk} C_k(G_S, \lambda)$, which in turn implies that the last expression in the marginal benefit expression and the first in the marginal cost expression.then terms are larger for l than for i.

Consider the last term in the marginal cost expression. Noting that:

$$C_k(G_S,\lambda) = \sum_{l=1}^{m} \tilde{c}_{kl} C_l(G_S,\lambda).$$

then $C_k(G_S,\lambda)$ can be written:

$$C_k(G_S,\lambda) = \bar{c}_{kj} C_j(G_S,\lambda).$$

where $\bar{c}_{kj} > 0$ is a parameter of the form:

$$\bar{c}_{kj} = \frac{z_1(\tilde{c}_{1j},...,\tilde{c}_{mj})}{z_2(\tilde{c}_{1k,...},\tilde{c}_{mk})}$$

where $z_1(.) : \mathbb{R}^m \to \mathbb{R}_+$ is a linear and increasing function and $z_2(.) : \mathbb{R}^m \to \mathbb{R}_+$ is a linear and decreasing function in all of its inputs except $\tilde{c}_{jk}$, as $\frac{\partial z_2(\tilde{c}_{1k,...},\tilde{c}_{mk})}{\partial \tilde{c}_{jk}} = 0$ by construction. Hence, the derivative of $C_k(G_S,\lambda)$ with respect to w_{ij} when $w = \tilde{w}$ can be written:

$$\frac{\partial C_k(G_S,\lambda)}{\partial w_{ij}} = \frac{\partial \bar{c}_{kj}}{\partial w_{ij}} C_j(\tilde{w}) + \bar{c}_{kj}(\frac{\partial C_j(G_S,\lambda)}{\partial w_{ij}}). \tag{1.21}$$

Noting that:

$$\frac{\partial C_j(G_S,\lambda)}{\partial w_{ij}} = \sum \frac{\partial \tilde{c}_{jl}}{\partial w_{ij}} C_l + \sum \tilde{c}_{jl}(\frac{\partial C_l(G_S,\lambda)}{\partial w_{ij}}),$$

Equation (1.21) implies:

$$\frac{\partial C_j(G_S,\lambda)}{\partial w_{ij}} = \sum_{l}^{m} \frac{\partial \tilde{c}_{jl}}{\partial w_{ij}} C_l + \sum_{l}^{m} \tilde{c}_{jl}(\frac{\partial \bar{c}_{lj}}{\partial w_{ij}} C_j(\tilde{w}) + \bar{c}_{lj}(\frac{\partial C_j(G_S,\lambda)}{\partial w_{ij}})). \tag{1.22}$$

Given that:

$$\frac{\partial \tilde{c}_{kj}}{\partial w_{ij}} = \frac{\frac{c}{a} w_{ik}}{\sum_i w_{ik}},$$

it is is clear that $\frac{\partial \tilde{c}_{kj}}{\partial w_{ij}}$ is increasing in $\sum_k w_{ik}$. Furthermore:

$$\frac{\partial \tilde{c}_{jk}}{\partial w_{ij}}|_{\boldsymbol{w_j}=\tilde{w}\boldsymbol{1_n}} = \frac{\frac{c}{a}w_{ik}(n\tilde{w}) - \sum_l w_{lk}\tilde{w}}{(n\tilde{w})^2} = \frac{\frac{c}{a}w_{ik}(n-1) - \sum_l (w_{lk} - w_{ik})}{\tilde{w}n^2}.$$

Hence, $\frac{\partial \tilde{c}_{jk}}{\partial w_{ij}}|_{\boldsymbol{w_j}=\tilde{w}\boldsymbol{1_n}}$ is increasing in w_{ik}, and so the first term in (1.22) is an increasing function of $\sum_k w_{ik}C_k$.

Combining the above observations together, it follows that $\mathrm{MB}_{ij} + \mathrm{MC}_{ij} \geq \mathrm{MB}_{lj} + \mathrm{MC}_{lj}$ if it is the case that both $C_i^B(G) < C_l^B(G)$ and $\sum_k w_{ik}(G) \leq \sum_k w_{lk}(G)$. Hence, in any solution to the platform owner's maximisation problem, it must be the case that $w_{ij}^P \geq w_{lj}^P$ for any solution to the platform's optimisation problem.

Proof of Corollary 1.4

The result follows immediately from Theorem 1.2: it is in fact just a restatement of it for the case where $\sum_k w_{ik}(G) = \sum_k w_{lk}(G)$.

Proof of Proposition 1.5

Recalling that $C(G_S, \lambda)\mathbf{1} = \sum_{k=0}^{\infty} \lambda^k R_S^k \mathbf{1}$ 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 $\boldsymbol{p}$ for a given m. At the same time, given the assumption that $x_{ij} > 0$, for each consumer i and seller j in a complete network, which implies that, holding price constant, i's expected consumer surplus is increasing in each weight w_{ij}. Hence, consumer surplus is maximises when the network is complete.

Proof of Proposition 1.6

When $c = 0$, the marginal cost associated with increasing w_{ij} is zero. Furthermore, the marginal benefit of increasing w_{ij} can be written as follows:

$$\text{MB}_{ij} = a\gamma p_j^* - b(p_j^*)^2.$$

As this expression is greater than zero, it follows that the optimal graph is a corner solution such that $w_{ij} = 1$ for each i, j pair, which is the complete graph.

Proof of Proposition 1.7

In the case where each observation probability is equal to υ, then a seller j's (unweighted) centrality in G_S is:

$$C_j(G_S, \lambda) = \sum_{n=0}^{\infty}(\frac{c}{b}\upsilon(m-1))^n.$$

As $j's$ price is strictly decreasing in their centrality for all j, it follows that $\frac{\partial^2 p_j}{\partial^2 \upsilon} < 0$. This implies the platform owner's maximisation problem has a unique global maximum that solves the first-order condition:

$$\frac{\partial \pi^P(\upsilon^*; \boldsymbol{p^*})}{\partial \upsilon} = 0.$$

By the envelope theorem, we can express the marginal effect of a change in υ as follows:

$$\frac{\partial \pi^P(\upsilon^*; \boldsymbol{p^*})}{\partial \upsilon} = n[map^*(\gamma - p^*) + 2c(m-1)^2 p^* \upsilon^*(p^* - \gamma) + c(\upsilon^*)^2(m-1)\sum_{l=1}^{m}\sum_{k=1}^{m}[\frac{\partial p^*}{\partial \upsilon}]p^*].$$

Hence, if c increases, then υ^* (weakly) decreases in order for the first-order condition above to hold. If c is sufficiently large, the the υ^* that solves the first-order condition is such that $\upsilon^* < 1$.

Denote $\pi_\upsilon^P = \frac{\partial \pi^P(\upsilon^*;\boldsymbol{p^*})}{\partial \upsilon}$. By the implicit function theorem:

$$\frac{\partial \upsilon^*}{\partial a} = -\frac{\pi_{\upsilon a}^P}{\pi_{\upsilon\upsilon}^P} > 0,$$

as $\pi_{\upsilon\upsilon}^P < 0$. As $\pi_{\upsilon c}^P < 0$, then by a similar logic $\frac{\partial \upsilon^*}{\partial c} < 0$.

Proof of Theorem 1.3

Suppose $\boldsymbol{\rho}^* = \frac{1}{m}\mathbf{1}_m$ and consider the case where ρ_j is marginally increased by reducing ρ_k. The expected equilibrium price vector of the game Γc can be written as follows:

$$\mathrm{E}[\boldsymbol{p^*}(G)] = \frac{1}{2}(\boldsymbol{\gamma} - \boldsymbol{\gamma}^T \mathrm{E}[C(G_S, \lambda)|\boldsymbol{\rho = \rho^*}]).$$

Noting that the expected price for each seller is equal in the case where $\boldsymbol{\rho} = \boldsymbol{\rho^*}$, let $p^* = \mathrm{E}[p_j^*(G)|\boldsymbol{\rho} = \boldsymbol{\rho}^*]$ for all j. By the envelope theorem, we can express the gross expected marginal benefit of increasing ρ_j in a similar way to above

$$\mathrm{E}[\mathrm{MB}_j] = n[a\xi p^*(\gamma - p^*) + (m-1)\xi c\upsilon p^*(p^* - \gamma)].$$

However, increasing ρ_j reduces ρ_k by the same amount and, as $\mathrm{MB}_j = \mathrm{MB}_k$, it follows that there is no change in the net marginal benefit associated with the proposed change in observation probabilities. There is a change in the net marginal cost, which can be expressed as follows:

$$\mathrm{E}[\mathrm{MC}_j - \mathrm{MC}_k] = n(m-1)\tilde{c}\sum_{i=1}^{m}(\frac{\partial \mathrm{E}[p_i^*]}{\partial \rho_j} - \frac{\partial \mathrm{E}[p_i^*]}{\partial \rho_k})p^*, \tag{1.23}$$

where $\tilde{c}$ is the expectation of $\tilde{c}_{il}$ for all i, l pairs when $\boldsymbol{\rho}^* = \frac{1}{m}\mathbf{1}_m$.

Consider an undirected graph, $\hat{G}_S$, where an edge between j and k is equal to $\hat{c}_{jk} = cw_{ij}w_{ik}$, recalling that the directed edge between j and k in G_S is equal to

$\tilde{c}_{jk} = \frac{\hat{c}_{jk}}{\beta_j}$. Also define $\hat{\lambda} = \max\{c_{ij}\}$ and $C(\hat{G}_S, \hat{\lambda})$ as a matrix whose jkth component is $\frac{\hat{c}_{jk}}{\hat{\lambda}}$. Note that the component of the centrality matrix $C(\hat{G}_S, \hat{\lambda})$ is equal to $\sigma_{ij} = \sum_l \hat{\lambda} r_{ij}^{[l]}$, where $r_{ij}^{[l]}$ is the weighted walks of length l that start at node i and end at node j. We can express $\hat{c}_{ij} = \hat{c}_{ji}$ as follows:

$$\mathrm{E}[\hat{c}_{ij}] = cn[\rho_i(\upsilon + \xi)\upsilon + \rho_j(\upsilon + \xi)\upsilon + (1 - \rho_i - \rho_j)\upsilon^2].$$

Compare the expected sum of all the walks of length 1 in $\hat{G}_S$ when $\boldsymbol{\rho} = \boldsymbol{\rho}^*$ and when $\boldsymbol{\rho} = \boldsymbol{\rho}'$, which is identical to $\boldsymbol{\rho}^*$ except that $\rho_j' = \rho_j^* + \epsilon$ and $\rho_k' = p_k^* - \varepsilon$. Given the expression for the expected value of $\mathrm{E}[\hat{c}_{ij}]$:

$$\sum_i \sum_j \mathrm{E}[\hat{c}_{ij} | \boldsymbol{\rho} = \boldsymbol{\rho}^*] = \sum_i \sum_j \mathrm{E}[\hat{c}_{ij} | \boldsymbol{\rho} = \boldsymbol{\rho}'].$$

Assume that $m \geq 3$. Consider the sum of the walks of length 2 in $\hat{G}_S$ from a seller $i \neq j, k$ which can be expressed: $\sum_l \sum_s c_{il} c_{ls}$. Again, compare this sum when $\boldsymbol{\rho} = \boldsymbol{\rho}^*$ and when $\boldsymbol{\rho} = \boldsymbol{\rho}'$. The following observations apply:

- any walk from i that does not run through j and k has the same value for both probability vectors;

- excluding walks that end at i, the difference between sum of walks though j but not k when $\boldsymbol{\rho} = \boldsymbol{\rho}'$ and when $\boldsymbol{\rho} = \boldsymbol{\rho}^*$ is equal to the absolute value of the difference between sum of walks though k but not j when $\boldsymbol{\rho} = \boldsymbol{\rho}'$ and when $\boldsymbol{\rho} = \boldsymbol{\rho}^*$;

- the increase in the value of the walk that runs from i through j and back to i is larger than the decrease in the value of the walk that runs from i through k and back to i for all i.

The last result holds because $\mathrm{E}[\hat{c}^2_{ij}(\boldsymbol{\rho}^*) - \hat{c}^2_{ij}(\boldsymbol{\rho}')] := \Delta\hat{c}^2_{ij} > |\Delta\hat{c}^2_{ik}|$ when $\upsilon > 0$. This inequality implies that the increase in the value of the sum of walks of length 2 starting at j is larger than the decrease in the sum of walks of length 2 starting at k. Hence, the sum of all the walks of length 2 starting at any $i \neq k$ is greater when $\boldsymbol{\rho} = \boldsymbol{\rho}'$ than when $\boldsymbol{\rho} = \boldsymbol{\rho}^*$. The same logic applies for walks of length $l > 2$, as $\frac{\partial^l C_j(\hat{G}_S, \hat{\lambda})}{\partial^l \hat{c}_{jk}} > 0$ for all l. It follows that:

$$\mathrm{E}[C_i(\hat{G}_S, \hat{\lambda})|\boldsymbol{\rho} = \boldsymbol{\rho}^*] < \mathrm{E}[C_i(\hat{G}_S, \hat{\lambda})|\boldsymbol{\rho} = \boldsymbol{\rho}'] \; i \neq k.$$

The effect on the centrality of k is ambiguous: the sum of shorter walks starting at k are smaller when $\boldsymbol{\rho} = \boldsymbol{\rho}'$ than when $\boldsymbol{\rho} = \boldsymbol{\rho}^*$, but the sum of longer walks are potentially larger. However, as $\Delta\hat{c}^n_{ij} > |\Delta\hat{c}^n_{ik}|$ for all $n \geq 2$, the decrease in k's centrality is less than the corresponding increase in j's centrality.

Furthermore, as the above argument does not rely on the value of $\boldsymbol{\rho}$, and therefore the above result applies to any increase in some probability ρ_j for any corresponding decrease in ρ_k where $\rho_j \geq \rho_k$. Hence, increasing the probability that a seller is prominent increases the sum of the sellers' centralities in the undirected graph $\hat{G}_S$, which can be thought of as capturing an increase in the concentration of competition on the network.

Now consider the directed graph G_S. Recall that $\tilde{c}_{ji} = \frac{\hat{c}_{ji}}{\beta_j}$ and thus $\tilde{c}_{ji}$ is convex in β_j for all $\underline{i}$ and j. Consider the sum of all walks of length 1, $\lambda r^{[1]}_{ij}$.

$$\lambda r^{[1]}_{ij} = \sum_{i=1} \sum_{j=1} \tilde{c}_{ij}.$$

Note that the $\sum \beta_i = \bar{\beta}$ for any realisation of the adjacency matrix $\boldsymbol{R}$. Hence, if $\upsilon > 0$ and $m \geq 2$, increasing ρ_j above $\frac{1}{m}$ shifts the probability mass realisations of $\boldsymbol{R}$ in which β_j is relatively large, which, given each $\tilde{c}_{ji}$ is convex, implies the following:

$$\mathrm{E}[\lambda r_{ij}^{[1]}|\boldsymbol{\rho} = \boldsymbol{\rho}^{*}] < \mathrm{E}[\lambda r_{ij}^{[1]}|\boldsymbol{\rho} = \boldsymbol{\rho}^{'}]\ \forall \boldsymbol{\rho}^{'} \neq \boldsymbol{\rho}^{*}.$$

walks of length 2 are all of the form: $\frac{\hat{c}_{ji}}{\beta_j}\frac{\hat{c}_{ij}}{\beta_i}$. Hence, each of the walk of length 2 starting at j is also convex in β_j. In fact, as walks starting at j of walk length $k > 2$ are functions of $\frac{1}{\beta_j^r}$ where $r \geq 2$, it follows that each walk of length k is also convex in β_j.

The preceding analysis indicates that increasing ρ_j above $\frac{1}{m}$ increases the probability of realisations of $\boldsymbol{R}$ in which both:

1. The sum of the unweighted walks $\sum_i \sum_j \hat{c}_{ij}^r$ are larger than $\mathrm{E}[\sum_i \sum_j \hat{c}_{ij}^r|\boldsymbol{\rho} = \boldsymbol{\rho}^{*}]$ for all $r \geq 2$ (assuming $m > 2$);
2. $\beta_j > \beta_k$ $k \neq j$ (an effect which occurs for all $m \geq 2$).

Hence, increasing ρ_j above $\frac{1}{m}$ shifts the probability mass towards outcomes in which $\sum_i \sum_j \lambda r_{ij}^{[q]}$ is larger than $\mathrm{E}[\sum_i \sum_j \lambda r_{ij}^{[q]}|\boldsymbol{\rho} = \boldsymbol{\rho}^{*}]$ for all $q \geq 1$. This implies that for all $q \in [1, \infty]$.:

$$\mathrm{E}[\sum_i \sum_j \lambda^{[q]} r_{ij}^{[q]}|\boldsymbol{\rho} = \boldsymbol{\rho}^{*}] < \mathrm{E}[\sum_i \sum_j \lambda^{[q]} r_{ij}^{[q]}|\boldsymbol{\rho} = \boldsymbol{\rho}^{'}]\ \boldsymbol{\rho} \neq \boldsymbol{\rho}^{*}.$$

The above expression in turn implies that:

$$\mathrm{E}[\sum_{i=1}^{m} C_i(\hat{G}_S, \hat{\lambda})|\boldsymbol{\rho} = \boldsymbol{\rho}^{*}] < \mathrm{E}[\sum_{i=1}^{m} \sum_{i=1}^{m} C_i(\hat{G}_S, \hat{\lambda})|\boldsymbol{\rho} = \boldsymbol{\rho}^{'}].$$

It follows that the sum of the price changes as a result of a change from $\boldsymbol{\rho}^{*}$ to any probability vector $\boldsymbol{\rho}^{'} \neq \boldsymbol{\rho}^{*}$ is negative. Note that, as argued above, the solution to the platform owner's first-order conditions must solve equation (23). Given that the sum of the price changes as a result of a change from $\boldsymbol{\rho}^{*}$ to any probability vector $\boldsymbol{\rho}^{'} \neq \boldsymbol{\rho}^{*}$ is negative, it follows that the the only solution to the first-order condition in (1.23) is $\boldsymbol{\rho} = \boldsymbol{\rho}^{*}$.

Proof of Proposition 1.8

Without loss of generality, let seller 1 be a seller with intercept parameter $\gamma_1 = \gamma$ and let $\rho_1 = 1 - \sum_{k=2}^{m} \rho_k$. For any solution $\boldsymbol{\rho}^{**}$ to the platform owner's maximisation problem that is interior, the set of first-order conditions for the platform owner's problem is:

$$\frac{\partial \pi^P(\boldsymbol{\rho}; \boldsymbol{p}^*)}{\partial \rho_i} = 0 \; \forall i.$$

Define $p_j^* := \mathrm{E}[p_j^*(G)|\boldsymbol{\rho}]$. By the envelope theorem, we express the gross expected marginal benefit of increasing ρ_j as follows:

$$\mathrm{E}[\mathrm{MB}_j|\boldsymbol{\rho} = \boldsymbol{\rho}^*] = n[a\xi p_j^*(\gamma_j - p_j^*) + \sum_{k=1}^{m} \frac{\partial \mathrm{E}[\hat{c}_{jk}]}{\partial \rho_j} p_j^*(p_k^* - \gamma)] > \mathrm{E}[\mathrm{MB}_k|\boldsymbol{\rho} = \boldsymbol{\rho}^*] \forall k.$$

As $\gamma_j > \gamma$ the net expected marginal benefit of increasing ρ_j, $\mathrm{E}[\mathrm{MB}_j|\boldsymbol{\rho} = \boldsymbol{\rho}^*] - \mathrm{E}[\mathrm{MB}_1|\boldsymbol{\rho} = \boldsymbol{\rho}^*] > 0$. Furthermore, note that $\mathrm{E}[\mathrm{MB}_j|\boldsymbol{\rho} = \boldsymbol{\rho}^*] - \mathrm{E}[\mathrm{MB}_1|\boldsymbol{\rho} = \boldsymbol{\rho}^*]$ is linearly increasing in a.

However, it is also the case that $\mathrm{E}[\mathrm{MC}_j|\boldsymbol{\rho} = \boldsymbol{\rho}^*] - \mathrm{E}[\mathrm{MC}_1|\boldsymbol{\rho} = \boldsymbol{\rho}^*] < 0$, which follows from the fact that:

$$\mathrm{E}[\mathrm{MC}_j - \mathrm{MC}_1] = -cn(m-1)(\gamma_j - \gamma) + n(m-1)\tilde{c} \sum_{i=1}^{m} (\frac{\partial \mathrm{E}[p_i^*]}{\partial \rho_j} - \frac{\partial \mathrm{E}[p_i^*]}{\partial \rho_1}) p^*.$$

The first term of this expression is trivially less than zero as $(\gamma_j - \gamma) > 0$. The second term is also negative because a marginal increase in ρ_j increases $\mathrm{E}[\sum_k \hat{c}_{jk}\gamma_j]$ more than a marginal decrease in ρ_1 decreases $\mathrm{E}[\sum_k \hat{c}_{1k}\gamma_1]$ because, again, $(\gamma_j - \gamma) > 0$. This implies that:

$$\frac{\partial \mathrm{E}[C(G_S, \lambda]}{\partial \rho_j} > \frac{\partial \mathrm{E}[C(G_S, \lambda]}{\partial \rho_1},$$

which implies that $\frac{\partial \mathrm{E}[p_i^*]}{\partial \rho_j} - \frac{\partial \mathrm{E}[p_i^*]}{\partial \rho_1} < 0$ for all i. However, note that the expression in (24) is not a function of a. As $\mathrm{E}[\mathrm{MB}_j|\boldsymbol{\rho} = \boldsymbol{\rho}^*] - \mathrm{E}[\mathrm{MB}_1|\boldsymbol{\rho} = \boldsymbol{\rho}^*]$ is increasing in a, there exists an $\tilde{a}$ such that if $a > \tilde{a}$, $\boldsymbol{\rho}^*$ cannot be a solution to the platform owner's maximisation problem.

Chapter 2

Rating the Competition: Seller Ratings and Intra-Platform Competition

2.1 Introduction

Many online platforms allow users to rate products on their platform. These ratings have many potential uses, including revealing to buyers their own valuation for products on offer. Another way in which product ratings are used is to determine the likelihood that sellers are observed by buyers. Given that buyers do not tend to observe every seller on a platform (see, for example, Ringer and Skiera, 2016 and Kim, Albuquerque and Bronnenberg, 2010), the choice of which sellers buyers observe will affect platform profits. Many online platforms use ratings as one of the inputs that determine how search results are ordered (Dinerstein et al, 2018). We analyse the incentive of a profit-maximising, monopolistic platform owner to use ratings as a means of determining the prominence of a seller on a platform.

Analysis of real-life online platforms reveals that there are differences in how important ratings are in default search settings. Platforms like Amazon, Airbnb, Jet.com, Wish.com, and Trivago appear to use these ratings as part of search algorithms, but are not used to directly rank products as default. Other platforms, like eBay, Etsy, Newegg and Zibbet do not to utilise ratings as part of their default search options (see Dinerstein et al, 2018 for a detailed account of eBay's change of search algorithm). Our analysis examines what drives the relative importance of ratings in the search algorithms of these online platforms.

To give an example of the problem under consideration, suppose there is a market with three products, X, $Y$and Z, where X is more highly-rated than Y and Y is more highly-rated than Z. A platform owner chooses two aspects of the search process in order to maximise joint network profits: (a) the baseline probability that any one product is observed and (b) the relative prominence of highly-rated products. For example, the platform owner can choose both the number of search results displayed for a given query and whether those search results are ordered according to ratings (i.e. X, Y, Z) or randomly (e.g. Y, Z, X).

We find that there is a trade-off inherent with choosing to positively weight the observation probabilities according to a product's rating. On the one hand, there is an incentive to bias the search process towards highly-rated sellers because doing so increases the probability that buyers observe high quality sellers. Buyers also demand more of the products sold by highly-rated sellers for a given set of prices, which implies that matching buyers to these sellers increases expected profits, if prices were independent of the search process chosen by the platform owner.

However, making it more likely that highly-rated sellers are observed causes competition to become more concentrated across the network. highly-rated sellers are more likely to compete with each other than in the case where matches are random. This fierce competition between a subset of the sellers in the network more than offsets the fact that sellers with a low rating are less likely to compete with one another. highly-rated sellers have a lower price than is the case were matching random, and all sellers on platform, including those with low ratings, respond by reducing their prices as well.

The optimal price-setting behaviour of sellers for a given network when sellers compete à la differentiated Bertrand is examined in Chapter 1. There, we characterise optimal price setting in terms of the sellers' Bonacich centrality in a secondary network in which the strength of connection between sellers is determined by the probability that they will competing for the same buyers once buyers and sellers are connected.

We use this differentiated Bertrand competition in a network framework to look at a case in which the preferences of buyers are stochastic with respect to the intercept term of their demand functions, capturing randomness in the valuation a set of buyers has for a seller's product. Each intercept term is assume to be composed of a "quality" component, which is reflected in average ratings, and an idiosyncratic component, which is not.

We examine the platform owner's optimal choice of information environment in such a setting, which is assumed to take place prior to the realisation of the consumers' preferences and therefore any rankings of the products on the platform. We also analyse how changes in parameters in the model affect that optimal choice and how the relative profitability of a purely ranking-based information environment compares to that as the variances of the quality and idiosyncratic components of the valuation parameter changes.

If products become less substituable, the concentration of competition effect reduces, increasing the profitability of information environments where buyers: (a) are more likely to observe sellers in general and (b) are in particular more likely to observe highly-rated products. Similarly, in a market where buyers are particularly sensitive to quality, there is a greater incentive to ensure that they are matched with high quality products. In such markets, the platform owner will increase the extent to which search is biased towards high quality products.

We examine the profitability of an environment where search results are biased in this way and compare it to the case where matches are random. The latter environment yields more profit when buyers are not particularly sensitive to quality, as the cost imposed by more concentrated competition outweighs any benefit associated with high quality matches.

Furthermore, as the variance of the component of the buyers' valuation that is idiosyncratic increases, the benefit of random match increases. Using ratings to bias search increases the probability that buyers miss out on high quality matches with a surplus largely composed of the buyer's idiosyncratic preference towards the product.

If buyer demand is sufficiently sensitive to quality, these effects are dominated by the benefits associated with matching buyers to high quality sellers. Our analysis therefore suggests that the nature of the market or markets that the platform hosts is crucial for determining the extent to which ratings are used to rank products. In markets where quality is a large driver in consumer demand, such as the hotel industry, ratings-based search increases profits. However, in markets driven largely by idiosyncratic demand, such as arts and crafts, random match may be preferred.

Consumer surplus is increasing in the probability that a consumer segment observes a seller and is decreasing in prices. As in Chapter 1, a benevolent central planner would set the baseline probability that buyers observe sellers to 1, so every consumer segment would observe every seller, generating the lowest possible price vector.

We also characterise the conditions under which the central planner would prefer the ratings-based environment over random matching. When the total probability of observation between the two environments is the same, the ratings-based environment is preferred because in expectation it creates higher-surplus matches and lower prices. When the platform owner does not use product ratings to bias search results, therefore, they likely do so at the expense of consumer welfare.

2.2 Literature review

Our modelling approach sits within the competition on networks framework, treatments of which include Ballester, Calvò-Armengol and Zenou (2006) and Bramoullé, Kranton and D'Amours (2014). Papers in this framework show the existence and nature of equilibrium actions in a network context when players actions directly affect the payoffs of players they are connected to. The games on network approach has been applied to competition in markets in, for example, Elliott and Galleotti (2019) and Bimpikis, Ehsani and Ilkiliç (2018).

There are a number of empirical studies that examine the effect of information design on the structure of competition in networks. The most prevalent of these simply attempt to capture the effect of product or seller ratings. For example, Chevalier and Mayzlin (2006), Vana and Lambrecht (2020) and Fradkin, Grewal and Holtz (2017) all examine the effect of ratings on demand on a range of online platforms, finding that ratings have a positive effect on demand, which validates our assumption that ratings are a means of identifying high-quality or highly demand products.

A more pertinent empirical study to this paper is Dinerstein et al (2018). The authors estimate a model in which platform owners choose the probability that buyers observe sellers on the basis of their price and quality. They assume that platform owners aim to maximise consumer surplus, and find that increasing the probability that low-price sellers are observed increases competition but reduces the probability that consumers do not observe high quality products. This trade-off is central to the model the authors use to estimate the effect of a change in eBay's search algorithm. They find that a reduction in posted prices is the result of observation probabilities being driven by seller price.

The search on platforms literature considers the effect the determinants of search have on the nature of competition. Choi, Dai and Kim (2018) analyse the case in which oligopolists compete on prices when prices partly determine the probability of being observed. In this context, a fall in search costs leads to a rise in prices because sellers are less concerned about not being observed, since their price is too high. Armstrong, Vickers and Zhou (2009) and Armstrong and Zhou (2011) show that search prominence affects competition and prices, with more prominent sellers setting lower prices and less prominent sellers setting higher prices, the overall effect being higher profits.

Our result that there are potential ex-ante gains for the platform owner generated by consumers observing high quality products shares features with the results found in work analysing the incentive of monopolists to reveal information relating to consumer valuations. Lewis and Sappington (1994) and Myatt and Johnson (2006) both show that a monopolist either has an incentive to full reveal information or reveal no information at all when the consumer is uninformed about their own preferences.

Recently, there been some analysis of the more specific problem of the incentive a platform owner has to reveal information to sellers and buyers when both are initially uninformed about buyer preferences. Armstrong and Zhou (2020) show that too much information revelation is potentially profit reducing ex ante, because it results in cases where, ex post, sellers compete aggressively with one another because they are not sufficiently differentiated. Elliott and Galleotti (2020) examine a case where the platform owner knows consumer preferences and chooses what information to pass on to firms, allowing the authors to characterise the conditions under which the platform owner can generate perfect segmentation.

The last two papers discussed fit into the wider Bayesian persuasion framework, which was introduced in Kamenica and Gentzkow (2011). The Bayesian persuasion literature examines how information revelation by one agent can be used to shape other agent's or agents' equilibrium behaviour in order to maximise some payoff. Applications of this framework to an industrial organisation context include Bergemann, Brooks, and Morris (2015), which examines third-degree price discrimination in the case where the firm is uninformed but consumers are not and Chen and Zhang (2020). While the basic nature of our analysis shares similarities with this class of problems, we examine the case where an uninformed agent is able to choose, ex-ante, how information can be used to shape the nature of the interaction between two, at least partially, informed agents.

2.3 The model

Suppose that there are a finite number of consumer segments, each composed of potentially a large but finite number of buyers. Each buyer within a consumer segment shares some trait, such as geographical location, age demographic, occupation etc. Call the set of these consumer segments B and suppose that $|B| = n$. Let S be a finite set of sellers and where $|S| = m$. 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 segments interact on a platform. A consumer segment, i, observes a seller j with probability w_{ij}, which is generated by an observation function discussed below. The observation process generates a weighted network $G = (B \cup S, E)$, where E is the set of weighted edges from S to B, where each consumer segment can be thought of as a node in the graph G. G has an adjacency matrix, $\boldsymbol{R}$, which is a zero-diagonal matrix with components w_{ij}.

We will assume that the following ex-post demand function captures the demand of a group of buyers i for a product j[1]:

$$x_{ij}(\boldsymbol{p}) = b(a - p_j) + \theta_i \gamma_{ij} + \sum_{k=1}^{m} c(p_k - a). \tag{2.1}$$

where $a, b, c \in \mathbb{R}_+$. We will suppose throughout that each γ_{ij} is stochastic. More specifically, assume that:

$$\gamma_{ij} = \gamma_j + \frac{1}{\theta_i} \epsilon_{ij}, .$$

where both γ_j and ϵ_{ij} are random iid variables with continuous and symmetric distribution supported on bounded intervals $[\gamma_L, \gamma_H]$ and $[\epsilon_L, \epsilon_H]$, with means $\bar{\gamma}$ and 0 and variances $\sigma_\gamma^2 > 0$ and $\sigma_\epsilon^2 \geq 0$ respectively. The first of these terms can be thought of capturing the relative "quality" of product j, whereas the second captures the idiosyncratic value consumer i derives from j. Goods are therefore both vertically and horizontally differentiated.

The parameter $\theta_i \in \{\theta_L, \theta_H\}$ then captures the extent to which consumer segment i values quality relative to their idiosyncratic assessment of the good. We will assume that all of the elements of the buyer's demand function, which are non-stochastic,

[1] While the settings in Chapters 1 and 2 are similar, there are differences in the notation used.

including θ_i, are common knowledge. As in Chapter 1, we also assume that $x_{ij} > 0$, which is guaranteed when b is sufficiently large, the conditions for which we discuss in more detail below.

Once the value of each γ_{ij} is realised, consumer segment i's ex-ante expected demand function for a good j can be represented by $\mathrm{E}[x_{ij}(\boldsymbol{p}, \gamma_{ij})] : \mathbb{R}^{m+1} \to \mathbb{R}$:

$$\mathrm{E}[x_{ij}(\boldsymbol{p})] = w_{ij}(b(a - p_j) + \theta_i\gamma_{ij} + \sum_{k=1}^{m} c(p_k - a)).$$

Define as j's expected demand function, $\mathrm{E}[x_j(\boldsymbol{p}, \boldsymbol{\gamma_j})] : \mathbb{R}^{m+n} \to \mathbb{R}$ where $\boldsymbol{\gamma_j}$ is a $n \times 1$ whose ith vector in the following way:

$$\mathrm{E}[x_j(\boldsymbol{p})] = \sum_{i=1}^{n} \mathrm{E}[x_{ij}].$$

Then, assuming the marginal cost of each seller is zero, the profit function for seller j is:

$$\mathrm{E}[\pi_j(\boldsymbol{p})] = p_j \mathrm{E}[x_j(\boldsymbol{p})].$$

Sellers are assumed to observe the entire graph G: they observe the matrix of observation probabilities. On the basis of this observation, sellers choose their prices. A seller, j, is assumed not to be able to price discriminate across buyers, and hence each sets a single price $p_j \in \mathbb{R}_+$. Sellers compete with one another on price, and set prices simultaneously. Therefore each seller's maximisation problem can be expressed as:

$$\max_{p_j} \pi_j(p_j,\ p_{\_j};\ \boldsymbol{G}).$$

Having characterised the game played by the sellers given the graph G, it is necessary to set out how G is determined. Define the rating of a product j as follows:

$$\hat{\gamma}_j := \frac{\sum \gamma_{ij}}{n}.$$

Ratings are then the averaging of the result of a process in which previous sets of buyer reveal their intercept parameters. By the law of large numbers, as $n \to \infty$, $\hat{\gamma}_j \to^p \gamma_j$. We will assume throughout that n is sufficiently large such that $\hat{\gamma}_j$ is very close to γ_j.

Having set out the demand function of the buyers, it is necessary to outline the observation process of the buyers. We will assume that the probability that a buyer i observes a seller j, w_{ij}, is determined in the following way:

$$w_{ij}(\upsilon, \kappa_i) = \begin{cases} \upsilon + \kappa_i(\hat{\gamma}_j - \frac{\sum_k \hat{\gamma}_k}{m}) & \text{if } \upsilon + \kappa_i(\hat{\gamma}_j - \frac{\sum_k \hat{\gamma}_k}{m}) \leq 1 \\ 1 & \text{if } \upsilon + \kappa_i(\hat{\gamma}_j - \frac{\sum_k \hat{\gamma}_k}{m}) > 1 \end{cases}, \tag{2.2}$$

where $\upsilon \in \mathbb{R}_+$ is the default probability that a consumer segment observes any of the sellers and $\kappa_i \in \mathbb{R}$ is the weight the observation function of i places on above average ratings. When $\kappa_i > 0$, if a product's rating is relatively high, the probability that they will be observed is greater than a product with a lower than average rating.

Note that (2.2) implies that the expected observation probability is equal for each seller, and the expected sum of every observation probability is also equal. Throughout we will assume that m is sufficiently large such that $\frac{\sum_k \hat{\gamma}_k}{m}$ is within ε of $\bar{\gamma}$ with high probability, and for simplicity we will write the observation function as if $\frac{\sum_k \hat{\gamma}_k}{m} = \bar{\gamma}$.

Let $\boldsymbol{\kappa}$ represent an $m \times 1$ whose ith component is κ_i. We define $\Gamma(\upsilon, \boldsymbol{\kappa})$ as the simultaneous move m-player game for the vectors υ and $\boldsymbol{\kappa}$ with payoffs as specified in (2.6) and strategy spaces $\mathbb{R}_+$. Section 2.4 will outline the equilibrium price vector for a given $\Gamma(\upsilon, \boldsymbol{\kappa})$. Given the equilibrium price setting behaviour, the platform owner is assumed to choose $\upsilon, \boldsymbol{\kappa}$ that maximise the following function:

$$\pi_P^* = \chi \sum_{j=1}^{m} \pi_j^*.$$

We will assume throughout that the platform chooses υ and $\boldsymbol{\kappa}$ prior to the realisation of the consumer preferences. Hence, the timing of the model can be summarised as follows:

1. The platform owner chooses υ and the vector $\boldsymbol{\kappa}$;
2. Buyer preferences are realised and the ratings vector $\hat{\gamma}$ is thus realised;
3. The realisation of $\hat{\gamma}$ generates a graph G;
4. Sellers set prices according to G.

In the next section we discuss the platform owner's choice of $\boldsymbol{\kappa}$ and υ in more detail.

2.4 Determining observation probabilities

Equation (2.2) in Section 2.3 defines an observation function that determines the probability that each consumer segment observes each seller. The platform owner's choice of υ and $\boldsymbol{\kappa}$ then determines the structure of competition in the network: the more observation probabilities are affected by a seller's rating, the more competition among highly valued sellers. This has implications for the competition structure of the network: competition is more concentrated, with some sellers competing much more fiercely than in a case where the observation function is unaffected by seller ratings, which (all else equal) will lead to lower prices.

We will largely restrict our analysis to the case where $\kappa_i = \kappa$ for all i. In this case, the probability that i observes j can be expressed:

$$w_{ij}(\upsilon, \kappa) = \begin{cases} \upsilon + \kappa(\hat{\gamma}_j - \bar{\gamma}) & \text{if } \upsilon + \kappa(\hat{\gamma}_j - \bar{\gamma}) \leq 1 \\ 1 & \text{if } \upsilon + \kappa(\hat{\gamma}_j - \bar{\gamma}) > 1 \end{cases} \quad \forall i. \tag{2.3}$$

The probability that a buyer observes a seller j is increasing linearly in both the informativeness parameter υ and the difference between j's rating and the sellers' average quality. Notice that $w_{ij} = w_{lj}$ for all i, l pairs, and it is therefore not possible for the platform to differentiate between different buyers in this setting.

In Section 6, we consider the case where the platform owner chooses the optimal observation function, in the sense that they maximise their profit by choosing the optimal level of υ_i and κ_i for a given set of model parameters. This captures the notion of the platform owner being able to optimally bias their search process towards sellers with high ratings.

In Section 7, the platform owner is unable to choose κ explicitly, due to e.g. technology constraints, and instead faces an observation function with a fixed κ and is only able to choose between making ratings public or not. If they make the information public, the information environment is "ratings-based", while if they do not, the environment is said to be a "random match" environment. These two environments generate different structures for a given realisation of the rankings vector $\hat{\boldsymbol{\gamma}}$, as Figure 2.1 shows.

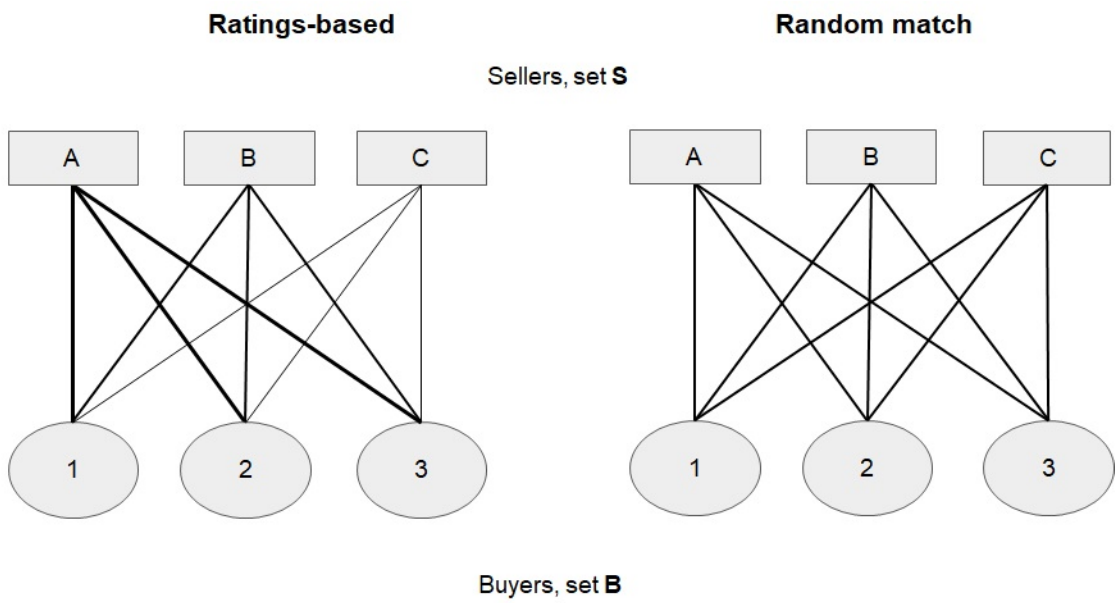

Figure 2.1: Ratings-based and random match environments where the product quality parameters are as follows: $\hat{\gamma}_A > \hat{\gamma}_B > \hat{\gamma}_C$. The thicker a line, the higher the probability of observation.

The graph generated by the ratings-based environment is such that buyers are more likely to observe highly-rated sellers, hence, highly-rated sellers are more likely to compete with one another than they are with sellers with a lower rating. In contrast, in the random match environment the probability of observation is the same for each buyer-seller pair, which implies that competition is less concentrated. The concentration of competition, as we will see, has an impact on the nature of competition and equilibrium prices, which in turn affects network profits.

2.5 The price setting equilibrium

To characterise the equilibrium price vector of the sellers competing in G, it should be noted that buyers are not strategic agents in the game, and therefore it is possible to define a network which is strategically equivalent to an original network G, but only includes seller nodes. Such a network can be used to characterise the equilibrium of the game $\Gamma(\upsilon, \boldsymbol{\kappa})$. Define:

$$\alpha_j := \sum_{i=1}^{n} [w_{ij}(\upsilon, \kappa_i)(ab + \theta_i \gamma_{ij}) - \hat{c}_{jk},$$

where:

$$\hat{c}_{jk} := \sum_{k=1}^{n} \sum_{i=1}^{n} acw_{ij}(\upsilon, \kappa_i) w_{ik}(\upsilon, \kappa_i).$$

Also, define:

$$\beta_j := a \sum_{i=1}^{n} w_{ij}(\upsilon, \kappa_i).$$

Transforming the profit function in (2.6) to put it in terms of a seller-only network yields the following payoff function for seller j:

$$\pi_j(\boldsymbol{p}) = p_j(\alpha_j - \beta_j p_j + \sum_{k=1}^{m} \hat{c}_{jk} p_k),$$

Rescaling (2.6) by $1/\beta_i$ and multiplying by $\frac{1}{2}$ generates the following:

$$\tilde{\pi}_j(\boldsymbol{p}) = p_j \tilde{\alpha}_j - \frac{1}{2} p_j^2 + \sum_{k=1}^{m-1} \tilde{c}_{jk} p_j p_k$$

where $\tilde{\alpha}_j = \frac{\frac{1}{2}\alpha_j}{\beta_j}$ and $\tilde{c}_{jk} = \frac{\frac{1}{2}\hat{c}_{jk}}{\beta_j}$. The maximisation problem $\max_{p_i} \tilde{\boldsymbol{\pi}}(\boldsymbol{p_i})$ has the same set of first-order conditions as the one which involves maximising a vector containing the profit functions in (2.6). Define $\lambda = \max\{\widetilde{c}_{jk} | j \neq k\}$ and R_S as a symmetric zero diagonal matrix of a network G_S with entries $\frac{\widetilde{c}_{jk}}{\lambda}$. As in Chapter 2, this transformation yields a *competition network*, G_S, which is a projection of G.

Let $\tilde{\boldsymbol{\alpha}}$ represent a $m \times 1$ vector with element j equal to $\tilde{\alpha}_j$. To ensure that there is a unique equilibrium in this setting, it is necessary to specify a condition under which demands are all positive for any equilibrium. Define $Z_{\tilde{\alpha}}$ as follows:

$$Z_{\tilde{\alpha}} := [\boldsymbol{I} - \lambda R_C]^{-1} \tilde{\boldsymbol{\alpha}}_{\boldsymbol{L}},$$

where R_C is the adjacency matrix of the complete seller-only graph, with $w_{ij} = 1$ for all i, j pairs and $\tilde{\boldsymbol{\alpha}}_L$ denotes the vector $\tilde{\alpha}$ when $\theta\gamma_{ij} = \theta\gamma_L + \epsilon_L$ for all i, j pairs.

Let $\boldsymbol{\gamma}_L$ denote the $m \times 1$ vector when the latter equality holds and Σ denote a $m \times m$ matrix with each diagonal entry equal to b and each of its off-diagonal components equal to $-c$. It is clear that $\boldsymbol{\Sigma}(\boldsymbol{a} - \boldsymbol{Z}_{\tilde{\boldsymbol{\alpha}}})$ is increasing in b. Moreover, there exists a $\bar{b} \in \mathbb{R}_+$ such that if $b > \bar{b}$, then $\boldsymbol{\Sigma}(\boldsymbol{a} - \boldsymbol{Z}_{\tilde{\boldsymbol{\alpha}}}) + \boldsymbol{\gamma}_L > 0$.

Define $C_{\tilde{\alpha}}(G_S, \lambda) := [\boldsymbol{I} - \lambda R_S]^{-1} \tilde{\boldsymbol{\alpha}}$, which is the weighted transformed Bonacich centrality measure of the network G_S. The following proposition, which characterises the equilibrium price vector then holds:

Proposition 2.1. *Suppose that $b > \bar{b}$. Then the game $\Gamma(\upsilon, \kappa)$ has a unique Nash equilibrium in pure strategies, which is the equilibrium price vector:*

$$\boldsymbol{p}^*(\upsilon, \boldsymbol{\kappa}) = \boldsymbol{C}_{\tilde{\boldsymbol{\alpha}}}(\boldsymbol{G_S}, \boldsymbol{\lambda}) \tag{2.4}$$

iff $(\boldsymbol{I} - \lambda R_S)$ is positive definite.

There exists a unique Nash equilibrium price vector for a game $\Gamma(\upsilon, \boldsymbol{\kappa})$ which is equal to the Bonacich centrality of the sellers in the network G_S multiplied by $\tilde{\boldsymbol{\alpha}}$, assuming that

$(\boldsymbol{I} - \lambda R_S)$ is positive definite and at the equilibrium vector of prices, each consumer demands a non-zero amount of each good.

The assumption after the "iff" statement in Proposition 2.1 guarantees that, for whatever the structure of the graph G and the realisation of the set of intercept parameters, $x_{ij} > 0$ for all i, j pairs. Specifically, this assumption states that, even if the graph structure results in the most competition possible (i.e. it is complete) and demand is as low as possible, $x_{ij} > 0$. This guarantees the existence of an equilibrium.

Define $\Gamma_c(\upsilon, \kappa)$ as a game in which $\theta_L = \theta_H = \theta$, $(\boldsymbol{I} - \lambda R_S)$ is diagonally dominant and $b > \bar{b}$. Also, define:

$$\omega_j := \frac{\sum_{i=1}^{n} w_{ij}\gamma_{ij}}{\beta_j}.$$

Let $\boldsymbol{\omega}$ denote a $m \times 1$ vector whose jth component is ω_j and, abusing notation slightly, let $\boldsymbol{C} = \boldsymbol{C(G_S, \lambda)}$. Let $\mathrm{E}[\boldsymbol{p}(\upsilon,\kappa)]$ denote the expected price vector for υ and κ. Under the assumption that there is no difference between types, it can be shown that equation (2.4) in Proposition 2.1 can be written as follows:

$$\mathrm{E}[\boldsymbol{p}^*(\upsilon, \kappa)] = \frac{1}{2}(\boldsymbol{a} - \mathrm{E}[\boldsymbol{C}]\boldsymbol{a} + \theta \mathrm{E}[\boldsymbol{C\omega}]), \tag{2.5}$$

where $\boldsymbol{a}$ is the vector $a\boldsymbol{1_m}$. Hence, there exists a $\bar{\theta} \in \mathbb{R}_+$ such that if $\theta < \bar{\theta}$, then seller j's expected price is decreasing in their weighted expected Bonacich centrality for all j. Throughout, we will assume that $\theta < \bar{\theta}$.

2.6 Weighting ratings

We will examine the platform owner's choice of optimal observation function, where the platform owner is able to optimise their payoff with respect to both the expected observation probability (υ) and the extent to which the observation function is biased towards highly-rated sellers (κ).

First, we will consider a case where the observation probabilities are of the form specified in (2.3). The probability that a buyer observes a seller j is increasing linearly in both the informativeness parameter υ and the difference between j's rating and the sellers' average quality. Notice that $w_{ij}(\upsilon, \kappa) = w_{lj}(\upsilon, \kappa)$ for all i, l pairs, and it is therefore not possible for the platform to differentiate between different buyers.

To understand some of the trade-offs inherent in the platform owner's maximisation problem, it is useful to suppose at first that the platform owner can only choose υ and that κ is fixed at 0. The platform owner then solves the following problem:

$$\max_{\upsilon} \pi_P^*(\upsilon, \kappa).$$

An increase in υ has two effects on profits. Increasing υ increases the probability that a given buyer is observed increases expected sales, which increases profits. However, there is a cost to increasing υ: it results in there being a higher probability that buyers observe more sellers, which creates more competition across the network. More competition reduces demand for any one good and reduces prices, decreasing profits as well.

To see the trade-off the platform owner faces more formally, we will split the marginal effect of υ into two parts. At the equilibrium value of υ, the marginal increase in profits associated with an increase in the probability that each buyer observes each seller can, by the envelope theorem, be expressed as follows:

$$\mathrm{MB}_\upsilon = n\chi \sum_j \mathrm{E}[p_j^*(ba + \theta\gamma_j - bp_j^*)|w_j \leq 1],$$

where p_j^* is seller j's price at the equilibrium. The expected marginal cost in terms of increased competition can similarly be defined:

$$\mathrm{MC}_\upsilon = \chi \sum_j \mathrm{E}[\sum_{k=1}^{m} 2\upsilon n c p_j^*(p_k^* - a) + \sum_{l=1}^{m}\sum_{k=1}^{m}[\hat{c}_{kl}\frac{\partial p_k^*}{\partial \upsilon}]p_l^*|w_j \leq 1].$$

When $\kappa = 0$ and the solution is interior, the platform owner's profit is maximised when the marginal benefit and marginal cost sum to zero. The following proposition holds:

Proposition 2.2. *Suppose $\kappa = 0$ and that the game $\Gamma_c(\upsilon, 0)$ is played after the platform chooses υ. If there is not a unique solution to the platform owner's above maximisation problem, then $inf\{max_\upsilon \pi_P^*(\upsilon, 0)\} = 1$. Otherwise, there exists a unique optimum, $0 < \bar{\upsilon} < 1$. Furthermore, $\bar{\upsilon}$ is decreasing in c, and hence there exists a $\varphi \in \mathbb{R}_+$such that if $c > \varphi$, then $\bar{\upsilon} < 1$.*

Seller, j's equilibrium price is decreasing in their Bonacich centrality, which is convex in υ. Hence, the marginal cost associated with increasing υ is convex, while the equivalent marginal benefit is linear. Hence, either a corner solution obtains (i.e. $\upsilon^* = 1$ is a solution to the platform owner's optimisation problem) or there exists a solution, where $\upsilon = \bar{\upsilon} < 1$, which is a unique solution to that problem.

Increasing υ increases the level of competition in the network. If c is sufficiently high, then products are so substituable that the platform owner is willing to forgo some of the profit associated with buyers observing sellers for certain in order to increase prices. This is a result that is consistent with that found in Chapter 1: the platform owner faces a trade-off between expected sales on the one hand and competition and lower prices on the other.

Now, consider the case where the platform owner chooses κ as well as υ. Increasing κ has the effect of increasing the relative weight placed on j's rating when determining the probability that j will be observed. Define the platform owner's profit maximisation problem as follows:

$$\max_{\kappa,\nu} \pi_P^*(\upsilon, \kappa) \tag{2.6}$$

Theorem 2.1 characterises the solution to this maximisation problem in terms of κ as follows:

Theorem 2.1. *Suppose $\bar{\upsilon} < 1$ and that the game Γ_c is played after the platform chooses υ and κ. Then there exists a $\bar{c} \in \mathbb{R}_+$ where if $c < \bar{c}$, then the first-order conditions of the above maximisation problem are uniquely solved by the vector (υ^*, κ^*) such that $\kappa^* > 0$ and $\upsilon^* < \bar{\upsilon}$.*

The assumption that $c < \bar{c}$ guarantees that the platform owner's profit maximisation problem is concave in κ and υ, and will be maintained throughout. Theorem 2.1 is a result of a trade-off the platform owner faces in making $\kappa > 0$. A positive κ^* results increases the probability that a high quality product is observed for a given level of υ. Buyers demand more high quality products for a given price level. Hence, if they are also more likely to observe those products, then the platform owner's profit is increasing in the variance of the quality parameter, holding prices constant.

At the same time, κ^* being greater than zero has implications for the structure of competition on the network. It implies that highly-rated sellers are more likely to be competing in more markets (ex-post) and are in particular more likely to be in competition with one another. Specifically, the higher κ, the more concentrated competition is between highly-rated players. If the probability that a seller competes with a set of sellers is high, then this has a disproportionately large effect on the prices of the other sellers. In the model, this is the result of the fact that a seller j's centrality in the graph G_S is convex in the weight between j and k, $\hat{c}_{jk}$.

Intuitively, there being a-higher-than average probability that two sellers compete with one another drives their own prices down, which propagates across the network. Weighting the observation function by rating results in high quality sellers competing with one another with relatively high probability, driving down prices. The effect of concentrated competition is depicted in Figure 2.2.

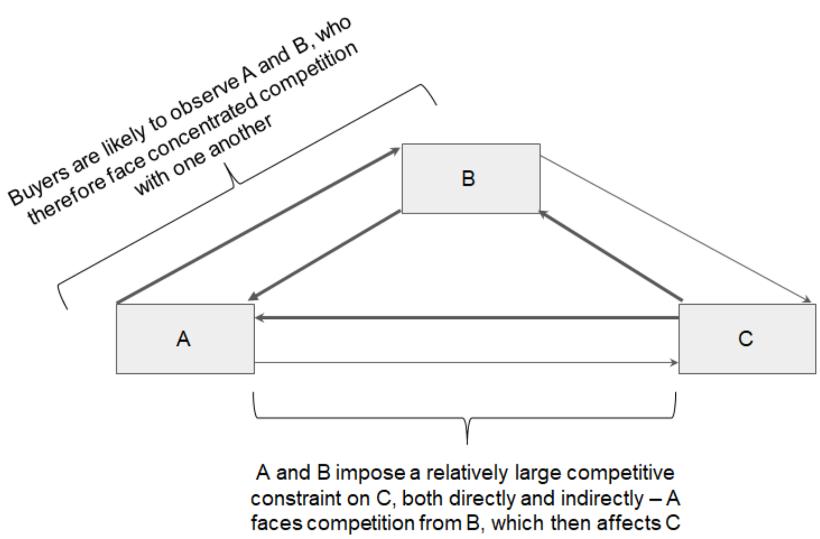

Figure 2.2: The effect increasing κ above zero has on competition, where $\hat{\gamma}_A > \hat{\gamma}_B > \hat{\gamma}_C$.

Hence, the platform owner faces a trade-off between increased revenue (holding prices constant) resulting from buyers observing high quality sellers and the fact that increasing the probability highly-rated sellers are observed increases competition. Assuming the optimal network structure when $\kappa = 0$ is not just the complete graph (i.e. assuming

$\bar{\upsilon} < 1$), then at $\kappa^* = 0$, the benefit of increasing the probability a high-quality seller is observed outweighs the cost associated with more concentrated competition and hence there is an incentive to set $\kappa^* > 0$.

To see the benefit of raising κ above 0 formally, it is useful to consider the marginal benefit of increasing κ on profit when $\kappa = 0$, defining this marginal benefit in a similar way to the way in which we defined MB_υ above:

$$\mathrm{MB}_\kappa := \chi \sum_j \mathrm{E}[(\hat{\gamma}_j - \bar{\gamma})(p_j^*(ba + \theta\gamma_j - bp_j^*)|w_j < 1].$$

When $\mathrm{E}[\hat{\gamma}_j^2] > \bar{\gamma}^2$ (i.e. when $\sigma_\gamma > 0$), $\mathrm{MB}_\kappa > 0$. As the marginal cost of increasing κ is close to zero when $\kappa = 0$, it follows that $\kappa^* > 0$.

When $\kappa > 0$, the marginal cost of increasing υ increases, while the marginal benefit remains constant. Hence, the optimal level of υ is falling in κ, which implies that $\upsilon^* < \bar{\upsilon}$. Given that the platform owner's profits are submodular in κ, ν, this is an application of Le Chatelier's principle (Milgrom and Roberts, 1996). Setting $\kappa^* > 0$ concentrates competition, and the platform owner therefore has an incentive to reduce the baseline observability of every seller in the network to reduce competition.

Prices and comparative statics

The above analysis raises the question: what is the effect of κ being strictly positive on prices? To illustrate the effect of biasing observation probabilities on prices, we simplify the analysis by assuming that υ is fixed at the optimal point when $\kappa = 0$, $\bar{\upsilon}$. Then suppose the platform owner chooses κ such that they solve the following problem:

$$\max_\kappa \pi_P(\bar{\upsilon}, \kappa). \tag{2.7}$$

Given that $\frac{\partial^2 \pi_P(\bar{\upsilon},\kappa)}{\partial^2 \kappa} < 0$, the first-order condition $\frac{\partial \pi_P(\bar{\upsilon},\kappa^*)}{\partial \kappa} = 0$ defines a unique solution to this problem, assuming $\bar{\upsilon} = 0$. Denote this solution $\tilde{\kappa}$. Then the following result holds:

Proposition 2.3. *Suppose υ is fixed at $\bar{\upsilon}$ and the platform owner solves the $max_\kappa \pi_P(\bar{\upsilon}, \kappa)$. If $\bar{\upsilon} < 1$, then $E[\boldsymbol{p}(\bar{\upsilon}, 0)] < E[\boldsymbol{p}(\bar{\upsilon}, \tilde{\kappa})]$.*

Increasing κ results in more concentrated competition in the network, reducing expected prices for a given $\tilde{\upsilon}$. High quality sellers are in competition with one another, which increases competition and lowers prices. Even in the case where prices are lower as a result of biasing buyer observations to high quality sellers, the platform owner is willing to lower prices in order to increase the probability that higher quality sellers are observed.

Let $\tilde{\kappa}$ again denote the solution to the platform owner's maximisation problem in (2.7) when $\upsilon = \bar{\upsilon} < 1$. The following results hold:

Proposition 2.4. *If $\tilde{\kappa}$ denotes the unique solution to the platform owner's maximisation problem stated in (2.7) and $\bar{\upsilon} < 1$, then $\frac{\partial \tilde{\kappa}}{\partial c} \leq 0$ and $\frac{\partial \tilde{\kappa}}{\partial \theta} \geq 0$.*

As the sellers' goods become more substituable, the cost of a given level of competitive overlap between the sellers increases. Hence, as c increases, prices are lower for a given level of κ. An increase in c increases the effect of the concentration in competition implied by $\kappa > 0$. highly-rated sellers compete with one another with higher probability when $\kappa > 0$, and the relative strength of this competition effect is increasing in product substitutability.

As θ increases, the marginal benefit of increasing κ also increases. This follows because (a) $\frac{\partial \mathrm{MB}_\kappa}{\partial \theta} = \chi \sum_j (\hat{\gamma}_j - \bar{\gamma})(\gamma_j) > 0$ and (b) $\frac{\partial p_j}{\partial \theta} > 0$ by equation (2.5). It follows that κ^* is increasing in θ. Intuitively, as the extent to which buyers value quality increases, the value of matching buyers with high quality products also increases, which implies that the optimal value of κ must rise.

The results thus far have considered the case in which the observation function is the same for all buyers. We relax this assumption in the section below.

Heterogeneous consumer segments

Thus far, we have assumed that the observation function is the same for each group of buyers. In this section we relax these assumptions, examining the effect doing so has on the platform owner's choice of observation function.

First, assume that, rather than choosing a single parameter κ, the platform owner chooses an observation function for both high-type buyers and low-type consumer segments. Recall that a seller for whom $a_i = a_L$ is a low-type consumer segment and a high-type segment is similarly defined. A high-type segment is more sensitive to quality than a low-type segment.

Let $i \in \{L, H\}$ denote a consumer segment's type, and let w_{Lj} denote the probability that a low-type segment observes a seller j. We write the observation probability function as follows:

$$w_{ij}(\upsilon, \kappa_i) = \begin{cases} \upsilon + \kappa_i(\hat{\gamma}_j - \bar{\gamma}) & \text{if } \upsilon + \kappa_i(\hat{\gamma}_j - \bar{\gamma}) \leq 1 \\ 1 & \text{if } \upsilon + \kappa_i(\hat{\gamma}_j - \bar{\gamma}) > 1 \end{cases}.$$

where $\kappa_i \in \{L, H\}$. We therefore allow the observation probabilities for a given seller to differ between types of consumer segments.

To illustrate the effect of introducing heterogeneous consumer segments, we suppose that $\theta_H = \theta_L + \varepsilon < \bar{\theta}$, so that high-types only have slightly more demand for a given set of products. Under this assumption, it is clear that each seller's price continues to be decreasing in their centrality. To simplify the problem further, we assume that υ is fixed at some point such that $\upsilon = \bar{\upsilon} < 1$.

Given the above assumptions, the platform owner's maximisation problem can be expressed:

$$\max_{\kappa_L,\kappa_H} \pi_P^*(\kappa_L, \kappa_H; \boldsymbol{\theta}),$$

where $\boldsymbol{\theta} = (\theta_L, \theta_H)^T$. The above maximisation problem has two first-order conditions such that:

$$\frac{\partial \pi_P^*(\kappa_L^*, \kappa_H^*; \boldsymbol{\theta})}{\partial \kappa_i} = 0 \text{ for } i = L, H.$$

The following proposition outlines the effect of an increase in θ_H on a solution to these conditions, κ_L^* and κ_H^*:

Proposition 2.5. *Suppose that $\bar{\upsilon} < 1$. Then $\frac{\partial \kappa_L^*}{\partial \theta_H} < 0$ and $\frac{\partial \kappa_H^*}{\partial \theta_H} > 0$.*

As θ_H increases, the difference in demand for high quality products from high-type and low-type consumer segments increases. Therefore, when θ_H increases, the platform owner has an incentive to increase the probability that high-types observe those products above the probability that low-types observe them by increasing κ_H^*.

However, by increasing the probability that high quality sellers are observed, the platform owner concentrates competition between high quality sellers. This has the effect of reducing prices. To partially compensate for this price effect, the platform owner has an incentive to reduce the probability that low-types observe high quality sellers by reducing κ_L^*.

In the profit-maximising solution, then, the platform owner is willing to sacrifice high value matches for low-types in order to reduce the price effect associated with increasing the probability that such matches take place for high-types. The optimal graph resulting from this trade-off is depicted in Figure 2.3.

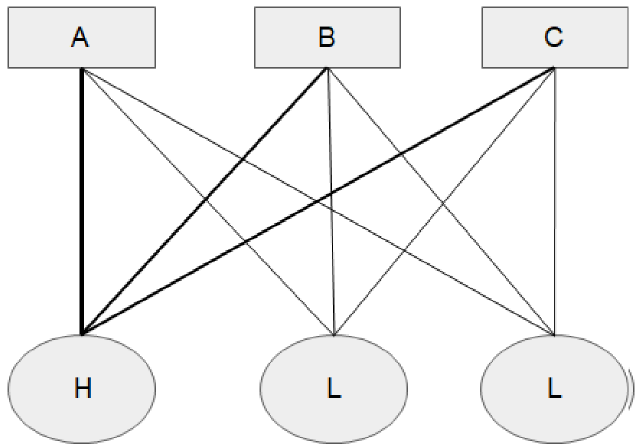

Figure 2.3: The optimal graph structure where $\hat{\gamma}_A > \hat{\gamma}_B > \hat{\gamma}_C$. The letter on each buyer group reflects their type.

The probability that the high-type segment observes the high quality product, A, is higher than for the same probability for low-types, who are less likely to observe A than they would be had the high-type been a low-type instead. More generally, this analysis implies that the platform owner has an incentive to ensure particular consumer segments observe high quality products and other segments are far less likely to observe these products, implying that consumer targeting is likely to be profitable.

2.7 The profitability of matching environments

Section 6 outlined the trade-offs faced by the platform owner when choosing the weight to place on ratings in order to maximise profit. We consider a constrained decision problem in which the platform owner chooses whether to make ratings public or not.

In doing so, it is possible to further illustrate the trade-off between matching and concentrated competition found previously, as well as chracterise the effect an increase in the idiosyncratic valuation variance has on the profitability of using ratings to shape the observation process.

We conduct a profitability comparison between two environments with observation functions of the form in (2.9), but with two different sets of values for κ and υ: the "random match" environment, where $\kappa = 0$, and another where $\kappa = \bar{\kappa} > 0$; the "ratings-based" environment. We will assume throughout that $\theta_i = \theta$ for all i.

Denote the platform owner's profit function $\pi_P(\upsilon, 0; \theta, \sigma_\epsilon, c)$. The profit generated by the random match environment can be defined as follows:

$$\pi_A(c, \sigma_\epsilon, \theta) = \max_{\upsilon \in R_+} \pi_P(\upsilon, 0; \theta, \sigma_\epsilon, c).$$

Define $\mathrm{argmax}_{\upsilon \in R_+} \pi_P(\upsilon, 0; \theta, \sigma_\epsilon, c) = \bar{\upsilon}$. Then the profit generated by the ratings-based environment can be defined:

$$\pi_R(c, \sigma_\epsilon, \theta) = \pi_P(\bar{\upsilon}, \bar{\kappa}; \theta, \sigma_\epsilon, c).$$

The random match environment generates the maximum profit for the platform, conditional on their only being able to choose υ when $\kappa = 0$. To make the comparison between two environments as tractable as possible, we assume that the rating-based environment is such that $\upsilon = \bar{\upsilon}$ and $\kappa = \bar{\kappa}$. The baseline between the two environments is therefore the same, but there is a fixed bias towards highly-rated products in the ratings-based environment which is not present in the random match environment.

The preceding analysis draws out the relative costs and benefits associated with the two information environments. A ratings-based environment results in buyers observing high quality sellers with a higher probability. However, such an environment results in more concentrated competition between highly-rated players, resulting in lower prices. The bias towards highly-rated sellers also increases the probability that a high-surplus match generated by a high idiosyncratic valuation is missed, as it increases the probability that a relatively lowly-rated seller is not observed by a buyer that values their product relatively highly.

A "random match" environment, in contrast, results in observation probabilities being equal across the network. This implies that the competition effect associated with one set of sellers being more central than others is not present. However, such a structure results in a lower probability of buyers observing a highly-ranked product, which reduces profits, all else equal.

To make a profitability comparison between the two environments meaningful, we will assume that there exists some θ^* such that both: (a) $\bar{\upsilon}(\theta^*) < 1$ and; (b) it is the case that $\bar{\kappa}$ is the optimal κ when $\theta = \theta^*$ and $\upsilon = \bar{\upsilon}$. That is, θ^* satisfies the following expression:

$$\text{argmax}_{\kappa \in R_+} \pi_P(\kappa, \bar{\upsilon}; \theta^*, \sigma_\epsilon, c) = \bar{\kappa}. \tag{2.8}$$

That is, when $\theta = \theta^*$, $\bar{\kappa}$ is the unique solution to the platform owner's optimisation problem if they were able to choose κ. We are therefore assuming that the vector $(\bar{\upsilon}, \bar{\kappa})'$ will be the unique solution to the unconstrained maximisation problem of the platform for some value of θ, θ^*. This is reasonable in that the platform owner's profit function is concave in κ, which implies that if $\upsilon = \bar{\upsilon}(\theta^*) < 1$, then there exists some κ that uniquely maximises the platform owner's profit by Theorem 2.1. Define:

$$\pi_R(c, \sigma_\epsilon, \theta) - \pi_A(c, \sigma_\epsilon, \theta) := \tilde{\pi}(c, \sigma_\epsilon, \theta).$$

We can use the preceding analysis to establish for what values of θ, if any, the random-match environment is preferred to the ratings-based environment. Theorem 2.2 shows that there are cases where not making ratings public (and therefore generating the random match environment) is preferred by the platform owner:

Theorem 2.2. *Suppose that $\sigma_\epsilon = 0$ and that $\Gamma_c(\upsilon, \kappa)$ is played in both environments. Then for a given $\bar{\kappa}$ there exists an $\tilde{\theta} \in R_+$ such that: if $\theta > \tilde{\theta}$ then $\tilde{\pi}(c, \sigma_\epsilon, \theta) < 0$; if $\theta < \tilde{\theta}$, then $\tilde{\pi}(c, \sigma_\epsilon, \theta) > 0$; and if $\theta = \tilde{\theta}$ then $\tilde{\pi}(c, \sigma_\epsilon, \theta) = 0$.*

Intuitively, Theorem 2.2 states that if $\bar{\kappa}$ is relatively large when θ is near zero then the concentration in competition associated with the ratings-based environment results in profit being lower than in the random match environment. When $\bar{\kappa}$ is large, relatively small differences in ratings result in large differences in observation probability, which has the effect of increasing competition, reducing the profit of the platform owner.

As θ increases, the value of matching a buyer with a high quality seller increases, which means that the incentive to bias the observation function to high quality products increases. The rating-based environment thus becomes more attractive to the platform owner. There is a point for a given $\bar{\kappa}$, $\tilde{\theta}$, at which the two environments yield the same pay-off. If $\theta > \tilde{\theta}$, then the optimal κ at θ is larger than $\bar{\kappa}$, and the platform owner always strictly prefers the profit associated with the rating-based environment.

Proposition 2.6 considers the effect of changes to c and θ:

Proposition 2.6. *Suppose that $c < c^{'}$ and $\theta < \theta^{'} < \tilde{\theta}(\bar{\kappa})$. Then $\tilde{\pi}(c, \sigma_\epsilon, \theta) \geq \tilde{\pi}(c^{'}, \sigma_\epsilon, \theta)$*

and $\tilde{\pi}(c, \sigma_\epsilon, \theta') \geq \tilde{\pi}(c, \sigma_\epsilon, \theta)$.

As goods become less substituable, the cost of the more concentrated competition associated with the ratings-based environment decreases, implying that the profit from that environment increases more than the profit from the random match environment.

The same is true of a change in how sensitive buyers are to changes in quality: as buyers value quality more, the higher the benefit associated with increasing the probability of buyers observing high-quality sellers. This implies that a ratings-based environment becomes more profitable relative to the random match environment, where high quality products are no more likely to be observed than low quality products.

An increase in σ_ϵ, however, has the effect of increasing the profit of the random match environment relative to the ratings-based environment.

Proposition 2.7. *Suppose that* $\sigma_\epsilon < \sigma_\epsilon'$ *and* $\frac{1-\bar{v}}{\bar{\kappa}} < (\gamma_H - \bar{\gamma})$. *Then* $\tilde{\pi}(c, \sigma_\epsilon, \theta) < \tilde{\pi}(c, \sigma_\epsilon', \theta)$.

As j's price and i's demand is increasing in ϵ_{ij}, j's expected profit is convex in ϵ_{ij}, a result consistent with Myatt and Johnson (2006). Thus, when the additional profit generated from a positive realisation of ε_{ij} is greater than the loss in profit associated with a negative realisation of ε_{ij}.

When $\kappa > 0$, the sum of the expected sum of the observation probabilities for any one consumer segment is smaller than the sum of the probabilities for $\kappa = 0$, assuming that $\bar{v} > 1 - \bar{\kappa}(\gamma_H - \bar{\gamma})$. As the distribution is symmetric, if a seller's quality is sufficiently large, then they will be observed with probability 1, but, clearly their observation probability cannot increase above that.

Given the assumption on $\frac{1-\bar{v}}{\bar{\kappa}}$, increasing κ from 0 to $\bar{\kappa}$ does not reduce the probability of observing any low-quality seller to 0. Hence, the total sum of observation probabilities is lower in the ratings-based environment than in the random match environment. This implies that by biasing observation towards high quality sellers, the platform owner increases the probability of missing out matching a consumer segment with a product that is a high-value idiosyncratic match.

The result in Proposition 2.7 then captures the notion that introducing ratings may result in buyers being less likely to observe (or pay attention to) lower rated products, resulting in them observing fewer products overall. As σ_ϵ increases, the reduction in the expected number of products observed by consumer segments is more costly because the probability of missing a high-value match also increases.

The preceding discussion is summarised in Figure 2.4 below. Figure 2.4 depicts the point at which $\tilde{\pi}(c, \sigma_\epsilon^{'}, \theta) = 0$ for given values of θ and σ_ϵ. As θ increases, the idiosyncratic variance in preferences required for the platform to be indifferent between the environments increases, as the rating environment becomes increasingly attractive to the platform owner.

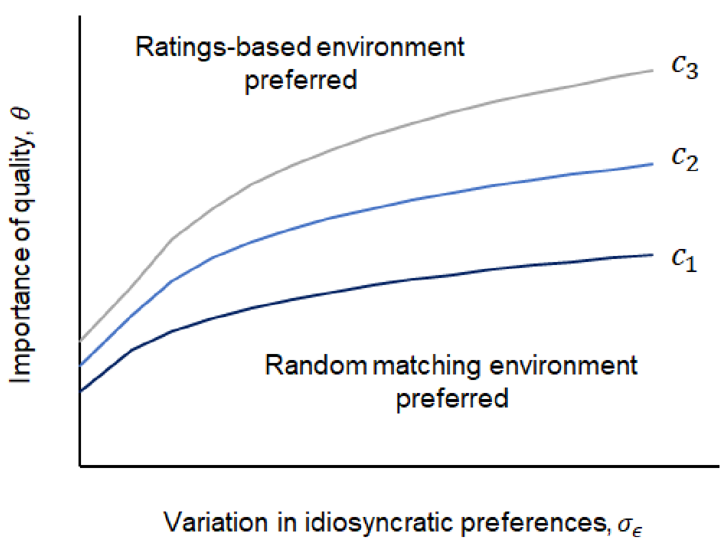

Figure 2.4: Isoprofit lines where $\tilde{\pi}(c, \sigma_\epsilon^{'}, \theta) = 0$ and $c_3 > c_2 > c_1$. The ratings-based environment is preferred at points below the relevant line.

2.8 Consumer welfare

Thus far we have analysed the extent to which biasing the search process towards high quality sellers increases profits. We now consider consumer welfare by assessing the effect of changes to the matching environment on consumer surplus. We define i's consumer surplus as follows:

$$\mathrm{E}[\mathrm{CS}_i(\boldsymbol{x_i^*}; \boldsymbol{p}^*)] = \sum_{j=1}^{m} \frac{1}{2}\mathrm{E}[x_{ij}^*(a + \frac{\theta}{b}(\gamma_{ij}) + \sum_{k=1}^{m} \frac{c}{b} w_{ik}(p_k^* - a)] - p_j^*)].$$

Define $\sum_i \mathrm{CS}_i(\boldsymbol{x_i^*}; \boldsymbol{p}^*)] = \mathrm{CS}(\boldsymbol{x}^*; \boldsymbol{p}^*(\upsilon, \kappa))$, where $\boldsymbol{x}^*$ is an $n \times m$ matrix whose ijth component is x_{ij}^*. Suppose that $\kappa = 0$, that the observation probabilities of the consumer segments take the form in equation (2.3) and a central planner chooses υ in order to maximise total consumer surplus. The central planner's maximisation problem can thus be stated:

$$\max_{\upsilon} \mathrm{CS}(\boldsymbol{x}^*; \boldsymbol{p}^*(\upsilon, 0)).$$

The following proposition holds:

Proposition 2.8. *Suppose the game $\Gamma_c(\upsilon^*, 0)$ is played after the central planner chooses υ^*. Any $\upsilon \geq 1$ is a solution to the central planner's above maximisation problem. Therefore, $E[w_{ij}] = 1$ for all i, j pairs.*

As in Chapter 1, increasing an observation probability w_{ij} increases consumer surplus because it both increases the probability of surplus-increasing sales and reduces prices across the network. The platform owner therefore harms consumer surplus if they reduce υ below 1, which they may have an incentive to do in order to increase prices.

Now, consider the case in which the central planner chooses between the two observation environments outlined in Section 7. Recall that the random match environment is one in which $\kappa = 0$ and $\upsilon = \bar{\upsilon}$, while in the ratings-based environment observation is biased towards highly-rated sellers. Define:

$$\tilde{\mathrm{CS}}(\bar{\upsilon}, \bar{\kappa}) = \mathrm{CS}(\boldsymbol{x}^{*}; \boldsymbol{p}^{*}(\bar{\upsilon}, \bar{\kappa})) - \mathrm{CS}(\boldsymbol{x}^{*}; \boldsymbol{p}^{*}(\bar{\upsilon}, 0)).$$

Theorem 2.3 compares the two observational environments in terms of their effect on consumer welfare:

Theorem 2.3. *If* $1 - \bar{\kappa}(\gamma_H - \bar{\gamma}) > \bar{\upsilon}$ *and* $\Gamma_c(\upsilon^{*}, \kappa_i)$ *is played in both environments, then* $\tilde{CS}(\bar{\upsilon}, \bar{\kappa}) > 0$.

As discussed above, increasing κ imposes a cost on the platform owner in that it concentrates competition and reduces price. For the central planner, however, the price effect does not represent a cost but rather an additional benefit associated with increasing the probability that high quality sellers are observed. The ratings-based environment is unambiguously associated with lower expected prices and a higher probability of high quality matches than the random match environment, assuming that $1 - \bar{\kappa}(\gamma_H - \bar{\gamma}) > \bar{\upsilon}$, both of which yield higher expected consumer surplus.

The ratings-based environment is therefore always preferred to the random match environment by the central planner. In contrast, we have shown that the platform owner may prefer the random match environment if θ is sufficiently low, because the lower prices associated with the ratings-based environment are detrimental to firm profits.

2.9 Discussion

The analysis in Sections 2.6 and 2.7 helps explain why online platform use ratings differently to determine which sellers are more likely to be observed by buyers. In general, the platform owner faces a trade-off between increased competition and lower expected prices on the one hand and an increase in profits due to matching on the other. Our analysis suggests that the extent to which the platform owner has an incentive to bias their search process towards highly-rated sellers depends on the nature of the goods being sold on the platform.

As products become more substitutable, the cost imposed by the concentrated levels of competition imposed by the observation function being dependent on seller ratings increases. This would imply that platforms selling products that are relatively substitutable are less likely to use ratings to determine the probability buyers observe products. Platforms wholly or largely selling products like groceries, second-hand books, electronics and clothing, for example, would then be less likely to use ratings as prominently in their search processes as platforms selling products like arts and crafts, where products are generally less substitutable.

The analysis in Section 2.6 also suggests that when buyers value product quality more highly, the platform owner has a larger incentive to ensure that their search process is biased towards high quality sellers. Doing so ensures that high quality matches are missed with lower probability. Hence, platforms offering products that have high degrees of quality differentiation are likely to bias their search results more towards high quality products.

Our analysis in Section 2.7 allows us to consider how the type of products offered on a platform can affect the relative profitability of the two information environments being analysed. Markets for different products are assessed in terms of the extent to which buyers value quality and the variance of the idiosyncratic preferences that buyers have for them in Figure 2.5.

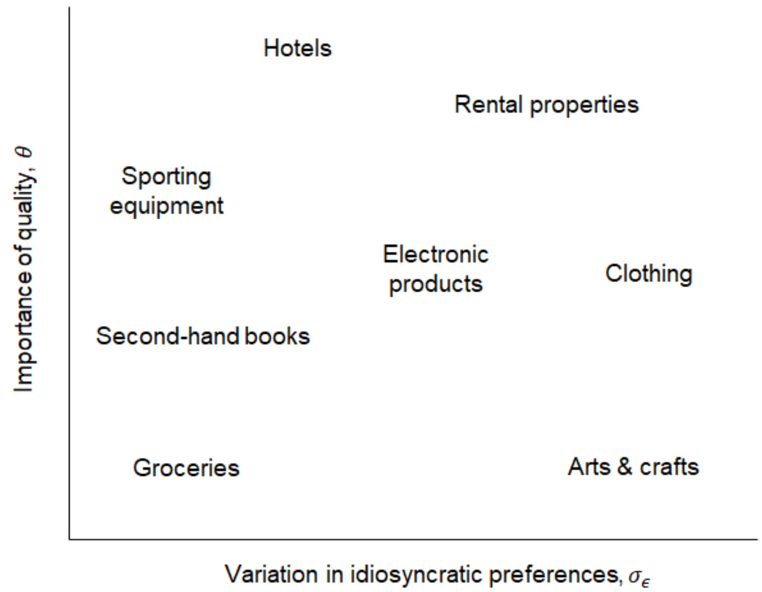

Figure 2.5: A comparison of different types of markets according to their values of importance of quality and the variance in idiosyncratic preferences

When platforms offer (or predominately offer) products with low values of θ, the concentrated competition effect identified in Section 2.6, whereby using ratings to determine observation probabilities increases the effective level of competition across the market, dominates. When the ex-post difference between the valuations of the sellers is low, the expected reduction in prices associated with biasing observation towards high-quality sellers outweighs any potential benefits associated with the ratings-based environment, which implies that it is relatively less profitable than the random match environment.

For products that have high idiosyncratic variance but relatively low expected differences in quality, the random match environment is relatively more profitable than the ratings-based environment. This is because there is an increasing probability of high-surplus matches driven by the convexity of the profit function. On the other hand, the relative profit associated with high-quality sellers being observed with a high probability increases as σ_γ increases, as a result of the same convexity in profit function. Hence, in such environments, the ratings-based environment would be more profitable.

These results are summarised in Figure 2.6, which depicts the isoprofit curve along which the two information environments yield equal profit, which is upward sloping and concave in $\sigma_\epsilon - \theta$ space.

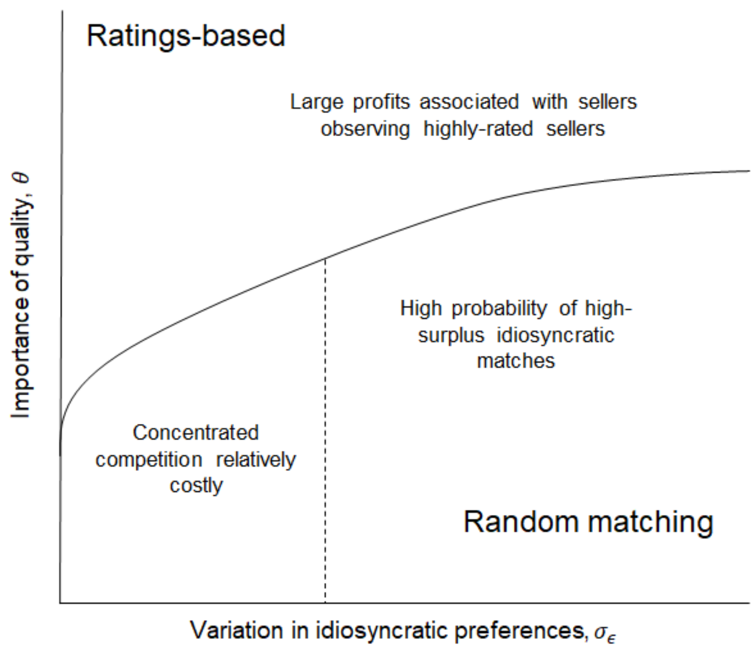

Figure 2.6: A comparison of the random match and ratings-based environments: for any market above the solid line, profits are higher in the ratings-based environments.

Figure 2.6 makes clear that platforms that (largely) host markets where product quality is less important to consumers should use an information environment that weights ratings less than platforms where product quality is more important to consumers. It is possible to approximate the positions of real-life platforms within the space shown in

Figure 2.6, which is shown in Figure 2.7.

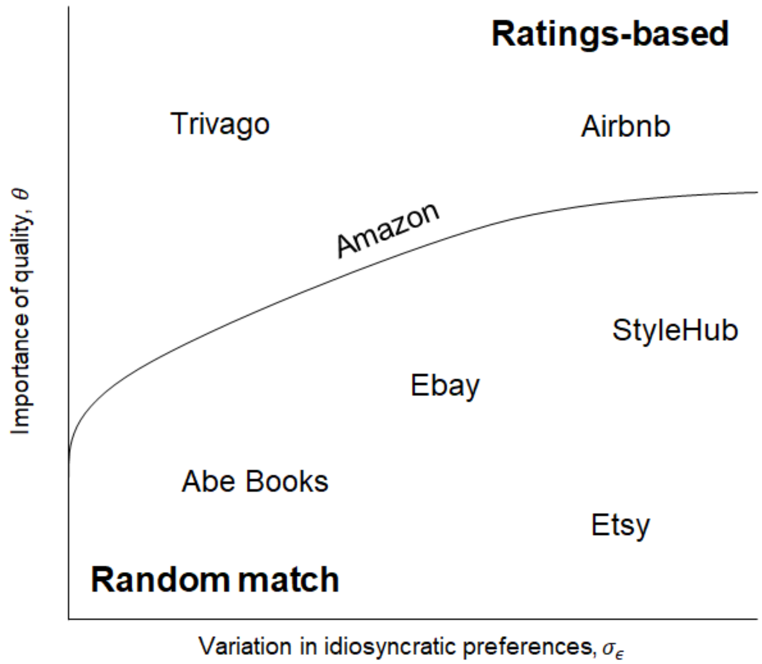

Figure 2.7: A comparison of the random match and ratings-based environments. For any market above the solid line, profits are higher in the ratings-based environments.

There are few studies that specifically examine the effect ratings have on the probability that buyers observe sellers have on these platforms. Much of the previous analysis relating to this topic finds that buyers are more likely to buy from highly-rated sellers. For example, Luca (2016) finds that an increase of 1 star on Yelp.com increases demand by 5-9%. Similarly, Ögut and Tas (2009) find that increases in customer ratings increase demand for hotel rooms in Paris and London.

These findings are often interpreted as capturing the effect that ratings have on influencing buyer valuations - a higher rating reveals information about the sellers on the platform, which in turn increases demand. It is likely that at least some of the effect of ratings captured by the empirical studies above is the result of the effect on consumer valuations. However, it is also clear that platforms use seller ratings as a part of their search process. If buyers are more likely to observe products, sales are more likely. The size of this effect, and the role it plays in determining competition, is relatively unexplored in the literature.

Our analysis generates several predictions that would be worthwhile exploring empirically:

1. Platforms selling products where perceived or actual quality is more important will use ratings as a larger component of their search process;

2. Platforms selling products with high idiosyncratic variance will use ratings as a smaller component of their search process;

3. Revealing product rankings may reduce prices within given platform or market;

4. Platforms selling relatively homogeneous products will use ratings as a smaller component of their search process.

There are to the author's knowledge no empirical studies that specifically examine these questions directly, and they would be important questions to consider in future empirical research.

2.10 Conclusion

The aim of this paper has been to examine the effect of using seller ratings to determine the probability that groups of sellers observe buyers on a large platform. If seller ratings are used to order search results on large platforms, then they determine which sellers are more likely to be observed, and therefore the structure of competition on the network. Empirical evidence makes clear that there the extent to which seller ratings are used to order search results differs considerably between platforms, and the aim of our analysis is to understand what drives these differences.

Our analysis indicates that there is a clear trade-off that the platform owner faces when using seller ratings to order search results. Increasing the extent to which sellers that are highly-rated are observed by buyers is potentially profit increasing because buyers are more likely to be matched with a high quality product, increasing willingness to pay.

At the same time, biasing search towards highly-rated sellers has the effect of concentrating competition among those sellers, which results in lower expected prices across the network. The platform owner resolves this trade-off in the case where they can freely choose the extent to which they bias their search results by limiting the use of ratings in the search process. As products become more substitutable, or consumers become less sensitive to product quality, the platform owner has an incentive to reduce the extent to which the search process is biased towards highly-rated products.

We also examine the case in which the platform owner chooses between an environment where they can either make ratings public or not. If they do so, then highly-rated sellers are more likely to be observed, but not necessarily in an optimal way: for example, the search process may be too biased towards highly-rated sellers, increasing competition and lowering prices. If the platform does not make ratings public, then consumers are less likely to observe high quality products, but competition is less concentrated.

If the variance of the idiosyncratic component of the buyers' valuation is large enough, then an environment where. matches are random is preferred because the cost associated with missing highly valuable matches driven by idiosyncratic preference gets larger. If the buyers' sensitivity is sufficiently high, the idiosyncratic match effect and the the reduction in prices due to concentrated competition are outweighed by the benefits of matching buyers to high quality sellers. Platforms in markets where buyers particularly value quality therefore have a greater incentive to reveal product ratings than in other markets.

Our analysis suggests that whether large, online platforms use seller ratings in their search process depends upon the characteristics of the market or markets they host. Markets that provide goods that are more tailored towards individual preferences, like arts and crafts, may be less willing to reveal ratings, or use them prominently in their search process, than in markets where buyer demand is sensitive to quality, like the hotel sector.

We show that it is optimal in some cases (e.g. when product substitutability is high) for the platform owner to choose random matching over the ratings-based environment, even when the expected number of sellers observed by consumers is the same in both cases. Consumers are always better off in the ratings-based environment, as in expectation prices are lower and matches have higher surplus than when matching is random. The platform owner's incentive to maintain high prices by choosing the random matching environment is therefore potentially detrimental to consumer welfare.

Appendix

Proof of Proposition 2.1

The proof of Proposition 2.1 follows directly from the proof of Theorem 1.1 in Chapter 1, noting that the assumption that $b > \bar{b}$ guarantees that $x_{ij} > 0$ if $w_{ij} > 0$ in the same way that $a > \bar{a}$ does in that chapter.

Proof of equation (2.5)

Recall that:

$$\boldsymbol{p}^*(\upsilon, \boldsymbol{\kappa}) = \boldsymbol{C_{\tilde{\alpha}}(G_S, \lambda)}$$

and:

$$\alpha_j = \sum_{i=1}^{n} [abw_{ij}(\upsilon, \kappa_i) + w_{ij}\theta_i\gamma_{ij}] - \hat{c}_{jk}.$$

Hence:

$$\mathrm{E}[C_{\tilde{\alpha}}(G_{\bar{S}}, \lambda)] = \mathrm{E}[\frac{1}{2}\sum_{k=0}^{\infty}\lambda^k R_S^k(\boldsymbol{a} + \theta\boldsymbol{\omega}) - \sum_{k=1}^{\infty}\lambda^k R_S^k \boldsymbol{a}].$$

The equality holds because each ϵ_{ij} is independent of R_S^k for all k, and hence $\mathrm{E}[R_S^k\epsilon_{ij}] = 0$ for all i, j, k.

$$\mathrm{E}[C_{\tilde{\alpha}}(G_{\bar{S}}, \lambda)] = \mathrm{E}[\frac{1}{2}\boldsymbol{a} - \frac{1}{2}\sum_{k=0}^{\infty}\lambda^k R_{\bar{S}}^k \boldsymbol{\gamma} + \sum_{k=0}^{\infty}\lambda^k R_S^k \theta\boldsymbol{\gamma}].$$

The last equality implies the result.

Proof of Proposition 2.2

Note that the ijth entry in the centrality matrix $C(G_{S,}\lambda)$ is equal to $\sum_l \lambda^l r_{ij}^{[l]}(G_S)$, where $r_{ij}^{[l]}(G_S)$ measures the weighted paths of length l that start at node i and end at node j. If $\kappa = 0$, then $C_j(G_{S,}\lambda)$ can be written as follows:

$$\mathrm{E}[C_j(G_S, \lambda)] = \sum_{k=0}^{\infty}(\frac{c}{b}\upsilon(m-1))^k \tag{2.9}$$

Each expected centrality is convex in υ, as $\mathrm{E}[C_j(G_S, \lambda)]_{\upsilon\upsilon} > 0$. As the expected equilibrium price vector is decreasing in the expected centrality of each seller j, this implies that the expected price vector is strictly concave in υ. Hence, if the solution of the first-order conditions of the platform owner's profit maximisation problem is interior, it represents a unique, global optimum, where $0 < \upsilon^* < 1$. If there is a corner solution, then any value of υ that results in $w_{ij} = 1$ for all i, j is optimal, which is just any $\upsilon \geq 1$, which implies that $\inf\{\max_\upsilon \pi_P^*(\upsilon, 0)\} = 1$.

It is clear from (9) that the absolute value of the derivative $\frac{\partial p_j}{\partial \upsilon}$ is increasing in c. Furthermore, the expression $2\upsilon n c p_j^*(p_k^* - a)$, which is a component of the marginal cost of increasing υ in the main text is also increasing in c. It follows that there exists a $\varphi \in \mathbb{R}_+$ such that if $c > \varphi$, then $\bar{\upsilon} < 1$.

Proof of Theorem 2.1

Note that the centrality matrix of G_S can be written:

$$C(G_S, \lambda) = \sum_{k=0}^{\infty} \lambda^k R_S^k.$$

Noting that given the assumptions on the observation probabilities, $w_{ik} = w_k$ for all i, then:

$$\mathrm{E}[\tilde{c}_{jk}] = c\mathrm{E}[\frac{nw_k w_j}{nw_j}] = c\mathrm{E}[w_k],$$

which follows from the fact that $\mathrm{E}[w_k w_j] = \mathrm{E}[w_k]\mathrm{E}[w_j]$, from the assumption that γ_j, γ_k are independent. Hence, an element of the kth column of the matrix R_S^3, r_k^3 is a function of $w_j^2 w_k$. Noting that the expected price vector is described by (2.5) and that:

$$w_j = \upsilon + \kappa(\hat{\gamma}_j - \bar{\gamma}),$$

then the second derivative of the weighted centrality with respect to κ is always positive. Consider the Hessian of the centrality of j:

$$H = \begin{bmatrix} \frac{\partial C_j}{\partial \upsilon \partial \upsilon} & \frac{\partial C_j}{\partial \upsilon \partial \kappa} \\ \frac{\partial C_j}{\partial \kappa \partial \upsilon} & \frac{\partial C_j}{\partial \kappa \partial \kappa} \end{bmatrix}.$$

Note that the expected adjacency matrix, $\mathrm{E}[R_S^n]$ where $n \geq 1$ includes a term that is a function of υ^n. Hence, $\frac{\partial C_j}{\partial \upsilon \partial \upsilon}$ features positive terms that are a function of c. This is not the case for the other components of H, which only have positive terms which are functions of c^n where $n \geq 3$. Hence, if c is sufficiently small, $\frac{\partial C_j}{\partial \upsilon \partial \upsilon}\frac{\partial C_j}{\partial \kappa \partial \kappa} > \frac{\partial C_j}{\partial \upsilon \partial \kappa}\frac{\partial C_j}{\partial \kappa \partial \upsilon}$ for all values of υ and κ and hence $|H| > 0$. This implies that the then centrality vector $C(G_S, \lambda)\mathbf{1}$ is convex in κ, υ, which implies that the platform owner's profit maximisation function is concave in (υ, κ). Thus, there exists a $\bar{c}$ such that if $c < \bar{c}$, then if there exists a interior solution to the platform owner's maximisation problem, then it is a unique, global optimum.

Note that seller j's profit function is twice differentiable and continuous, and $|H| < 0$. Hence, $\mathrm{E}[p_j(\upsilon, \kappa)]$ exists and is continuously differentiable, by the implicit function theorem. As per Milgrom and Segal (2002), these are jointly sufficient conditions for the envelope theorem to apply, and thus $\frac{d\pi_P(\upsilon^*,\kappa^*)}{d\kappa} = \frac{\partial\pi_P(\upsilon^*,\kappa^*)}{\partial\kappa}$.

To see that for any solution to the platform owner's maximisation problem it must be the case that $\kappa^* > 0$ when $\bar{\upsilon} < 1$, consider the fact that the expected marginal benefit, which we define as the first term of the summation inside the above expectation, associated with an increase in κ can be stated, by the envelope theorem, as follows:

$$\mathrm{E}[\mathrm{MB}_\kappa] = \sum_j \mathrm{E}[n(\hat{\gamma}_j - \bar{\gamma})(p_j^*(ba + \theta\gamma_j - bp_j^*))|0 \leq w_{ij} \leq 1] \tag{2.10}$$

If $\bar{\upsilon} < 1$ and $\sigma_y > 0$ then the expected marginal benefit of increasing κ is strictly positive when $\kappa = 0$, as $\mathrm{E}(\theta(\hat{\gamma}_j^2)] > \mathrm{E}[\theta(\hat{\gamma}_j\bar{\gamma})]$ because $\mathrm{E}[\hat{\gamma}_j^2] > \bar{\gamma}^2$. The expected marginal cost associated with increasing κ is as follows:

$$\mathrm{E}[\mathrm{MC}_\kappa] = \mathrm{E}[\sum_{k=1}^{m}(\sum_{j=1}^{m}\frac{\partial\hat{c}_{jk}}{\partial\kappa}p_j^*(p_k^* - a) + \sum_{l=1}^{m}[\hat{c}_{kl}\frac{\partial p_k^*}{\partial\kappa}]p_l^*)|0 \leq w_{ij} \leq 1] = 0. \tag{2.11}$$

Both $\frac{\partial\mathrm{E}\hat{c}_{jk}}{\partial\kappa}|_{\kappa=0}$ and $\frac{\partial\mathrm{E}p_k^*}{\partial\kappa}|_{\kappa=0}$ equal zero. To see this, note that $\mathrm{E}[\hat{c}_{jk}] = \mathrm{E}[(\upsilon + \kappa(\gamma_j - \bar{\gamma}))(\upsilon + \kappa(\gamma_k - \bar{\gamma}))]$ and, as $\mathrm{E}[\gamma_j\text{-}\bar{\gamma}] = 0$ and γ_j,γ_k are independent, it follows that for $\frac{\partial\mathrm{E}\hat{c}_{jk}}{\partial\kappa}$ to be non-zero, it must affect the total probability that both k and j are observed, and it does not do so at $\kappa = 0$ and $\bar{\upsilon} < 1$.

We can also show that $\frac{\partial\mathrm{E}p_k^*}{\partial\kappa}|_{\kappa=0} = 0$. By the analysis above, any entry of the expected centrality matrix that is a function of κ is a function of κ^n where $n \geq 2$. This follows as only entries that are functions of $\mathrm{E}[\hat{\gamma}_j^n]$ where $n \geq 2$ are functions of κ, and such entries are also functions of κ^n.

As $\mathrm{E}[\mathrm{MB}_\kappa|\kappa = 0] > 0$ and $\mathrm{E}[\mathrm{MC}_\kappa|\kappa = \bar{\kappa}] = 0$, it follows that $\kappa^* > 0$ as required.

To see that $\upsilon^* < \bar{\upsilon}$, note that for $n > 4$, there exists a least one entry, r_n, of the matrix R_S^n which is a function of $\upsilon w_j w_k^2$. Hence, if $\kappa > 0$, the marginal cost of increasing υ at $\bar{\upsilon}$ is larger than when $\kappa = 0$. It follows that the υ that solves the first-order condition in (2.16) is decreasing in κ^*. Alternatively, notice that the game is submodular in υ and κ. When $\sigma_\gamma = 0$, the vector $(0, \bar{\upsilon})'$ is an optimum. Hence, if $\sigma_\gamma > 0$, then, by Theorem 2 of Milgrom and Roberts (1996), it must be the case that $\upsilon^* < \bar{\upsilon}$. This is then an application of the global Le Chatelier's principle.

Proof of Proposition 2.3

From the results of Theorem 2.1, we know that if $\bar{\upsilon} < 1$, then $\mathrm{E}[\mathrm{MB}_\kappa | \kappa = 0] > 0$ and $\mathrm{E}[\mathrm{MC}_\kappa | \kappa = 0] = 0$ if $\sigma_\gamma > 0$. Hence $\tilde{\kappa} > 0$. From the functional form of $\mathrm{E}[C_j(\upsilon, \tilde{\kappa})]$, it is clear that $\mathrm{E}[C_j(\upsilon, \tilde{\kappa})] > \mathrm{E}[C_j(\upsilon, 0)]$, which implies the result.

Proof of Proposition 2.4

Consider the first-order condition of the platform owner's problem $\max_{\kappa,\nu} \pi_P^*(\bar{\upsilon}, \kappa)$, which we express as follows:

$$\pi_\kappa := \frac{\partial \pi_P(\tilde{\kappa}; \theta, c)}{\partial \kappa} = 0.$$

By the implicit function theorem:

$$\frac{\partial \tilde{\kappa}}{\partial \theta} = -\frac{\pi_{\kappa\theta}}{\pi_{\kappa\kappa}}.$$

To find the sign of $\pi_{\kappa\theta}$, we inspect (2.10) and (2.11) and note that these expression are equal at $\tilde{\kappa}$. By (2.10), $\mathrm{E}[\mathrm{MB}_\kappa]$ is clearly increasing linearly in θ. Furthermore, $\frac{\partial p_i(\upsilon,\kappa)}{\partial \theta} > 0$ for all i, which implies that $\mathrm{E}[\mathrm{MC}_\kappa]$ is increasing in θ as well. Therefore,

$\pi_{\kappa\theta} > 0$. As $\tilde{\kappa}$ is a maximum, $\pi_{\kappa\kappa} < 0$. Hence $\frac{\partial\tilde{\kappa}}{\partial\theta} > 0$. Noting that $\pi_{\kappa c} < 0$, then by similar reasoning, $\frac{\partial\tilde{\kappa}}{\partial c} < 0$.

Proof of Proposition 2.5

As $\theta_H = \theta_L + \varepsilon < \bar{\theta}$ by assumption, $\frac{\partial \mathrm{E}[\boldsymbol{p}^*]}{\partial \mathrm{E}[C(G_S,\lambda)]} < 0$. Denote the first-order conditions of the platform owner's maximisation problem as follows:

$$\pi_L := \frac{\partial \pi_P(\kappa_L^*, \kappa_H^*; \boldsymbol{\theta})}{\partial \kappa_L} = 0,$$

and:

$$\pi_H := \frac{\partial \pi_P(\kappa_L^*, \kappa_H^*; \boldsymbol{\theta})}{\partial \kappa_H} = 0.$$

These two first-order conditions provide a mapping from $\mathbb{R}^4 \rightarrow \mathbb{R}^2$. Denote the Hessian of the platform owner's maximisation problem as follows:

$$H = \begin{bmatrix} \pi_{HH} & \pi_{HL} \\ \pi_{LH} & \pi_{LL} \end{bmatrix}.$$

At any maximum (κ_L^*, κ_H^*) of the platform's first-order condition, $|H| > 0$. It is clear from equations (2.10) and (2.11) that:

$$\frac{\partial \pi_L}{\partial \theta_H} = 0,$$

and:

$$\frac{\partial \pi_H}{\partial \theta_H} > 0.$$

By the implicit function theorem:

$$H^{-1}\begin{bmatrix} \frac{\partial \pi_H}{\partial \theta_H} \\ \frac{\partial \pi_L}{\partial \theta_H} \end{bmatrix} = \begin{bmatrix} \frac{\partial \kappa_H^*}{\partial \theta_H} \\ \frac{\partial \kappa_L^*}{\partial \theta_H} \end{bmatrix}.$$

Therefore:

$$-\frac{1}{|H|}\begin{bmatrix} \pi_{LL} & -\pi_{LH} \\ -\pi_{HL} & \pi_{HH} \end{bmatrix}\begin{bmatrix} \frac{\partial \pi_H}{\partial \theta_H} \\ \frac{\partial \pi_L}{\partial \theta_H} \end{bmatrix} = \begin{bmatrix} \frac{\partial \kappa_H^*}{\partial \theta_H} \\ \frac{\partial \kappa_L^*}{\partial \theta_H} \end{bmatrix}.$$

It follows from the above expression that $\frac{\partial \kappa_H^*}{\partial \theta_H} > 0$ and $\frac{\partial \kappa_L^*}{\partial \theta_H} < 0$.

Proof of Theorem 2.2

First, it is necessary to check that the function $\tilde{\pi}(c, \sigma_\epsilon, \theta)$ is continuous in θ. To see this, recall that a seller j solves the problem:

$$\pi_j^*(\upsilon, \kappa) = \max\nolimits_{p \in R_+} \pi_j(\boldsymbol{p}).$$

Hence, if the conditions of the envelope theorem are met, we know that the following condition holds:

$$\frac{d\pi_A(c, \sigma_\epsilon, \theta)}{d\theta} = \chi \sum_j \frac{\partial \mathrm{E}[\pi_j^*]}{\partial \theta},$$

with the equivalent result holding for π_R. By inspection, π_A and π_R are then linear in θ, and hence both profit functions are continuous in θ. It follows that $\tilde{\pi}(c, \sigma_\epsilon, \theta)$ is a continuous function in θ as well.

Next, we show that $\tilde{\pi}(c, \sigma_\epsilon, \theta)$ is an increasing function on $\theta \in [0, \theta^*]$. Note first that when $\theta = 0$, $\kappa^* = 0$, which follows from the fact that:

$$\mathrm{E}[\mathrm{MB}_\kappa] = \sum_j \mathrm{E}[n(\hat{\gamma}_j - \bar{\gamma})(p_j^*(ba + \theta\gamma_j - bp_j^*))|0 \leq w_{ij} \leq 1] = 0$$

when $\theta = 0$ and $\mathrm{E}[\mathrm{MC}_\kappa|\upsilon = \bar{\upsilon}] > 0$. This then implies that there exists a value of θ such that $\tilde{\pi}(c, \sigma_\epsilon, \theta) < 0$. This is because $\kappa = 0$ satisfies the condition that $\mathrm{E}[\mathrm{MB}_\kappa|\kappa = 0] = \mathrm{E}[\mathrm{MC}_\kappa|\kappa = 0] = 0$ when $\theta = 0$.

Consider an increase in θ from 0. We can show that $\frac{\partial \pi_R}{\partial \theta} > \frac{\partial \pi_A}{\partial \theta}$ if $\sigma_\gamma > 0$. To see this, note that for all j the following expression holds:

$$\frac{\partial \mathrm{E}[\pi_j^*|\kappa = \kappa_i]}{\partial \theta} = \mathrm{E}[p_j^*((\bar{\upsilon} + \kappa_I(\gamma_j - \bar{\gamma})\gamma_j + \sum_{k=1}^{m} \hat{c}_{jk} \frac{\partial \mathrm{E}[p_k^*|\kappa = \kappa_i]}{\partial \theta})]|0 \leq w_{ij} \leq 1], \quad (2.12)$$

where $\kappa_A = 0$ and $\kappa_R = \bar{\kappa}$. As $\sigma_\gamma > 0$, the first term is greater when $\kappa = \bar{\kappa}$ than when $\kappa = 0$.

Now consider the second term in (2.12) and note that $\frac{\partial \mathrm{E}[p_k^*|\kappa = \kappa_i]}{\partial \theta} = \mathrm{E}[\boldsymbol{C}(\bar{\upsilon}, \kappa_i)\boldsymbol{\omega}|\kappa = \kappa_i]$. When $\kappa = 0$ the vector $\boldsymbol{\omega}$ and the matrix $\boldsymbol{C}$ are independent, as a seller's centrality is independent of the realisation of each γ_j. Hence:

$$\mathrm{E}[\boldsymbol{C}(\bar{\upsilon}, 0)\boldsymbol{\omega}|\kappa = 0] = \mathrm{E}[\boldsymbol{C}(\bar{\upsilon}, 0)]\mathrm{E}[\boldsymbol{\omega}].$$

However, in the case where $\kappa = \bar{\kappa} > 0$, $\mathrm{cov}(\mathrm{E}[\boldsymbol{C}(\bar{\upsilon}, \bar{\kappa}), \boldsymbol{\omega}) > 0$. This follows from the fact that each entry of the adjacency matrix, R_S can be expressed $\hat{c}_{jk} = \bar{\upsilon} + \bar{\kappa}(\gamma_k - \bar{\gamma})$ and thus $\boldsymbol{C}(\bar{\upsilon}, \bar{\kappa})$ is a function of γ_j^n for all j and $n \in [1, \infty]$. Hence, $\mathrm{E}[\boldsymbol{C}\boldsymbol{\omega}|\kappa = \bar{\kappa}] > \mathrm{E}[\boldsymbol{C}\boldsymbol{\omega}|\kappa = 0]$. It follows that $\frac{\partial \pi_R}{\partial \theta} > \frac{\partial \pi_A}{\partial \theta}$. It follows that $\theta^* > 0$.

By assumption, $\pi_R(c, \sigma_\epsilon, \theta^*)$ is equal to the result of an unconstrained maximisation problem where the platform owner freely chooses κ when $\theta = \theta^*$ and $\upsilon = \bar{\upsilon}$. Given that the platform is also free to set $\kappa = 0$ and $\upsilon = \bar{\upsilon}(\theta^*)$, but chooses not to, it follows that

$\pi_R(c, \sigma_\epsilon, \theta^*) > \pi_A(c, \sigma_\epsilon, \theta^*)$. Therefore $\tilde{\pi}(c, \sigma_\epsilon, \theta^*) > 0$. $\tilde{\pi}(c, \sigma_\epsilon, \theta)$ is thus continuous in θ and takes the value $\tilde{\pi}(c, \sigma_\epsilon, \theta) > 0$ and $\tilde{\pi}(c, \sigma_\epsilon, \theta) < 0$. By the intermediate value theorem, there exists an $0 < \theta < \theta^*$ such that $\tilde{\pi}(c, \sigma_\epsilon, \theta) = 0$.

Proof of Proposition 2.6

The fact that $\frac{\partial \pi_R}{\partial \theta} > \frac{\partial \pi_A}{\partial \theta}$ is proven in the proof of Theorem 2.2 above. This directly implies $\tilde{\pi}(c, \sigma_\epsilon, \theta') \geq \tilde{\pi}(c, \sigma_\epsilon, \theta)$.

Now, let κ_A denote the value of κ for the random match environment and κ_R denote the value of κ for the ratings-based environment. By the envelope theorem the effect of c increasing on profit can be expressed:

$$\frac{\partial \mathrm{E}[\pi^P|\kappa = \kappa_i]}{\partial c} = \chi \mathrm{E}[\sum_{j=1}^{m}\sum_{k=1}^{m} \frac{\partial \hat{c}_{jk}}{\partial_c} p_j^*(p_k^* - a) + \sum_{l=1}^{m}\sum_{k=1}^{m}(\hat{c}_{kl}\frac{\partial \mathrm{E}[p_k^*|\kappa_i]}{\partial c}]p_l^*)].$$

Recall that $\upsilon_R = \upsilon_A = \bar{\upsilon}$, but $\bar{\kappa}_R > 0$ and $\bar{\kappa}_A = 0$. As $|\frac{\partial \mathrm{E}[p_k^*]}{\partial c \partial \kappa}| > 0$, it follows that:

$$|\frac{\partial \mathrm{E}[p_k^*|\kappa = \bar{\kappa}]}{\partial c}| > |\frac{\partial \mathrm{E}[p_k^*|\kappa = 0]}{\partial c}|$$

which implies the first result in the proposition.

Proof of Proposition 2.7

Recall that $\gamma_{ij} = \theta(\gamma_j + \frac{1}{\theta})\epsilon_{ij}$. As ϵ_{ij} and γ_j are independent, ϵ_{ij} does not affect the centrality vector of the graph generated by υ, κ for all i and j. Hence:

$$\frac{\partial \boldsymbol{p}^*(\upsilon, \kappa)}{\partial \epsilon_{ij}} = \frac{\partial \tilde{\alpha}}{\partial \epsilon_{ij}} > 0.$$

Given the definition of $\tilde{\alpha}$, it is clear that this expression is postie and does not include ϵ_{ij}, and hence $\boldsymbol{p}^*(\upsilon, \kappa)$ is linear in each ϵ_{ij}. Consider the profit function of a seller j:

$$\pi_j(\boldsymbol{p}) = p_j(\upsilon, \kappa; \epsilon_{ij})(\alpha_j(\epsilon_{ij}) - \beta_j p_j(\upsilon, \kappa; \epsilon_{ij}) + \sum_{k=1}^{m} \hat{c}_{jk} p_k(\upsilon, \kappa; \epsilon_{ik}), \tag{2.13}$$

Each seller's profit, and therefore the platform owner's profit, is thus increasing in $\sigma_\epsilon^2 = \mathrm{E}\epsilon_{ij}^2$, as the platform owner's expected profit is equal to $\chi \sum_j \mathrm{E}[\pi_j(\boldsymbol{p})]$.

Note that $\sum_j \mathrm{E}[w_j|\kappa = \bar{\kappa}] \leq \sum_j \mathrm{E}[w_j|\kappa = 0]$ when $\bar{\kappa} > 0$, and $\frac{1-\bar{\upsilon}}{\bar{\kappa}} < (\gamma_H - \bar{\gamma})$. Note that as observation probabilities are bounded at 1, implying that:

$$\mathrm{E}[w_j|\kappa = \bar{\kappa}] = \bar{\upsilon} + \int \bar{\kappa}(\gamma_j - \bar{\gamma})\xi(\gamma_j)d\gamma_j = \bar{\upsilon} + \int_{\gamma_L}^{X} \bar{\kappa}(\gamma_j - \bar{\gamma})\xi(\gamma_j)d\gamma_j$$

Where $X < \gamma_H$ solves the expression $\bar{\upsilon} + \bar{\kappa}(X - \bar{\gamma}) = 1$ and $\xi(\gamma_j)$ is the probability density function of γ_j. It follows that the final part of the summation is negative, thus:

$$\mathrm{E}[w_j|\kappa = \bar{\kappa}] < \bar{\upsilon} = \mathrm{E}[w_j|\kappa = 0].$$

From (2.13) and the definition of α, it is clear that:

$$\frac{\partial \mathrm{E}[\pi_j(\boldsymbol{p})]}{\partial \sigma_\epsilon \partial w_j} > 0.$$

Assuming then that $\frac{1-\bar{\upsilon}}{\bar{\kappa}} < (\gamma_H - \bar{\gamma})$, it follows that: $\pi(c, \sigma_\epsilon^{'}, \sigma_\gamma, 0) - \pi(c, \sigma_\epsilon, \sigma_\gamma, 0) > \pi(c, \sigma_\epsilon, \sigma_\gamma, \bar{\kappa}) - \pi(c, \sigma_\epsilon^{'}, \sigma_\gamma, \bar{\kappa})$.

Proof of Proposition 2.8

When $\kappa = 0$, i's consumer surplus can be written as follows:

$$\mathrm{E}[\mathrm{CS}_i(\boldsymbol{x_i^*};\boldsymbol{p}^*)] = \sum_{j=1}^{m} \frac{1}{2}\upsilon \mathrm{E}[x_{ij}^*|w_{ij}=1](a+\frac{\theta}{b}\bar{\gamma}+\frac{c}{b}\upsilon(\sum_k p_k^* - a) - p_j^*)],$$

which follows because the realisation of each γ_{ij} is independent of the realisation of the observation probability and prices are non-random when $\kappa = 0$.

Suppose $\upsilon^* < 1$. First, note that $\mathrm{E}[x_{ij}^*]$ is increasing as $\upsilon < 1$ increases. Noting that $p_j^* = p^*$ for all j in this setting, the effect on expected demand of increasing υ can be written:

$$\frac{\partial \mathrm{E}[x_{ij}^*|w_{ij}=1]}{\partial \upsilon} = -b[\frac{p^*}{\partial \upsilon}] + c(m-1)(p^*-a) + c\upsilon(m-1)\frac{\partial p^*}{\partial \upsilon})]$$

As shown in Proposition 2, prices are decreasing in υ. Given that the matrix $(\boldsymbol{I} - \lambda R_S)$ is diagonally dominant:

$$b\frac{\partial p^*}{\partial \upsilon} > c\upsilon(m-1)\frac{\partial p^*}{\partial \upsilon}. \tag{2.14}$$

When $\theta < \bar{\theta}$, it must be that $p < a$. This follows because by definition if $\theta < \bar{\theta}$ then the following expression is decreasing in centrality:

$$\mathrm{E}[\boldsymbol{p^*}(\upsilon,\kappa)] = \frac{1}{2}(\boldsymbol{a} - \mathrm{E}[\boldsymbol{C}]\boldsymbol{a} + \theta \mathrm{E}[\boldsymbol{C\omega}]).$$

As $\mathrm{cov}(\boldsymbol{C},\boldsymbol{\omega})$ and $\mathrm{E}[\boldsymbol{\omega}] > 0$, it must be the case that $p^* < \frac{1}{2}a < a$. Hence, $(p^* - a) < 0$. It follows from this and (2.14) that $\frac{\partial \mathrm{E}[x_{ij}^*|w_{ij}=1]}{\partial \upsilon} > 0$.

Now consider the second component of the term inside the brackets of the above expression, $(a + \frac{\theta}{b}\bar{\gamma} + \frac{c}{b}\upsilon(\sum_k p_k^* - a) - p_j^*)$. By (2.14), the derivative of this expression with respect to υ is positive. Hence, $\mathrm{E}[CS_i(\boldsymbol{x_i^*};\boldsymbol{p}^*)]$ is increasing in $\upsilon < 1$ for all i, which proves the result.

Proof of Theorem 2.3

Recall that if $\theta < \bar{\theta}$, the expected equilibrium price vector $\mathrm{E}[\boldsymbol{p}^*]$ is decreasing in the centrality vector $\boldsymbol{C1_m}$. Suppose that $\frac{1-\bar{v}}{\bar{\kappa}} < (\gamma_H - \bar{\gamma})$ and $\frac{\bar{v}}{\bar{\kappa}} > (\bar{\gamma} - \gamma_L)$. In this case, $\sum_i \sum_j \mathrm{E}[w_{ij}|\kappa = \bar{\kappa}] = \sum_i \sum_j \mathrm{E}[w_{ij}|\kappa = 0]$.

Consider the expectation of the sum of i's demand for each seller's product conditional on κ_i, $\mathrm{E}[\sum_j x^*_{ij}(\boldsymbol{p^*})|\kappa = \kappa_i]$:

$$\mathrm{E}[\sum_j x^*_{ij}(\boldsymbol{p^*})|\kappa = \kappa_i] = \mathrm{E}[\sum_{j=1}(\bar{v}+\kappa_i(\gamma_j-\bar{\gamma}))(b(a-p^*_j(\bar{v},\kappa_i)+\theta\gamma_j+\sum_{k=1}\hat{c}_{jk}(p^*_k(\bar{v},\kappa_i)-a))]. \tag{2.15}$$

Define the random variable:

$$\tilde{x}_{ij} = (b(a - p^*_j(\bar{v},\kappa_i) + \theta\gamma_j + \sum_{k=1}\hat{c}_{jk}(p^*_k(\bar{v},\kappa_i) - a).$$

When $\kappa = 0$, (2.15) can be expressed:

$$\mathrm{E}[\sum_j x^*_{ij}(\boldsymbol{p^*})|\kappa = 0] = \bar{v}\mathrm{E}[\sum_{j=1}^{m}\tilde{x}_{ij}|\kappa = 0]$$

as cov($\bar{v}$,$\tilde{x}_{ij}$) $= 0$. Note that:

$$\mathrm{E}[\sum_{j=1}^{m}\tilde{x}_{ij}|\kappa = 0] = (m-1)[b(a - p^*(\bar{v},0) + \theta\bar{\gamma} + (m-1)\hat{c}(p^*(\bar{v},0) - a)].$$

$\mathrm{E}[p^*_j|\kappa = \bar{\kappa}] = \mathrm{E}[p^*_k|\kappa = \bar{\kappa}]$ in this setting as each seller is identical prior to the realisation of the matrix of the preference parameters $\boldsymbol{\gamma}$ and γ_j, γ_k are independent.

When $\kappa = \bar{\kappa} > 0$, cov($\sum_j \bar{v} + \bar{\kappa}(\gamma_j - \bar{\gamma}), \sum_j \tilde{x}_{ij}$) > 0. This holds for two reasons:

1. As $\sigma_\gamma > 0$, $\bar{\upsilon} + \kappa(\gamma_{ij} - \bar{\gamma})(\theta\gamma_{ij})$ is increasing in κ when $\frac{1-\bar{\upsilon}}{\bar{\kappa}} < (\gamma_H - \bar{\gamma})$;

2. As $(\boldsymbol{I} - \lambda R_S)$ is diagonally dominant, $\sum_j x^*_{ij}(\boldsymbol{p}^*)$ is decreasing in $\boldsymbol{p}^*$, which is decreasing in κ (assuming $\theta < \bar{\theta}$), as per the proof of Theorem 2.1.

Furthermore

$$\mathrm{E}[\sum_{j=1}^{m} \tilde{x}_{ij} | \kappa = \bar{\kappa}] = \sum_{j=1}^{m} [b(a - \mathrm{E}[p^*_j | \kappa = \bar{\kappa})]) + \theta\bar{\gamma} + \sum_{k=1}^{m} \mathrm{E}[\hat{c}_{jk}(p^*_k(\bar{\upsilon}, \bar{\kappa})] - a)],$$

which is greater than $\mathrm{E}[\sum_{j=1}^{m} \tilde{x}_{ij} | \kappa = 0]$ as $\mathrm{E}[p^*_j | \kappa = \bar{\kappa}] < \mathrm{E}[p^*_j | \kappa = 0]$ and $(\boldsymbol{I} - \lambda R_S)$ is diagonally dominant. Hence, $\mathrm{E}[\sum_{j=1}^{m} \tilde{x}_{ij} | \kappa = \bar{\kappa}] > \mathrm{E}[\sum_{j=1}^{m} \tilde{x}_{ij} | \kappa = 0]$and $\mathrm{cov}(\bar{\upsilon} + \bar{\kappa}(\gamma_j - \bar{\gamma}), \sum_j \tilde{x}_{ij}) > 0$, and thus:

$$\mathrm{E}[\sum_j x^*_{ij}(\boldsymbol{p}^*) | \kappa = \kappa_i] > \mathrm{E}[\sum_j x^*_{ij}(\boldsymbol{p}^*) | \kappa = 0].$$

Now consider the second term in the expression forCS_i:

$$\tilde{y}_{ij} := a + \frac{\theta}{b}(\gamma_{ij}) + \sum_{k=1}^{m} \frac{c}{b} w_{ik}(p^*_k - a)] - p^*_j.$$

By the same reasoning as used to analyse the function $\mathrm{E}[\sum_{j=1}^{m} \tilde{x}_{ij} | \kappa = \bar{\kappa}]$, it is clear that:

$$\mathrm{E}[\sum_{j=1}^{m} \tilde{y}_{ij} | \kappa = \bar{\kappa}] > \mathrm{E}[\sum_{j=1}^{m} \tilde{y}_{ij} | \kappa = 0].$$

Defining $y^*_{ij} := [\upsilon + \kappa(\gamma_j - \bar{\gamma})]\tilde{y}_{ij}$ and noting that $\mathrm{cov}(\bar{\upsilon} + \bar{\kappa}(\gamma_j - \bar{\gamma}), \sum_j \tilde{y}_{ij}) > 0$ by the same logic as above, we conclude that:

$$\mathrm{E}[\sum_j y^*_{ij}(\boldsymbol{p^*})|\kappa = \kappa_i] > \mathrm{E}[\sum_j y^*_{ij}(\boldsymbol{p^*})|\kappa = 0].$$

Finally, we note that:

$$\mathrm{cov}(\sum_j x^*_{ij}(\boldsymbol{p^*}), \sum_j y^*_{ij}(\boldsymbol{p^*})|\kappa = \bar{\kappa}) > \mathrm{cov}(\sum_j x^*_{ij}(\boldsymbol{p^*}), \sum_j y^*_{ij}(\boldsymbol{p^*})|\kappa = 0).$$

This holds because $\mathrm{cov}(C_j, C_k|\kappa = \bar{\kappa}) > 0$ for all j, k, while $\mathrm{cov}(C_j, C_k||\kappa = 0) > 0$ for all j, k. Hence the following two statements both hold:

$$\text{i) } \mathrm{E}[\sum_j x^*_{ij}|\kappa = \bar{\kappa}]\mathrm{E}[\sum_j y^*_{ij}|\kappa = \bar{\kappa}) > \mathrm{E}[\sum_j x^*_{ij}|\kappa = 0]\mathrm{E}[\sum_j y^*_{ij}|\kappa = 0]$$

$$\text{ii) } \mathrm{cov}(\sum_j x^*_{ij}, \sum_j y^*_{ij}|\kappa = \bar{\kappa}) > \mathrm{cov}(\sum_j x^*_{ij}, \sum_j y^*_{ij}|\kappa = 0).$$

Hence, it must be the case that:

$$\mathrm{E}[\sum_j x^*_{ij} y^*_{ij}|\kappa = \bar{\kappa}] > \mathrm{E}[\sum_j x^*_{ij} y^*_{ij}|\kappa = 0],$$

and so:

$$\mathrm{E}[\mathrm{CS}_i|\kappa = \bar{\kappa}] > \mathrm{E}[\mathrm{CS}_i|\kappa = 0].$$

This result holds for all i, and hence the result is proved for $\frac{1-\bar{v}}{\bar{\kappa}} < (\gamma_H - \bar{\gamma})$ and $\frac{\bar{v}}{\bar{\kappa}} > (\bar{\gamma} - \gamma_L)$.

Now consider the case where $\frac{\bar{v}}{\bar{\kappa}} > (\bar{\gamma} - \gamma_L)$. The preceding analysis still holds, but in this case, $\sum_j \mathrm{E}[w_{ij}|\kappa = \bar{\kappa}] > \sum_j \mathrm{E}[w_{ij}|\kappa = 0]$ for all i. As $\mathrm{E}[\mathrm{CS}_i]$ is increasing directly in $\sum_j \mathrm{E}[w_{ij}]$, and prices are decreasing in the probability that each seller is observed, it follows that the result that $\mathrm{E}[\mathrm{CS}_i|\kappa = \bar{\kappa}] > \mathrm{E}[\mathrm{CS}_i|\kappa = 0]$ for all i continues to hold when $\frac{\bar{v}}{\bar{\kappa}} > (\bar{\gamma} - \gamma_L)$.

Chapter 3

To Lead or To Follow? The Incentives of Party Leadership and Legislators in a Multi-Principal Environment

3.1 Introduction

A large body of literature (see, for example Fenno, 1978, Levitt and Synder, 1997 and Gowrisankaran, Mitchell, and Moro, 2004) documents and models how the incentive to be re-elected affects the effort allocation of members of Congress. Voters are more likely to re-elect politicians who have a reputation for putting effort into tasks in their constituency, such as town-hall meetings. Legislators therefore have an incentive to build a reputation for being such a politician by exerting more effort on tasks related to their constituency at the beginning of their time in office than they do at the end of it, an effect known as "frontloading" constituency effort.

However, less attention has been paid to the fact that legislators serve two masters:

the voter and leader of their party in Congress (e.g. the Senator Majority or Minority leader or the Speaker of the House of Representatives). Both principals are able to reward the politician. Voters incentivise politicians by voting them out if they do not spend time in their constituency, while party leaders are able to offer politicians rewards like committee positions or access to the executive branch. If the latter benefits are sufficiently valued by the politician, then they backload their constituency effort so as to enhance their reputation as an effective operator in the legislature, an effect explored in Charlson (2014) and supported empirically by e.g. Johannes (1980) and Sarbaugh-Thompson et al. (2004).

Modelling the party leader's role in the politician's time allocation problem identifies the complex nature of the party leader's incentives in this environment. The politician's trade-off between time allocations is internalised by the party leader in that, while they wish to reward politicians skilled at legislative work, they also want politicians of their own party to be re-elected. As a result, there are cases where the party leader would prefer to bind themselves to not providing the legislative benefit to ensure the politician increases their probability of re-election.

Moreover, the party leader would prefer to be able to commit to not offering the legislative benefit quite aside from their desire to see the politician of their own party re-elected. The party leader is a second-mover in the sense that they are only able to meaningfully reward the politician if the politician is re-elected. As a result, it is possible that providing the politician with a legislative reward results in an increase in constituency effort. To receive the legislative benefit, the politician needs to be re-elected, and hence such a benefit can incentivise increased constituency effort, even if, conditional on re-election, the probability of receiving that benefit is reducing in constituency effort allocation. The party leader therefore suffers from both a lack of commitment power and, potentially, second-mover disadvantage as well.

To model these effects, we assume that politicians allocate their time across constituency-based tasks (for example town hall meetings, district surgeries, etc.) and tasks that take place in the legislature (e.g. networking, voting etc).[1]

There are two types of politician: one who prefers to allocate their time to legislative business and one who prefers to allocate it to constituency-based tasks. Politicians have a term limit of two and are therefore a lame duck in their final period. Voters prefer politicians who spend effort on constituency-based tasks. The party leader values both the re-election of politicians of their party and the election of legislators who innately value putting effort into legislative projects, as this allows them to more easily pass high-quality legislation.

We show that whether a politician frontloads or backloads their effort depends on the relative value the politician places on the incentive offered by the two principals. If the value a politician places on the legislative rewards that the party leader offers is sufficiently large compared with the benefit of being re-elected, then they allocate more of their time to legislative tasks in their first term than in their second, which implies a backloading of constituency effort. A politician that values legislative rewards highly may therefore risk the greater chance of not being re-elected in order to increase the likelihood that they receive legislative benefits if they are re-elected.

We then explore the incentives of the party leader. As the utility the party leader gains from the politician's re-election increases, the optimal level of legislative incentive they wish to offer the politician decreases. Furthermore, the effectiveness of the legisla-

[1]There could potentially be some complementary between the two tasks. For example, including attempting to pass pork barrel legislation would benefit constituents and the party leader as well. The possibility of complementarity is ignored in this paper so that it can capture the trade-off that undoubtedly exists in general.

tive incentive decreases in the level of constituency effort exerted by the politician.

Our analysis therefore suggests that the party leader is strictly worse off because they lack commitment power. The party leader's lack of commitment power means that their choice of legislative reward does not take into account the fact that it distorts the politician's time allocation. As a result, the party leader tends to over-reward politicians who generate a low signal and under-reward politicians who generate a high signal. We find that the party leader is strictly worse off as a result of not being able to contract ex-ante with the politician.

We also find that the party leader may suffer from second-mover disadvantage. The case where the politician can only be rewarded if they are re-elected is compared with the case where they can be rewarded independently of re-election. If the benefit associated with high levels of legislative effort is large enough relative to the cost of inducing it, then the then party leader suffers from second-mover disadvantage. When the politician can only receive legislative benefits when they are re-elected, as legislative effort gets higher (and therefore the re-election probability gets smaller), the politician needs to be incentivised with a larger amount of legislative benefits. Providing such benefits is costly, which implies that the party leader is worse-off than in the case where they can provide benefits unconditionally.

However, if the costs of providing legislative benefits are high enough, the party leader actually prefers the case in which they move second. When the optimal level of legislative effort is relatively small, moving second can make the party leader's action more effective at increasing legislative effort than it would be in the case where the party leader does not move second.

3.2 Literature review

The distinction Fenno (1978) puts forward between the "expansionist phase" of a legislator's career, in which they build a reputation among voters in their constituency, and the "protectionist phase", which involves engaging in legislative tasks in Washington, is an early example of a reputational model of politician time allocation. The expansionist phase increases the probability of victory in future elections, allowing the politician to focus on their career in Congress as they gain seniority.

Formal treatments of political reputation include Besley and Case (1995), who show that when voters are uncertain about a politician's type, the legislator has an incentive to distort their action away from its optimal level in order to gain a reputation as a conscientious representative. Banks and Sundaram (1998) extend this model to the case where a long-lived principal interacts with a series of short-lived agents, who are retained only if they generate a cut-off level of surplus and find that the effort level of agents is reducing over time.

Political reputation models build on empirical research relating to constituency effort frontloading. Levitt and Synder (1997) find that politicians who provide constituency services early in their careers increase the likelihood that they are re-elected. They term the increase in re-election probability generated in this way as a "selection effect": incumbents are more likely to be re-elected if they are seen as a politician who values constituency effort. Cain, Freejohn and Fiorina (1987) and Gowrisankaran, Mitchell, and Moro (2004) both also provide evidence for this selection effect.

However, there is also evidence that some politicians backload their constituency effort. For example, Sarbaugh-Thompson et al. (2004) find that state legislators increase the amount of time they spend on helping constituents with problems in the period in which they are a lame duck. Johannes (1980) shows that constituency casework increases with seniority, suggesting that at least some members of Congress increase their constituency efforts later in their terms.

Another approach to examining the way in which re-election incentives distort the legislator's effort or time allocation builds on Holmstrom's (1999) "career concerns" framework. In the career concerns framework, the principal's payoff depends upon a noisy signal that in turn depends on the agent's type and their action, both of which are private information for the agent. The agent has an incentive to "signal-jam" to raise the principal's expectation of their type in future periods. Such a model differs from reputational models in that the principal's payoff depends directly on the agent's type, rather than on their action alone.

Persson and Tabellini (2000) and Ashworth (2005) apply the signal-jamming approach to a political reputation framework. Politicians distort their level of effort upwards in order to increase the voter's expectation about their type. Voters condition their voting decision on observed outcomes, and therefore as the number of periods of play increases, the fewer low-type politicians survive. High-type politicians engage in fewer costly constituency services, as the probability of producing a high signal is correlated with type.

Our analysis involves two principals interacting with a single agent who allocates their effort across two tasks. There are a number of more general analyses of multi-principal, single-agent models, including Bernheim and Whinston (1986), who characterise the conditions under which there exists a mechanism that induces an efficient action choice by the agent, Mastrolia and Ren (2018), who expand that result to a continuous time setting and Baverman and Stiglitz (1982), who apply the common agency framework to sharecropping. Lohmann (1998), Dixit, Grossman and Helpman (1997) and Gailmard (2009) apply the multi-principal, single-agent framework to a political environment.

A general treatment of multi-task principal-agent problems can be found in Holmstrom and Milgrom (1991). Hatfield and Padro i Miquel (2006) examines a case where a politician allocates their time across multiple tasks and find that because the only incentive voters can offer a politician is re-election, voters cannot incentivise a politician to allocate their effort efficiently across two tasks.

Charlson (2014) examines the time allocation of a politician whose action is distorted both by the re-election incentive and a fixed incentive offered by the party leader who has commitment power. The analysis there focused on the extent to which the multi-principal framework can explain why some politicians backload constituency effort. The model here examines the case where the party leader's incentive is continuous, and examines the extent to which a lack of commitment power and the fact that the party leader moves second affects the party leader's utility. The voter and politician's utility functions and the nature of the signals the politician's action generates are similar across the two environments.

3.3 The model

Politicians

Following Charlson (2014), the politician engages in two types of task: tasks which are beneficial to their party in the legislature (referred to here as "legislative tasks") and tasks which are beneficial to their constituents ("constituency tasks"). Politicians allocate their effort across up to two terms in office. As we will see, the politician's effort allocation will affect their re-election probability and the legislative rewards bestowed on them by the party leader. After choosing their effort allocation in the first term, the politician faces re-election before they are able to choose their action profile again in their second term. The challenger they face in the election is assumed to be of another party.

Politicians are either of type p ("populist") or type l ("legislative entrepreneur") with equal probability and are privately informed of their type. Politicians allocate effort between legislature tasks and constituency tasks, with types differing in respect to their first-best effort allocation.

Let $x_{i,t} \geq 0$ denote the level of constituency effort exerted in period t by a politician i and $y_{i,t} \geq 0$ be i's legislative effort level. A politician i's utility function in period t is as follows:

$$u_{i,t}(x_{i,t}, y_{i,t}; K_{i,t}, E_t, a_i) = \frac{1}{2}a_i x_t + \frac{1}{2}(1-a_i)y_t - \frac{1}{4}(x_t)^2 - \frac{1}{4}(y_t)^2 + zk_{i,t} + E_t$$

where $a_l < 0.5 < a_p < 1$ and $a_l + a_p \geq 1$. The value of E_t and K_t are contingent on the actions of the party leader. E_t is the election incentive, and can be expressed:

$$E_t = \begin{cases} e & \text{if } s_{t-1} = 1 \\ 0 & \text{otherwise} \end{cases},$$

where $s_t = 1$ indicates that the politician has been elected by the voter in period t. e is thus the extra utility the politician receives by being in office. $k_t \in \mathbb{R}_+$ is the amount of legislative benefit the politician is provided by the party leader and $z \in \mathbb{R}_+$ is the utility politician i receives from a unit increase in k_t.

We normalise the maximum amount level of a politician's effort to 1. Hence:

$$x_{i,t} + y_{i,t} \leq 1 \; \forall t.$$

In periods in which the politician does not move, $x_t = y_t = 0$ by definition. If the i loses the election, then they are removed from office, and a new politician, j is elected. The new politician is of either type with equal probability.

Voter

The voter chooses whether to vote for the incumbent or for a new politician. As in Charlson (2014), the voter observes their own utility and their utility function in period t can be written as follows:

$$m_t = \begin{cases} 1 & \text{if } q_t \ > q^* \\ \\ 0 & \text{if } q_t \leq q^* \end{cases},$$

where $q_t \sim U(0, \frac{1}{1-x_t})$. The voter observes the politician's effort only if the signal q exceeds some threshold q^*. As the politician exerts more effort in their constituency, the greater the probability that the threshold is exceeded. Specifically:

$$\Pr(q_t > q^*) = 1 - q^*(1 - x_t).$$

The voter therefore imperfectly observes the actions of the politician in the periods prior to the one in which they will vote, capturing the notion that voters are not perfectly attentive.

The voter's choice variable at election period t is s_t. If $s_t = 1$, then the voter re-elects the politician and if $s_t = 0$, they elect the politician from the other party, and remove the incumbent. There is only one election period in model (period 2), and so $s_t \equiv 0$ for all $t \neq 2$.

Party leader

The party leader is the leader of the party to which the incumbent politician belongs (e.g. the Speaker of the House of Representatives). The party leader values politicians of their own party being re-elected, the passage of legislation of good quality and the provision of legislative benefits to politicians who exert high legislative effort. The party leader's utility in period t is as follows:

$$n_t(k_t; \Psi_t, \Upsilon_t) = dy_{i,t}k_t - \frac{1}{2}ck_t^2 + \Psi_t + \Upsilon_t,$$

where $k_t \in \mathbb{R}_+$ is the amount of legislative benefit (e.g. the provision of a committee position, etc.) that the party leader provides to a politician of their own party, and $d > 0$ is the marginal benefit the party leader receives a from an extra unit of legislative effort when the legislative benefit level is k_t. The more legislative effort the politician engages in, the greater the utility the party leader gains from providing legislative benefits to them, capturing the notion that the party leader gains utility from politicians who put in legislative effort towards, for example, the committee position allocated to them.

Υ_t denotes the utility obtained by the party leader if the politician is re-elected:

$$\Upsilon_t = \begin{cases} \upsilon & \text{if } s_{t-1} = 1 \\ 0 & \text{otherwise} \end{cases},$$

and Ψ_t is a binary signal of the form:

$$\Psi_t = \begin{cases} \theta & \text{if } r_t > r^* \\ 0 & \text{if } r_t \leq r^* \end{cases},$$

where $r_t \sim U(0, \frac{1}{1-y_t})$. The functional form of the signal Ψ_t captures the notion that more legislative effort increases the likelihood the passage of good quality legislation but does not necessarily induce this outcome. As θ increases, the utility the party leader receives from the passage of high-quality legislation increases. The party leader only observes the outcome of the legislative process, and not the level of legislative effort directly.

If the party leader provides a politician with a non-zero level of legislative benefits, they do so after the incumbent has been re-elected. Furthermore, it will be assumed initially that neither principal has any commitment power. The party leader is therefore unable to credibly commit to providing the benefit upon the observation of a given signal ex-ante and only provides a benefit once they have observed the actions of the politician.

The party leader is assumed not to be able to reward politicians of a different party, so by definition if $s_t = 0$, $k_{t+1} = 0$. It is useful to define the party leader's overall level of utility, which can be expressed:

$$N(k_t; \Psi_t, \Upsilon_t) := \sum_{t=1}^{3} n_t(k_t; \Psi_t, \Upsilon_t).$$

Timing of the model

At the start of the game, the politician i's type is assigned, and they are type l or p with equal probability. The politician's type is private information. The model is assumed to have three periods, after which the game ends. These periods take the following form:

Period 1: The politician i chooses a level of $x_{i,1}$ and $y_{i,1}$, which generate q_1 and r_1.

Period 2: The voter whether to re-elect the incumbent or elect the challenger, i.e. they set $s_2 = 1$ or $s_2 = 0$.

Period 3: The party leader chooses whether to give the politician the legislative benefit and the elected politician chooses x_3 and y_3 and generates q_3 and r_3, after which, the game ends.

Note that we assume that the game ends in period 3 whether or not a new politician is elected in period 2. As the focus of our analysis is how the incentives of the politician, party leader and voter are affected when the politician can be re-elected, we assert that this abstraction is reasonable.

Equilibrium characterisation

We solve the game via backwards induction in the normal manner.

Period 3: The politician's strategy

The politician cannot be re-elected as the game ends in period 3. The principals are not able to affect the effort allocation of the politician in period 3. The politician therefore chooses the effort allocation that optimises their instantaneous utility, which, noting that the their time constraint binds as $a_l + a_p \geq 1$, implies that the optimal effort allocation is as follows:

$$x^*_{i,3} = a_i$$

$$y^*_{i,3} = 1 - a_i.$$

The optimal level of constituency effort for type p politicians is greater than for type l politicians, while type l politicians exert more legislative effort in equilibrium.

Period 3: The party leader's strategy

The party leader knows that the politician has been re-elected in period 2, and therefore updates their prior belief about the politician's type. Define:

$$\sigma_i := \Pr(\text{politician is of type } i | s_2 = 1).$$

By Bayes rule, the party leader's belief that the politician is of type i when $r_1 > r^*$ can be written:

$$\frac{\Pr(r > r^*, s_2 = 1 \text{ and type } i)}{\Pr(r > r^* \text{and } s_2 = 1)} = \frac{(1-r^*(1-y_{1,i})\sigma_i}{\sum_i \sigma_i(1-r^*(1-y_{1,i}))}$$

Noting that the belief that politician is of type i if $r_1 < r^*$ can be expressed similarly, the party leader's belief that the politician's type is l, α, can be denoted:

$$\alpha = \begin{cases} \frac{(1-r^*(1-y_{1,l})\sigma_i}{\sum_i \sigma_i((1-r^*(1-y_{1,i}))} & \text{if } r_1 > r^* \text{ and } s_2 = 1 \\ \frac{r^*(1-y_{1,l})\sigma_i}{\sum_i \Pr\sigma_i(1-r^*(1-y_{1,i}))} & \text{if } r_1 \leq r^* \text{and } s_2 = 1 \\ 0.5 & \text{if } s_2 = 0 \end{cases} .$$

We assume that the party leader is only able to provide legislative incentives to an incumbent politician, motivating this assumption by noting that in general if there is new politician elected, they are of a different party to the party leader. Assuming the incumbent is re-elected, and given the politician's strategy in period 3, the party leader maximises their utility by choosing k (for ease of notation we drop the t subscript here as the party leader only sets k once in the model) that solves the following:

$$\max_k\{dk(\alpha(1-a_l) + (1-\alpha)(1-a_p)) - \frac{1}{2}ck^2\}.$$

Hence:

$$k^* = \begin{cases} \frac{d}{c}(\alpha(1-a_l) + (1-\alpha)(1-a_p)) & \text{if } s_2^* = 1 \\ 0 & \text{if } s_2^* = 0 \end{cases} .$$

We define as k_H^* as the optimal level of k when the politician is re-elected and $r_1 > r^*$

and k_L^* as the optimal level of k when the politician is re-elected and $r_1 < r^*$.

Period 2: voter's strategy

The voter's signal in period 1 is informative about the politician's type, and they update their belief accordingly. As there is initial probability that the politician is of either type is equal, by Bayes rule the voter's belief that a politician is of type p if $q_1 > q^*$ is as follows:

$$\Pr(\text{type}{=}p \mid q_1 > q^*){=}\frac{\Pr(q > 1 \text{ and type } p)}{\Pr(q > 1)} = \frac{0.5(1-q^*(1-x_{1,p}))}{0.5(1-q^*(1-x_{1,p}))+0.5(1-q^*(1-x_{1,p}))}$$

It follows that, if $\mathrm{E}[x_1^*|p] \geq \mathrm{E}[x_1^*|l]$ (which holds in equilibrium; see below), the voter's belief that the politician is of type p is above 0.5 iff q_1 is greater than the threshold q^*. If $q > q^*$, the incumbent is more likely to be of type p than than their challenger (who is type p with probability 0.5). The voter's utility of re-electing the incumbent (i.e. setting $s_2 = 1$) is $\mu_2 a_p + (1 - \mu_2)a_l$. This is weakly preferred to setting $s_2 = 0$ when $\mu_2 \geq 0.5$, which implies that the voter's optimal strategy is as follows:

$$s_2^* = \begin{cases} 1 & \text{if } q_1 > q^* \\ 0 & \text{if } q_1 \leq q^* \end{cases}.$$

Period 1

Given the strategies outlined in periods 2 and 3, a politician i's chooses to optimise their utility in period 1. Their optimisation problem is as follows:

$$\max_{x_{i,1},y_{i,1}} \{\frac{1}{2}a_i x_{i,1} + \frac{1}{2}(1-a_i)y_{i,1} - \frac{1}{4}(x_{i,1})^2 - \frac{1}{4}(y_1)^2 +$$
$$[1 - q^*(1 - x_{i,1})](1 - r^*(1 - y_{i,1})(zk_H + e + f_i) + r^*(1 - y_{i,1})(zk_L + f_i + e))\},$$

where f_i is the utility a type i politician receives from their optimal effort allocation in period 3, and is therefore equal to$\frac{1}{4}a_i^2 + \frac{1}{4}(1-a_i)^2$. The above maximisation problem is subject to the resource constraint:

$$x_{i,1} + y_{i,1} \leq 1,$$

which must bind in equilibrium, as $a_l + a_p \geq 1$. Thus, i's optimisation problem can be restated:

$$\max_{x_{i,1}}\{\frac{1}{2}a_i x_{i,1} + \frac{1}{2}(1-a_i)(1-x_{i,1}) - \frac{1}{4}(x_{i,1})^2 - \frac{1}{4}(1-x_{i,1})^2 +$$
$$[1 - q^*(1-x_{i,1})](zk_H + V_i - r^* x_{i,1}\tilde{k})\},$$

where $\tilde{k} := k_H - k_L$ and $V_i := e_i + f_i$. Define the maximand of this problem as $U_i(x_{1.i})$. As $U_i(x_{1.i})$ is concave in $x_{1,i}$, it follows that the solution to the first-order condition:

$$U_i^{'}(\bar{x}_{1.i}^*) = 0$$

defines a unique equilibrium if the solution is interior (i.e. if $0 < x_{i,1}^* < 1$). Let $\bar{x}_{i,1}^*$ and $\bar{y}_{i,1}^*$ denote this interior solution to the party leader's optimisation problem. Then the following must hold for the optimal allocation in period 1:

$$x_{i,1}^* = \min\left\{\bar{x}_{i,1}^*,\ 1\right\};$$

$$y_{i,1}^* = \max\left\{1 - \bar{x}_{i,1}^*,\ 0\right\},$$

where:

$$\bar{x}_{i,1}^* = \frac{a_i + q^*(zk_H + V_i)}{1 + 2z\tilde{k}r^*}.$$

Call the above game Γ. Theorem 3.1 summarises the analysis in the previous sections formally.

Theorem 3.1. *The game Γ outlined above has the following unique Perfect Bayesian Equilibrium:*

***Period 1**: The politician sets $x_{i,1} = x^*_{i,1}$ and $y_{i,1} = y^*_{i,1}$*

***Period 2:** The voter re-elects the politician iff $q > q^*$, with the belief that the politician is of type p equal to μ.*

***Period 3**: If the politician is re-elected, the party leader's belief that the politician is of type l is α and $k^* = \frac{d}{c}(\alpha(1-a_l) + (1-\alpha)(1-a_p))$ if $s_2 = 1$ and 0 otherwise. If the politician is re-elected sets $x^*_{i,3} = a_i$ and $y^*_{i,3} = 1 - a_i$. If a new politician of type j is elected, they set $x^*_{j,3} = a_j$ and $y^*_{j,3} = 1 - a_j$.*

An example, Γ_1

We will use a slightly simplified form of the game Γ in Section 5, which we call Γ_1. We assume that Γ_1 has the same basic set-up as Γ, but that some of the parameters take on specific values. Specifically, we assume that: $q^* = r^* = 1$, $\upsilon = 0$, $e_i = e$ for all i and $a_p = 1 - a_l$.

These assumptions imply the following about the probability that the two binary signals, r_t and q_t are above their threshold level:

$$\Pr(r_t > 1) = y_1$$

and:

$$\Pr(q_t > 1) = x_1.$$

The voter's belief that the politician is of type p when $q_1 > 1$ is as follows:

$$\Pr(\text{type=}p \mid q_1 > q^*) = \frac{\Pr(q > 1 \text{ and type } p)}{\Pr(q > 1)} = \frac{x^*_{p,1}}{x^*_{p,1} + x^*_{l,1}}.$$

The above expression implies that the party leader's belief if they observe $r_1 > 1$ and the politician is re-elected is as follows:

$$\Pr(\text{type}{=}l \mid r_1 > r^*) = \frac{y^*_{l,1}\frac{x^*_{l,1}}{x^*_{p,1}+x^*_{l,1}}}{y^*_{l,1}\frac{x^*_{l,1}}{x^*_{p,1}+x^*_{l,1}} + y^*_{p,1}\frac{x^*_{p,1}}{x^*_{p,1}+x^*_{l,1}}.}.$$

Then the solution to the politician's maximisation problem yields:

$$x^*_{i,1} = \min\left\{\frac{a_i + zk_H + V}{1+2z\tilde{k}},\ 1\right\};$$

$$y^*_{i,1} = \max\left\{1 - x^*_{i,1},\ 0\right\},$$

where $V = e + f$ and $f = \frac{1}{4}a + \frac{1}{4}(1-a) = \frac{1}{4}$. The politician's constituency effort is increasing in their innate preference for re-election (e) and the benefit to re-election associated with being able to choose their optimal level of effort allocation in their second term (f). Whether constituency effort is increasing or decreasing in the legislative benefit is ambiguous, which discuss in more detail below.

3.4 Analysis

The time allocation of the politician

As noted in the introduction, some politicians frontload their constituency effort, whereas others backload it. The model set out in Section 3 can explain both of these observations. A legislator frontloads their constituency effort effort if $x^*_{i,1} > x^*_{i,4}$ and backloads it if the reverse inequality holds. Whether politicians backloads or frontloads legislative effort depends on the legislative benefit, z. The following proposition relates the relative level of legislative effort across the two periods:

Proposition 3.1. *Suppose Γ is played. For any finite value of the electoral benefit,*

$V_i > 0$, there exists a $\bar{z}(V_i)$ such that if $z \leq \bar{z}(V_i, r^*)$, then $x^*_{i,1} \geq x^*_{i,4}$ and if $z > \bar{z}(V^*_i)$, then $x^*_{i,1} \leq x^*_{i,4}$.

When the legislative benefit is large it is preferable for the politician to frontload their legislative effort in order to increase the probability that they will be provided with the legislative benefits. If the politician values legislative benefits relatively highly, then they will attempt to build a reputation as a legislative entrepreneur by engaging in more legislative effort in their first term in office.

On the other hand, if the utility associated with re-election is relatively large compared with the value the politician receives from legislative benefits, then the politician will allocate more constituency effort to their first term, in order to increase the probability of re-election. This result is consistent with Charlson (2014), which explores the politician's optimal time allocation when the party leader is able to commit to offering a fixed legislative benefit.

The legislator exerts high legislative effort in their first term if the value they place on legislative benefits is relatively large compared with the re-election incentive, occurs if z is large enough relative to V_i. This implies the following proposition:

Proposition 3.2. *Suppose Γ is played. The threshold marginal value on legislative benefits, $\bar{z}(V_i, r^*)$ is increasing in V_i and decreasing in r^*.*

As V_i increases, the incentive to allocate resources towards re-election straightforwardly increases, which implies that for a politician to backload their constituency effort they would need value legislative benefits more to offset the direct benefit to re-election.

As r^* increases, there is a greater incentive for the politician to engage in legislative effort because the party leader requires there to be a greater level of legislative effort to receive a positive signal. Hence, the threshold level of the marginal value of legislative benefits must be higher.

The party leader without commitment

One key element of the model and the real-life interaction it captures is the fact that the party leader is only able to provide the politician with a reward if they are re-elected. This shapes the party leader's ability to influence the politician's decision making and the utility they receive from being able to offer legislative rewards.

The party leader has an incentive to ensure that politicians of their own party are re-elected, an incentive known in the literature as "electoral maintenance" (Pearson, 2005). At the same time, the party leader also has a desire to pass high-quality legislation, which is more likely when politicians who exert relatively high levels of legislative effort are in office.

The desire for electoral maintenance may mean that the party leader's utility would increase if they were able to commit to not providing a legislative benefit to politicians at all. To see this, define $\hat{x}_{i,1}$ such that it solves the politician's first-order condition when $k_H^* = k_L^* = 0^*$ and hence $\tilde{k} = 0$:

$$a_i + V_i = \hat{x}_{i,1}.$$

For now, assume $x_{1,i}^* < \hat{x}_{1,i}$. When this holds the party leader would face a trade-off in the case where they can commit not to provide legislative benefits. When the politician spends on legislative tasks in period 1, the party leader reaps immediate benefits allows the party leader to give electoral benefits to politicians who provide a high signal. At the same time, the politician is less likely to be re-elected if they spend more time on legislative work, which is has a direct utility cost to the party leader.

The above trade-off leads to the following proposition relating to v, the benefit to the party leader associated with the re-election of the politician, and the party leader's desire to commit to not providing a legislative benefit at all. Let $N(k_L, k_H)$ denote the party leader's expected utility when they commit to a payment schedule such that if $r < r^*$ they give the politician k_L and if $r > r^*$ they give the politician k_H. The following proposition holds:

Proposition 3.3. *Suppose Γ is being played, $d > 0$ and that $x^*_{1,i} < \hat{x}_{1,i}$. There exists a threshold level of υ, $\bar{\upsilon}$, the benefit to the party leader associated with the re-election of the politician, such that if $\upsilon > \bar{\upsilon}$ then $N(0,0) > N(k^*_L, k^*_H)$.*

If the benefit to the party leader associated with an incumbent being re-elected is high enough, then the party leader's outweighs their motive to incentivise legislative effort, and they would prefer to be able to commit to not providing the legislative benefit. In addition to the effects above, there are set-ups of the game in which the legislative benefit does not incentivise increased legislative effort for all types of politician. To see an example of this, consider the game Γ_c at the end of Section 2 and note that $x^*_{1,i} \leq \hat{x}_{1,i}$ when:

$$\frac{a_i + (zk^*_H + V)}{1 + 2z\tilde{k}} \leq a_i + (V + k^*). \tag{3.1}$$

The inequality in (3.1) implies the following proposition:

Proposition 3.4. *Suppose the game Γ_1 is played. If $2a_l + \frac{1}{2} < 1$, then there exists a value of $\bar{e} \in \mathbb{R}_+$, such that if $e < \bar{e}$ then the party leader's offer of k_H reduces the amount of legislative effort of the l type in period 1 compared to the case where k_H is not offered; that is, $x^*_{1,i} \leq \hat{x}_{1,i}$.*

Proposition 3.4 implies that if the benefit of re-election is sufficiently low, legislative benefits incentivises the l type politician to increase their effort because they can only receive the benefit if they are re-elected. As the direct benefit of re-election becomes larger, the legislative reward becomes less influential in the politician's decision to get re-elected, and more influential in allocating their time towards legislative tasks in order to receive the electoral benefit.

Hence, even if the party leader values incentivising legislative effort, it may be optimal for them to commit to not offering a lower level of legislative benefit than is implied by the equilibrium of Γ_1. However, there is no mechanism that allows them to do so. Once the incumbent is re-elected, and if their signal, $r_1 > r^*$, then the party leader only has an incentive to give the politician the legislative benefit. Any promise not to provide the benefit would be a non-credible commitment and therefore cannot be a part of equilibrium play.

Propositions 3 and 4 suggest that the party leader suffers from both a lack of commitment power and from second-mover disadvantage. In terms of the latter, the party leader can only provide a benefit to the politician if they are re-elected. Hence, there are cases where they would prefer not to incentivise the politician to put more weight on legislative tasks if it were possible to do so. As Proposition 3.4 makes clear, this may be true even if the party leader does not explicitly put value on the incumbent's re-election: as the politician's payoff is contingent on re-election, the party leader's incentive may actually lower legislative effort, and therefore the party leader's payoff.

3.5 Commitment

In this section, we will consider the case where the party leader is able to commit to a contract such that, prior to period 1, the party leader commits to rewarding the politician based on the realisation of the party leader's signal. We will use this abstraction to examine two ways in which the party leader's decision is constrained. We will examine the fact that the party leader: (a) lacks the ability to commit; and (b) can only reward politicians who are re-elected.

It will be assumed throughout that different versions of game Γ_1 are played. The parameter values will be the same as this game, but we will assume that the commitment power of the party leader and/or their ability to reward the politician after they lose an election are different to that game.

Contracting between the party leader and the politician

To examine the party leader's commitment problem, we will consider the case where the party leader is able to commit to rewarding the politician contingent on their signal, but only if the politician is re-elected. Formally, the party leader is able to contract with the politician such that they will provide the politician with a contract of the following form: (if $r < r^*$ and $s_2^* = 1,$ then pay k_L); (if $r \geq r^*$ and $s_2^* = 1,$ then pay k_H). We assume that the parameters of the model are the same as the equivalent parameters in game Γ_1. We will call this version of the game Γ_c.

We first consider the case where $d = 0$. Then, the party leader only values the outcome of their signals Ψ_1 and Ψ_4, and does not innately value providing legislative benefits to the l type. The party leader's optimal strategy solves the following problem:

$$\max_{k_L,k_H}\{N(\boldsymbol{k};\boldsymbol{x}(\boldsymbol{k},\Gamma_c)\}$$

where $\boldsymbol{k} = (k_L, k_H)'$ and $\boldsymbol{x} = (x_{l,1}, x_{p,1})'$. The derivative of the party leader's utility function with respect to k_j, $j \in \{L, H\}$can be stated as follows:

$$\text{marginal effect on effort}_{k_j} + \text{marginal effect on selection}_{k_j} + \\ \text{marginal cost of re-election}_{k_j} - ck_j.$$

The marginal effect on effort term captures the fact that changing k_H and k_L changes the optimal level of effort the politician allocates to legislative tasks. The marginal effect on selection is tied to the first effect, in that by changing the equilibrium effort the politician allocates to constituency tasks, the party leader changes the probability that a given type will be re-elected. This effects the party leader's utility in period 3, as the l type puts in more legislative effort in this period than the p type. The marginal cost of re-election reflects the fact that re-electing a politician requires the party leader to provide them with legislative benefits, which, in this setting, provides no utility to the party leader directly and is thus costly.

Note that the derivative of $x_{i,1}$ with respect to k_H is reducing in V. Define $\tilde{V}$ as follows:

$$\tilde{V} := a_l - \frac{z}{2}.$$

It is easy to verify that if $V > \tilde{V}$, then $\frac{\partial x_{l,i}}{\partial k_H} < 0$ for both i. The following proposition characterises the party leader's optimal strategy when $d = 0$:

Proposition 3.5. *Suppose Γ_c is being played, that $V > \tilde{V}$ and that $V + a_p < 1$. When $d = 0$, in any solution to the party leader's maximisation problem, $k_L^{'} = 0$ and $k_H^{'} > 0$.*

The assumption that $\tilde{V} + a_p < 1$ guarantees that $x_{l,i}$ is reducing in k_H when $k_H = 0$, as it rules out $x_{1,i}$ being a corner solution for both types if no legislative benefit is offered.

Setting $k_L^{'} > 0$ is costly for three reasons. Firstly, it reduces the incentive to engage in legislative effort. The politician receives at least $k_L^{'}$ upon re-election. Hence, as $k_L^{'}$

increases, the re-election incentive increases and the cost to putting in less legislative effort decreases. This increases the politician's optimal level of constituency effort, reducing their legislative effort.

Secondly, increasing k_L increases the constituency effort level of p types more than it does for l types. Therefore, increasing $k_L^{'}$ above 0 increases the probability that a p type will be re-elected, which reduces the party leader's payoff in period 3. As the party leader's payoff is decreasing in k_L due to both of these effects, it follows that $k_L^{'} = 0$.

Finally, increasing k_L increases the probability of re-election, which is straightforwardly costly when $d = 0$, as the party leader must reward the politician without receiving any direct benefit.

There is an incentive to set $k_H^{'}$ above zero for similar reasons to $k_L^{'} = 0$. Increasing $k_H^{'}$ increases legislative effort because it incentivises the politician to increase the probability that $r > r^*$. Furthermore it is clear that:

$$|\frac{\partial x_{p,1}}{\partial k_H}| > |\frac{\partial x_{l,1}}{\partial k_H}|. \tag{3.2}$$

Increasing the benefit associated with a high r decreases the constituency effort (and therefore increases the legislative effort) of p type politicians more than l types. Hence, the probability that a p type will be re-elected reduces relative to the probability that an l type will be re-elected.

Recall that in game Γ_1, where commitment is not possible, the following holds:

$$k^* = \frac{d}{c}(\alpha(1 - a_l) + (1 - \alpha)(1 - a_p)). \tag{3.3}$$

Therefore, as $d = 0$ here, $k_L^* = k_H = 0$. Define the equilibrium of the game where the party leader has no commitment power as the "no commitment equilibrium". Define:

$$N_i = \max_{k_L, k_H}\{N(\boldsymbol{k}; \boldsymbol{x}(\boldsymbol{k}, \Gamma_i)\} \; i \in \{1, c\}.$$

Proposition 3.6 relates the payoffs in of the no commitment equilibrium and the commitment equilibrium:

Proposition 3.6. *If $d = 0$ and $V > \tilde{V}$, then $N_c > N_1$.*

When $d = 0$, the party leader always prefers the case where they are able to commit, assuming that increasing k_H increases the legislative effort of the politician. Commitment allows the party leader to take into account the fact that their legislative benefits decision influences the action profile of the politician. In the no-commitment case, the party leader can only provide the legislative benefit according to the expected benefit resulting from the appointment, which in this case is zero. Thus, when $d = 0$, the party leader underprovides k_H in the no-commitment case.

Now consider the case where $d > 0$. The derivative of the party leader's utility function with respect to each k_j has three additional terms, which are as follows:

$$\text{marginal effect on re-election}_{k_j} + \text{marginal effect on screening}_{k_j} + \\ \text{marginal effect of appointment}_{k_j}.$$

The marginal effect on re-election term captures the party leader's only being able to receive benefits to appointment if the politician is re-elected. As has been shown, changing k_H and k_L affect the probability that the incumbent is re-elected. The party leader values re-election when $d > 0$ because they can only provide legislative benefits if the incumbent is re-elected. The probability of re-election is increasing in k_L and ambiguous in k_H: if $V < \tilde{V}$, then increasing k_H increases the probability of re-election; otherwise increasing k_H decreases the probability of re-election.

The party leader's choice of k_H and k_L also has a screening effect, in that the choice of these variables has an effect on their ability to distinguish between l type and p type politicians, conditional on re-election. The party leader only wants to provide legislative benefits to l types, and hence prefers the case where the signal they receive is maximally informative, all else equal. As shown above, an increase in k_H (and k_L) has different marginal effects on $x_{1,p}$ and $x_{1,l}$, which implies that changing k_j affects the party leader's ability to screen.

More specifically, conditional on re-election, increasing k_H increases the relative likelihood that a signal of $r > r^*$ comes from a p type rather than an l type. Hence, conditional on a politician being re-elected, increasing k_H reduces the party leader's ability to screen politicians and makes it more likely that they provide the p type politician with a high level of legislative benefits, which is costly. For the opposite reason, increasing k_L is beneficial for screening: the absolute value of the derivative of $y_{1,p}$ with respect to k_L is larger than the same derivative of $y_{1,l}$, and both are negative, implying that the relative likelihood that a positive signal is from a l type increases.

The last of the additional terms in the first-order condition relates to the benefit associated with providing legislative benefits to the politician. The party leader receives an expected payoff of providing a benefit of k to the politician equal to dk_j times the probability that k_j is offered. Notice that this term is positive in both first-order conditions - the party leader has an incentive to provide some legislative benefits to politicians even if they produce a low signal. Low signal producing politicians still may be of type l, and hence there is a payoff associated with providing them with some legislative benefit.

It is clear from the functional form of the party leader's payoff function that there exists a $\tilde{c} \in \mathbb{R}_+$ such that if $c > \tilde{c}$ then $N(\boldsymbol{k}; \boldsymbol{x}(\boldsymbol{k}, \Gamma_i))$ is concave in the vector $\boldsymbol{k}$. The analysis above implies the following theorem, recalling that k_j^* denotes the optimal level of k_j when the party leader does not have commitment power:

Theorem 3.2. *Suppose that $V > \tilde{V}$, $a_p + V < 1$ and $c > \tilde{c}$. Then there exists a unique solution to the party leader's maximisation problem, $(k_L^{'}, k_H^{'})$. Furthermore, then if $k_j^* = k_j^{'}$ then $k_i^* \neq k_i^{'}$ for $i \neq j$. Hence, $N_c > N_1$.*

If increasing k_H from zero reduces the each type's level of constituency effort (which is implied by $a_p + V < 1$), the party leader strictly prefers to have commitment power, assuming $V > \tilde{V}$. Given the trade-offs highlighted by the discussion of the first-order conditions of the party leader's maximisation problem, it is possible that, for example, $k_H^{'} = k_H^*$. However, if this is the case then it must be that $k_L^* \neq k_L^{'}$. Hence, the no-commitment equilibrium is never identical to the the commitment equilibrium. As the party leader could choose (k_L^*, k_H^*) in the commitment case but chooses not to, it follows that the commitment equilibrium is always strictly preferred to the no-commitment equilibrium.

Second-mover disadvantage

To focus the discussion on the second-mover disadvantage faced by the party leader, we assume they have commitment power and $d = 0$. We will compare Γ_c, which was discussed in the previous section, with the case where the party leader is able to provide legislative benefits to the politician whether or not the voter re-elects them.[2] The contract that the party leader offers to the politician in this case can be stated: (if $r < r^*$, then pay k_L); (if $r \geq r^*$, then pay k_H). Call the game where the party leader can offer k_H and k_L to the politician regardless of the election outcome Γ_2.

[2] This abstraction makes it possible to understand cases where the party leader has a second-mover disadvantage - clearly the party leader would be unlikely to be able to provide such a benefit if the incumbent were to lose the election.

By the same argument as the one made for Proposition 3.6, the party leader will only provide the politician a non-zero level of legislative benefit if $r_1 > r^*$, that is, $k_L(\Gamma_i) = 0$ for $i \in (c, 2)$. The solution of the optimisation problem of the incumbent is as follows:

$$x_{i,1}(0, k_H, \Gamma_2) = \max\{a_i - zk_H + V, 0\}. \tag{3.4}$$

The party leader's optimisation problem in game Γ_i can then be stated:

$$N_i = \max_{k_H}\{N(\boldsymbol{k}, \boldsymbol{x}(\boldsymbol{k}; \Gamma_i))\}\ i \in \{c, 2\}.$$

The analysis above showed that the party leader's first-order conditions when game i is played can be stated:

$$\text{Effort}_{k_j^*(\Gamma_i)} + \text{Selection}_{k_j^*(\Gamma_i)} + \text{Re-election}_{k_j^*(\Gamma_i)} = ck_j^*(\Gamma_i).$$

It can be shown that the selection effect and the re-election effect in Γ_2 are both equal to zero. It is clear from (3.5) that the marginal effect of k_H is not dependent on a_i, that is:

$$\frac{\partial x_{l,1}(k_L, k_H; \Gamma_2)}{k_H} = \frac{\partial x_{p,1}(k_L, k_H; \Gamma_2)}{k_H}.$$

The selection effect identified above is driven by the party leader's ability to decrease the probability that a politician of type p is elected more than they decrease the probability that a type l is re-elected. In the case where the party leader is able to reward the politician is either way, the marginal effect of legislative benefits is the same for both types and therefore the selection effect is zero.

Similarly, as the party leader does not move second in game Γ_2, they must provide the politician with legislative benefits whether or not they are re-elected. Hence, the re-election effect, which is driven by the party leader's incentive to increase the probability that the politician fails to get re-elected so that they do not have to provide them with the legislative, is zero in the case where the party leader does not move second.

Hence, there are benefits to the party leader from moving second, in that they are able to affect the selection of candidates by the electorate with their choice of k_H, and there is some probability that they can commit to providing the politician with a legislative benefit they do not then provide.

However, by not moving second, the party leader is potentially able to affect the politician's effort allocation at a reduced cost compared to in game Γ_c. The derivative of $x_{i,1}$ with respect to k_H in the case where the party leader moves second is as follows:

$$\frac{\partial x_{i,1}(0, k_H; \Gamma_c)}{\partial k_H} = \frac{z(1 - 2(a_i + V))}{(1 + 2zk_H)^2}.$$

When V is low, or k_H^* is sufficiently high, the marginal effect of increasing k_H on legislative effort in Γ_c is lower than z, the marginal effect of k_H on legislative effort in Γ_2. In fact it is possible that if a_l and V are low enough, increasing k_H actually reduces the legislative effort of type ls, though this is not possible for type p politicians. Second-mover disadvantage arises in cases where the party leader is able to induce a high probability of receiving a positive signal at a lower cost than in the case they move second.

The party leader is thus more likely to suffer from second-mover disadvantage in cases where the optimal level of k_H is relatively high in both games, which arises because, for example, the benefit associated with a positive realisation of the party leader's signal, θ, is relatively high compare with the cost of inducing effort, c. When the benefit to a positive realisation of the signal is large, the party leader wishes to induce a high level of $y_{i,1}$ for both types. In the case where they move second, the party leader's ability to increase $y_{i,1}$ becomes increasingly costly, because it becomes increasingly unlikely that the politician will receive the reward being offered to them.

The same is not true in game Γ_2, where the party leader can induce a legislative effort level of 1 for both types by setting $k_H^* \geq \frac{a_p - V}{z}$.

Theorem 3.3 formalises some of the above discussion:

Theorem 3.3. *Suppose that $d = 0$. For all V and $c > 0$, there exists a $\bar{\theta} \in \mathbb{R}_+$ such that if $\theta > \bar{\theta}$, then $N_c < N_2$: when these assumptions hold, the party leader suffers from second-mover disadvantage.*

As the payoff the party leader receives from a positive signal increases, the greater the loss associated with the fact that when the party leader moves second, $y_{i,1}$ is concave in k_H. If the politician can only be rewarded if they are re-elected, it becomes increasingly more costly for the party leader to increase the probability of receiving a positive signal because as $y_{i,1}$ increases the probability that the reward is received by the politician decreases. As θ increases, this disadvantage of moving second dominates the selection and re-election effects, and thus the party leader suffers from second mover disadvantage.

Hence if the party leader values the passage of high-quality legislation highly enough, then they suffer from second-mover disadvantage. As they are not necessarily able to reward a politician who has exerted the effort required to pass legislation, the party leader is unable to incentivise the politician to exert such effort as much as they would like to.

However, if the optimal level of k_H is relatively low in both games, then the party leader may actually prefer that they move second. For a start, when k_H is low, the selection and re-election effects are relatively large. In addition, $\frac{\partial x_{1,i}(0,k_H;\Gamma_c)}{\partial k_H}$ is relatively large when k_H is small, which implies that k_H is relatively effective at increasing $y_{i,1}$. In fact, if $V > 1 - a_l$, then moving second is actually more effective at decreasing the legislative effort of both types for some positive values of k_H.

The above analysis implies that as c increases, the party leader's second-mover disadvantage is reducing, and in fact the party leader may actually prefer moving second if c is sufficiently large. We do not examine this possibility explicitly here.

3.6 Discussion

Introducing a second principal into the canonical model of political reputation not only explains why politicians may frontload or backload time allocated to constituency work, but also illustrates the incentive structure of the party leader. The party leader generally lacks commitment power and moves second, both of which potentially leaves them worse off.

If the party leader is unable to contract with the politician, they are unable to internalise the distortive effect that legislative incentives have on the politician's time allocation. The time allocation of the politician affects the probability that the politician is re-elected, which in turn affects the party leader's ability to gain utility from providing politicians with legislative benefits. We find that the party leader is always strictly worse off as a result of not having the ability to commit as a result of these effects.

A further avenue of research would therefore be to examine the extent to which legislative leaders and parties more generally are able to commit to offering legislative appointments or other benefits to politicians. While party leader's are not able to directly contract with politicians, the fact that they repeatedly interact with politicians in a dynamic setting suggests that they have reputational incentives of their own with regards to promises made to politicians of their party. Analysing what shapes incentives and how they may affect the party leader's ability to in turn incentivise time allocation towards legislative tasks would be an fruitful line of inquiry.

Moving second means that the party leader's ability to incentivise the politician is tempered by the fact that the politician can only reap any benefits provided by the party leader if they are re-elected. This implies that, assuming the politician's initial level of time allocated to the constituency is not very high, the party leader is less able to distort effort towards legislative tasks, which makes the party leader worse off when the benefit to receiving a positive signal increases. The signal in this model captures the utility the party leader receives from the passage of high-quality legislation. Hence, as the desire for high-quality increases, the more likely it is that the party leader suffers from second-mover disadvantage.

However, moving second does have potential benefits to the party leader, in that they are able to affect the time allocation of the populist candidate more than the legislative entrepreneur, so their action has a differential effect on re-election probabilities. As well as this, the party leader only has to reward the politician if they are re-elected. As there is a cost to providing legislative benefits, potentially forgoing this cost while still incentivising the politician to engage in legislative effort is beneficial.

Furthermore, when the politicians gain large amounts of utility from being re-elected, the party leader's legislative incentive may actually be more effective at changing the politician's action for some levels of the legislative benefit. The party leader therefore may prefer moving second in the case where c is sufficiently large.

Several extensions to the general framework considered here would be fruitful. A formal characterisation of the conditions under which the party leader prefers to move second would provide a deeper understanding of the set-up. As well, considering the extent to which there being two principals rather than one affects optimal effort allocation would further illuminate the extent to which the party leader's moving second distorts incentives.

3.7 Conclusion

When faced with two principals, the voter and the leadership of their party, politicians need to decide what kind of reputation they wish to build. Depending on their preference for legislative benefits, politicians may wish to be seen as being either a legislative entrepreneur or a populist. If they value the benefits on offer from their party leadership, politicians may risk not being re-elected in order to increase the probability of receiving those benefits.

Our approach also captures the incentives of the party leader. The party leader wants to select politicians who are interested in working on legislative tasks, while at the same time also wants to have politicians of their party re-elected. If the latter incentive is large, the party leader may prefer to be able to commit to not providing legislative rewards, so as to ensure that the politician maximises their probability of being re-elected. In general, however, they do not have the means of doing so.

Even if the party leader does not value re-election, we find that by offering the legislative incentive they may actually reduce the legislative effort of some politicians. If the benefit the politician receives from being re-elected is small enough, the legislative reward acts as an incentive to increase effort in the constituency in order to increase the probability of re-election, whereupon they can receive the benefits bestowed on them by the party leader.

Both of the above points suggest that the party leader may suffer from both a lack of commitment power and second mover disadvantage, both of which we examine.

We show that the party leader is always strictly worse off as a result of lacking commitment power. As a result of not being able to commit to a reward profile, the party leader is unable to take into account the fact that the legislative benefit affects the probability that the politician will be re-elected, affecting the probability of its payment. Furthermore, the level of legislative benefit affects the probability that different types will be re-elected differently: increasing the legislative reward decreases the probability that the populist type will be re-elected. The party leader's action therefore has a selection effect, which they do not take into account when they are unable to commit to an action profile.

Constructing a game in which the party leader can choose the level of legislative reward and is able to reward the politician even if they are not re-elected makes it clear that the party leader's second-mover disadvantage is increasing as the expected benefit the party leader receives from the passage of high-quality legislation.

Our analysis implies that second-mover disadvantage is more likely in periods of time (or countries) in which partisanship is less salient. In such circumstances, re-election is relatively less valuable, and good quality legislation is more highly valued. As moving second makes re-election more likely and a positive realisation of the party leader's signal more likely, it follows that moving second is more costly than in cases where party leadership is more focused on the re-election of members of their own party.

However, if the cost of rewarding the politician is sufficiently high, the party leader actually prefers the case where they move second. Moving second makes it more likely that the party leader incentivises the politician to engage in legislative effort without having to provide the legislative reward, which is costly. Furthermore, when the optimal level of legislative benefit is sufficiently low (for example, because it is costly to reward the politician), the party leader's action may actually be more effective at increasing legislative effort when they move second than when they do not.

Appendix

Proof of Theorem 3.1

The characterisation equilibrium follows from the analysis in the main text of the paper. To see that PBE is unique, consider each of the periods in turn:

Period 3: Given the party leader's preferences, the stated strategy is uniquely optimal for every possible history of the game, and the belief profile α is pinned down by Bayes rule. The stated strategy in the proposed equilibrium strategy in period 3 is uniquely optimal for each type of politician given any profile of past actions by the politician and the two principals.

Period 2: Given the voter's preferences, the stated strategy is uniquely optimal for every possible history of the game, and the belief profile μ is pinned down by Bayes' rule.

Period 1: Given the fact that the action profiles in each future period are unique for a given set of parameter values, it follows that the action profile described in Theorem 3.1 is unique.

As the strategy profile for each player at each period period is unique for a given set of parameters, it follows that the uniqueness statement in Theorem 3.1 holds.

Proof of Proposition 3.1

From Section 3, if the optimal effort allocation of the politician is interior, then it is expressed:

$$\bar{x}_{i,1} = \frac{a_i + q^*(zk_H + V_i)}{1 + 2z\tilde{k}r^*}.$$

This expression is equal to $a_i$when the following holds:

$$q^*(z_i k_H + V_i) > 2z\tilde{k}r^*,$$

which implies that the threshold value of $\bar{z}_i$ can be expressed as follows:

$$\bar{z} = \frac{q^* V_i}{2\tilde{k}r^* - q^* k_H}. \tag{3.5}$$

Hence, if $z < \bar{z}$, then the politician's first period effort is (weakly) greater than their effort upon re-election, with the inequality strict if the politician's effort allocation is interior.

Proof of Proposition 3.2

The proof follows immediately from equation (3.5) in the proof of Proposition 3.1.

Proof of Proposition 3.3

By the definition of α and k^*, it is clear that $k_L < k_H$. Furthermore if $d > 0$, then $0 < k_L < k_H$.

Call the probability of the politician being re-elected $p_i(\tilde{k}, k_H)$, where $p_i(.) : \mathbb{R} \rightarrow [0, 1]$ is a decreasing function because we have assumed that $x^*_{1,i} < \hat{x}_{1,i}$. The total expected benefit to the party leader receives directly re-election is $p_i(\tilde{k})v$. As $v \rightarrow \infty$, the benefit of increasing $p_i(\tilde{k})$ approaches infinity, which would imply that there exists a $\bar{v}$ that is sufficiently large such that $N(\tilde{k}, \bar{v}) = N(0, \bar{v})$ where $\tilde{k} > 0$.

Combined, the two paragraphs above imply that there trivially exists some threshold value of v, $\bar{v}$, above which the party leader's utility is strictly higher in the equilibrium where they can commit to not offering a legislative benefit.

Proof of Proposition 3.4

Recall that inequality (3.1) above is as follows:

$$\frac{a_i + (zk_H^* + V)}{1 + 2z\tilde{k}} \geq a_i + (V + k^*).$$

Re-arranging (3.1) then yields the following expression:

$$\frac{k_H^*}{\tilde{k}} \geq 2(a_i + e + \frac{1}{4}).$$

As the left hand side is greater than 1 in equilibrium, it follows that if $2a_i + \frac{1}{4} < 1$, then there exists an $\bar{e} \in \mathbb{R}_+$ such that if $e < \bar{e}$, then the above inequality holds.

Proof of Proposition 3.5

Let "type 1" denote type l and type 2 denote type p. Define $\boldsymbol{x}$ as the vector $(x_{1,p}(k_H, k_l), x_{1,l}(k_H, k_l))$. The party leader's utility function in this case can be stated:

$$N(\boldsymbol{x}(k_L, k_H; \Gamma_j)) = \sum_{i=1}^{2} 0.5(y_{i,1}\theta + x_{i,1}((1-a_i)\theta - y_{1,i}(\frac{1}{2}ck_H^2) - x_{1,i}(\frac{1}{2}ck_L^2)) + 0.5\theta(1 - x_{i,1})$$

with the last term resulting from the fact that $0.5(1-a_l)\theta + 0.5(1-a_p)\theta = 0.5\theta$ given that $a_l = 1 - a_p$ by assumption. As argued in the main text, if $V < \tilde{V}$,then $\frac{\partial N(\boldsymbol{x}(k_L,k_H)}{\partial k_L} < 0$. This follows from the following observations:

1. $\frac{\partial y_{i,1}}{\partial k_L} < 0$ for both i if $a_p + V < 1$;

2. $\frac{\partial x_{i,1}}{\partial k_L} > 0$ for both i if $a_p + V < 1$ and $c > 0$ so $y_{1,i}(\frac{1}{2}ck_H^2)$ and $(1 - y_{1,i})(\frac{1}{2}ck_L^2)$ are both less than zero;

3. $\frac{\partial x_{i,1}}{\partial k_L \partial a_i} > 0$ and $0.5 < a_p$. Hence, $\sum_{i=1}^{2} 0.5(x_{1,i}a_i + 0.5(1 - x_{1,i}))$ is decreasing in k_L if $a_p + V < 1$.

Observations (1)-(3) indicate that $N(\boldsymbol{x}(k_L, k_H))$ is falling in k_L, hence $k_L^{'} = 0$. Note that if $V > \tilde{V}$, then $\frac{\partial y_{1,i}}{\partial k_H} > 0$.

Define: $\sum_{i=1}^{2} 0.5x_{i,1}a_i + 0.5(1 - x_{i,1}) := Y$. By point 3 above, $\frac{\partial Y}{\partial k_H} > 0$. Hence, when $k_H = 0$, it must be that $\frac{\partial N(\boldsymbol{x}, k_L, k_H)}{\partial k_H} > 0$. This holds as:

$$\frac{\partial y_{i,1}}{\partial k_H}|_{k_H=0} + \frac{\partial Y}{\partial k_H}|_{k_H=0} > 0.$$

This result implies that in any solution to the party leader's maximisation problem, $k_H^{'} > 0$.

Proof of Proposition 3.6

As shown in the main text, in the no-commitment equilibrium, $k_H^* = 0$ when $d = 0$. However, as shown in the proof of Proposition 3.5, if $V > \tilde{V}$, then $k_H^{'} > 0$ and $\frac{\partial N(\boldsymbol{x}(k_L, k_H))}{\partial k_H}|_{k_H=0} > 0$. Hence, any equilibrium in the commitment case yields a higher utility than in the no-commitment case.

Proof of Theorem 3.2

Suppose that $d > 0$. Then the party leader's payoff can be expressed:

$$N(k_L, k_H; \Gamma_j) = \sum_{i=1}^{2} (0.5\theta(y_{i,1}(\Gamma_j) + x_{i,1}(\Gamma_j)((1 - a_i)\theta + y_{i,1}(\Gamma_j)(d(1 - a_i)k_H - \frac{1}{2}ck_H^2)$$
$$+ (1 - y_{i,1}(\Gamma_j))(d(1 - a_i)k_L - \frac{1}{2}ck_L^2)) + 0.5\theta y_{i,1}(\Gamma_j).$$

If $c > \bar{c}$ and $a_p + V < 1$, then it follows that maximisation problem $\max_{k_L, k_H} N(k_L, k_H)$ has a unique interior optimum, which must solve the first-order conditions:

$$\frac{\partial N(k_L^{'}, k_H^{'}; \Gamma_c)}{\partial k_j} = 0 \text{ for } j = L, H. \tag{3.6}$$

Note that:

$$\sum 0.5x_{i,1}(k_j^*)(y_{1,i}(k_j^*)(d(1-a_i)-ck_j^*)) = 0 \text{ for } j = L, H$$

by the definition of k_j^*. The term on the right hand side of the equation is present in both the first-order condition in (3.7). Hence for $k_L^{'} = k_L^*$ and $k_H^{'} = k_H^*$, it must be that the other terms in the first-order conditions in (3.6) are equal to zero for the vector (k_L^*, k_H^*) . Define:

$$\beta_{k_j}(k_L, k_H) := \frac{\partial N(k_L^{'}, k_H^{'}; \Gamma_c)}{\partial k_j} - \sum_{i=1} 0.5x_{i,1}(k_j)(y_{i,1}(k_j)(d - ck_j))$$

For (k_L^*, k_H^*) to be the solution to the party leader's first-order condition it must be the case that:

$$\beta_{k_j}(k_L^*, k_H^*) = 0 \text{ for } j = L, H \tag{3.7}$$

Given the above functional form of the party leader's utility function, $\beta_{k_j}(k_L, k_H)$ can be written:

$$\beta_{k_j}(k_L, k_H) = g(\frac{\partial x_{l,1}}{\partial k_j}, \frac{\partial x_{p,1}}{\partial k_j})$$

where $g(.,.)$ is a linear function that is non-constant in both of its inputs. When $x_{1,i}^*$ is interior for both i, it must be the case that:

$$\frac{\partial x_{i,1}}{\partial k_H} \neq \frac{\partial x_{i,1}}{\partial k_L}. \tag{3.8}$$

Given the result in (3.8), it must be the case that no interior equilibrium can satisfy the two conditions implied by (3.7) simultaneously. Hence for all $d > 0$, if $k_L^* = k_L^{'}$ then $k_H^* \neq k_H^{'}$ and vice versa. Noting that when $d = 0$, $k_L^* \neq k_L^{'}$ and $k_H^* \neq k_H^{'}$,this result implies that the party leader strictly prefers the commitment equilibrium to the no-commitment equilibrium for all $d \in [0, \infty)$.

Proof of Theorem 3.3

Again, let "type 1" denote type l and type 2 denote type p, as in Proposition 3.5. The party leader's payoff function when Γ_c is played can then be stated:

$$N(0, k_H; \Gamma_c) = \sum_{i=1}^{2} 0.5(y_{i,1}(\Gamma_c)\theta + x_{i,1}(\Gamma_c)((1-a_i)\theta - y_{i,1}(\Gamma_c)(\frac{1}{2}ck_H^2)) + 0.5\theta y_{i,1}(\Gamma_c).$$

The equivalent payoff function for Γ_2 is as follows:

$$N(0, k_H; \Gamma_2) = \sum_{i=1}^{2} 0.5(y_{i,1}(\Gamma_2)(\theta - (\frac{1}{2}ck_H^2)) + x_{i,1}(\Gamma_2)(1-a_i)\theta + 0.5\theta y_{i,1}(\Gamma_2).$$

Define:

$$\tilde{N}(\boldsymbol{k}; \theta, c, V, z) = \sum_{i=1}^{2} 0.5(\tilde{y}_{i,1}[\theta - (1 - x_{i,1}(\Gamma_c)](\frac{1}{2}ck_H^2) + \tilde{x}_{i,1}(1-a_i)\theta) + 0.5\theta\tilde{y}_{i,1}.$$

where $\tilde{y}_{i,1} = y_{i,1}(\Gamma_2) - y_{i,1}(\Gamma_c)$ and $\tilde{x}_{i,1} = x_{i,1}(\Gamma_2) - x_{i,1}(\Gamma_c)$. $\tilde{N}(\boldsymbol{k}; \theta, c, V, z)$ then denotes the different between the party leader's utility in Γ_2 and game Γ_c. Note that if:

$$k_H^*(\Gamma_2) \geq \frac{a_p + V}{z}$$

then the probability that the party leader receives a positive signal is equal to 1 for both types. Given the expression for $N(0, k_H; \Gamma_2)$, it is clear that for any c there exists a $\tilde{\theta} > 0$ such that if $\theta > \tilde{\theta}$ the above inequality holds.

Note that, given the expression for $x_{i,1}^*(\Gamma_c)$, there is no finite value of k_H for which $x_{i,1}^*(\Gamma_c) = 0$. Hence, $\tilde{N}(\boldsymbol{k}; \theta, c, V, z)$ is increasing in θ function, and $\tilde{N}(\boldsymbol{k}; \theta, c, V, z) \to \infty$ as $\theta \to \infty$. It follows that there exists a $\bar{\theta} \in \mathbb{R}_+$ such that if $\theta > \bar{\theta}$, then $\tilde{N}(\theta) > 0$, and hence the result is proved.

Conclusion

In Chapters 1 and 2, we consider the incentives of a profit-maximising platform owner who is able to affect which sellers groups of consumers observe. We consider several ways in which the platform owner could be constrained in matching buyers to sellers, including the owner only being able to choose: the probability that a seller is prominent to buyers; the probability that an entrant is observed, keeping the rest of the network fixed; and to what extent product ratings determine the probability that a seller can be observed.

A consistent observation throughout both chapters is that the platform owner faces a trade-off when increasing the probability that buyers and sellers are matched: a match generates sales and therefore profit, but at the same time, the more sellers a buyer observes, the more competition there is across the entire network. The nature of price competition on networks, and specifically the fact that indirect connections affect prices, means that concentrating competition by making a subset of sellers prominent decreases prices.

The platform owner resolves their trade-off by reducing the extent to which sellers are prominent. Doing so results in fewer high-quality matches but generates higher prices, and therefore yields higher aggregate profits.

The platform owner's incentive to reduce seller prominence and seller observability more generally is harmful to consumer welfare. Consumer surplus is highest when buyers observe as many products as possible and prices are low. The complete network, where buyers observe sellers with probability 1, therefore always maximises consumer surplus.

Platform owners in general thus have an incentive to design platforms in a manner that is harmful to consumer welfare. They tend to prefer networks where competition is less concentrated, which in turn implies that buyers observe only a subset of the total products on the market, including those that are high-quality.

In the last decade, competition authorities have taken an increasing interest in ensuring that large online markets function such that there is enough competition between platforms. However, as platforms across the digital commerce sector consolidate, the extent to which platforms that link buyers and sellers can control the market environment in which competition takes place becomes increasingly important.

Our analysis indicates that policymakers should ensure that platform owners are not distorting intra-platform competition to increase prices and harm consumer welfare. To do so, competition authorities would need to examine the search processes that large online retail platforms use to match buyers with sellers. Specific recommendations on how to improve competition would be platform- and market-specific, as our analysis in Chapter 2 suggests.

In Chapter 3, we examine the optimal effort allocation of a politician who is incentivised to engage in constituency effort by voters and legislative effort by their party's leader in Congress. Doing so explains why some politicians choose to frontload their legislative effort, while others backload it: if a politician values legislative rewards highly, then they have an incentive to build a reputation as someone who cares about legislative work at the expense of their constituency.

Our analysis also sheds light on the party leader's incentive to provide the politician with legislative benefits. As the party leader is only able to reward the politician when the politician is re-elected, the effectiveness of the party leader's action is potentially constrained, and they may in some cases prefer to actually increase constituency effort, as in doing so the politician is more likely to be re-elected and thus be able to potentially claim the legislative benefit.

If the party leader cannot commit to a schedule of legislative rewards prior to the politician's performing their action, they are unable to take into account the fact that their reward affects the likelihood that the politician will be re-elected. This implies that the party leader is always strictly worse off in the case where they lack commitment power.

Furthermore, the party leader loses out because they move second. Since their action becomes less effective as the probability that the politician is re-elected becomes low, if the cost of rewarding the politician is low compared to the benefit associated with high legislative effort, the hypothetical case in which they would be able to reward the politician whether they are re-elected or not would be better for the party leader than the case in which they move second. If, however the cost of providing legislative rewards is high enough, then the party leader may actually prefer that they move second.

Bibliography

Amir, R., P. Ericksonz, and J. Jin (2015, July). On the microeconomic foundations of linear demand for diferentiated products. Discussion Paper Series, School of Economics and Finance 201507, School of Economics and Finance, University of St Andrews.

Armstrong, M. and J. Zhou (2011). Paying for prominence. *The Economic Journal 121*(556), F368–F395.

Ashworth, S. (2005, 08). Reputational Dynamics and Political Careers. *The Journal of Law, Economics, and Organization 21*(2), 441–466.

Austen-Smith, D. and J. Banks (1990). Stable governments and the allocation of policy portfolios. *The American Political Science Review 84*(3), 891–906.

Ballester, C., A. Calvò-Armengol, and Y. Zenou (2006). Who's who in networks. wanted: The key player. *Econometrica 74*(5), 1403–1417.

Banks, J. and R. K. Sundaram (1998). Optimal retention in agency problems. *Journal of Economic Theory 82*(2), 293–323.

Barro, R. J. (1973). The control of politicians: An economic model. *Public Choice 14*, 19–42.

Bergemann, D., B. Brooks, and S. Morris (2015, March). The limits of price discrimination. *American Economic Review 105*(3), 921–57.

Bernheim, B. D. and M. D. Whinston (1986). Common agency. *Econometrica 54*(4), 923–942.

Besley, T. and A. Case (1995, March). Incumbent Behavior: Vote-Seeking, Tax-Setting, and Yardstick Competition. *American Economic Review 85*(1), 25–45.

Bimpikis, K., S. Ehsani, and R. Ilkiliç (2018). Cournot competition in networked markets. *Management Science (forthcoming).*

Bonacich, P. (1972). Factoring and weighting approaches to status scores and clique identification. *The Journal of Mathematical Sociology 2*(1), 113–120.

Bonacich, P. (1987). Power and centrality: A family of measures. *American Journal of Sociology 92*(5), 1170–1182.

Bramoullé, Y., R. Kranton, and M. D'Amours (2014, March). Strategic interaction and networks. *American Economic Review 104*(3), 898–930.

Braverman, A. and J. Stiglitz (1982). Sharecropping and the interlinking of agrarian markets. *American Economic Review 72*(4), 695–715.

Bruce Cain, J. F. and M. Fiorina (1987). *The Personal Vote: Constituency Service and Electoral Independence.* Cambridge, MA: Harvard University Press.

Candogan, O., K. Bimpikis, and A. Ozdaglar (2012). Optimal pricing in networks with externalities. *Operations Research 60*(4), 883–905.

Charlson, G. (2014). A simple single-agent-multi-principal model of political reputation. Master's thesis, University of Oxford.

Chawla, S., F. Niu, and T. Roughgarden (2009, 07). Bertrand competition in networks. *SIGecom Exchanges 8*.

Chen, Y. and J. Zhang (2020, 05). Signalling by bayesian persuasion and pricing strategy. *The Economic Journal 130*, 976–1007.

Cominetti, R., J. R. Correa, and N. E. Stier-Moses (2009, December). The Impact of Oligopolistic Competition in Networks. *Operations Research 57*(6), 1421–1437.

Corominas-Bosch, M. (2004). Bargaining in a network of buyers and sellers. *Journal of Economic Theory 115*(1), 35–77.

Cox, G. and J. Katz (2002). *Elbridge Gerry's Salamander: The Electoral Consequences of the Reapportionment Revolution*. Cambridge University Press.

Cox, G. and M. Mccubbins (2005, 01). Setting the agenda: Responsible party government in the us house of representatives. *Setting the Agenda: Responsible Party Government in the U.S. House of Representatives*, 1–336.

De Francesco, M. and N. Salvadori (2013). Bertrand-edgeworth competition in an almost symmetric oligopoly. *Studies in Microeconomics 1*(2), 213–219.

Dinerstein, M., L. Einav, J. Levin, and N. Sundaresan (2018, July). Consumer price search and platform design in internet commerce. *American Economic Review 108*(7), 1820–59.

Dixit, A., G. Grossman, and E. Helpman (1997). Common agency and coordination: General theory and application to government policy making. *Journal of Political Economy 105*(4), 752–69.

Elliott, M. and A. Galeotti (2019, 12). The role of networks in antitrust investigations. *Oxford Review of Economic Policy 35*(4), 614–637.

Fenno, R. (1978). *Home Style: House Members in Their Districts*. Little, Brown & Co.

Gailmard, S. (2009). Multiple principals and oversight of bureaucratic policy-making. *Journal of Theoretical Politics 21*(2), 161–186.

Galeotti, Andrea, G.-B. and S. Goyal (2019). Targeting interventions in networks.

Gowrisankaran, G., M. F. Mitchell, and A. Moro (2008, January). Electoral Design and Voter Welfare from the U.S. Senate: Evidence from a Dynamic Selection Model. *Review of Economic Dynamics 11*(1), 1–17.

Hatfield, J. and G. P. i Miquel (2006). Multitasking, limited liability and political agency. Technical report, LSE.

Holmstrom, B. (1999). Managerial incentive problems: A dynamic perspective. *The Review of Economic Studies 66*(1), 169–182.

Holmstrom, B. and P. Milgrom (1987). Aggregation and linearity in the provision of intertemporal incentives. *Econometrica 55*(2), 303–328.

Huber, G., G. Sanford, and D. Landa (2007). Challenger entry and voter learning. *American Political Science Review 101*(2), 303â320.

Ireland, N. J. (1993). The provision of information in a bertrand oligopoly. *The Journal of Industrial Economics 41*(1), 61–76.

Janssen, M. and E. Rasmusen (2002). Bertrand competition under uncertainty. *The Journal of Industrial Economics 50*(1), 11–21.

Jeffreys, H. and B. S. Jeffreys (1988). *Methods of Mathematical Physics, 3rd ed.*, Chapter "Uniform Convergence of Sequences and Series", pp. 37–43. Cambridge University Press.

Johannes, J. R. (1980). The distribution of casework in the u. s. congress: An uneven burden. *Legislative Studies Quarterly 5*(4), 517–544.

Johnson, J. P. and D. P. Myatt (2006, June). On the simple economics of advertising, marketing, and product design. *American Economic Review 96*(3), 756–784.

Kim, J. B., P. Albuquerque, and B. J. Bronnenberg (2010). Online demand under limited consumer search. *Marketing Science 29*(6), 1001–1023.

Kranton, R. E. and D. F. Minehart (2001, June). A theory of buyer-seller networks. *American Economic Review 91*(3), 485–508.

Levitt, S. D. and J. M. Snyder (1997). The impact of federal spending on house election outcomes. *Journal of Political Economy 105*(1), 30–53.

Lewis, T. and D. Sappington (1994). Supplying information to facilitate price discrimination. *International Economic Review 35*(2), 309–27.

Lohmann, S. (1998). An information rationale for the power of special interests. *The American Political Science Review 92*(4), 809–827.

Luca, M. (2016). Reviews, reputation, and revenue: The case of yelp.com. Harvard Business School Working Papers 12-016, Harvard Business School.

Marjorie Sarbaugh-Thompson, Lyke Thompson, C. E. J. S. and R. Elling (2004). *The Political and Institutional Effects of Term Limits*. Palgrave Macmillan.

Mastrolia, T. and Z. Ren (2018). Principal-agent problem with common agency without communication. *SIAM Journal on Financial Mathematics*.

McAfee, R. (1994, 02). Endogenous availability, cartels, and merger in an equilibrium price dispersion. *Journal of Economic Theory 62*, 24–47.

Melo, E. (2017). Imperfect price competition in networks. *Job market paper*.

Milgrom, P. and I. Segal (2002). Envelope theorems for arbitrary choice sets. *Econometrica 70*(2), 583–601.

Ogut, H. and B. K. O. Tas (2012). The influence of internet customer reviews on the online sales and prices in hotel industry. *The Service Industries Journal 32*(2), 197–214.

Pearson, K. (2005). *Party Discipline in the Contemporary Congress: Rewarding Loyalty in Theory and in Practice.* Ph. D. thesis, University of California, Berkeley.

Persson, T. and G. Tabellini (2000). *Political Economics: Explaining Economic Policy.* MIT Press.

Polanski, A. (2007). Bilateral bargaining in networks. *Journal of Economic Theory 134*(1), 557–565.

Ringel, D. M. and B. Skiera (2016). Visualizing asymmetric competition among more than 1,000 products using big search data. *Marketing Science 35*(3), 511–534.

Snyder, J. M. and S. Ansolabehere (2002, 2002). "the incumbency advantage in u.s. elections: An analysis of state and federal offices, 1942-2000". *Election Law Journal 1*(3), 315–338.

Spiegler, R. (2006). The market for quacks. *The Review of Economic Studies 73*(4), 1113–1131.

Wawro, G. (2000). *Legislative Entrepreneurship in the U.S. House of Representatives.* University of Michigan Press.

Zaller, J. (1998). *Politicians as Prize Fighters: Electoral Selection and Incumbency Advantage. In Politicians and Party Politics.* John Hopkins University Press.