

Product Differentiation and Oligopoly: a Network Approach

Bruno Pellegrino

UNIVERSITY OF MARYLAND

January 21, 2022

Abstract

I present a theory of oligopoly and markups in general equilibrium, and use it to study the welfare impact of monopoly power. I develop a model populated by granular firms that compete in a network game of oligopoly, which co-exist with a fringe of atomistic competitors that enter and exit endogenously. To model product substitution, I introduce a Generalized Hedonic-Linear (GHL) demand system, which I identify using a publicly-available dataset that measures product similarity among all US public corporations. Using my model, I decompose firm-level markups into two components: one related to productivity and one related to product market centrality (which measures competition): both have contributed to the secular rise of markups. I estimate, for the year 2019, a large deadweight loss from oligopolistic behavior (about 11.5% of total surplus). Even larger is the effect on consumer surplus, which would double if firms behaved atomistically. My analysis also suggests that both these numbers are significantly larger than they were in the mid-90s, implying that oligopoly power exerts an increasingly large effect on welfare. To conclude, I present counterfactual analysis which suggests a potentially important role of startup acquisitions in generating these trends.

JEL Codes: D2, D4, D6, E2, L1, O4

Keywords: Competition, Concentration, General Equilibrium, Market Power, Markups, Mergers, Monopoly, Networks, Oligopoly, Startups, Text Analysis.

This paper received the following awards: EEA Best Job Market Paper Award, WFA Best Corporate Finance Paper, UCLA Xavier Drèze Prize for Outstanding Research Paper and Best Presentation Award at the Unicredit Job Market Bootcamp. I gratefully acknowledge financial support from the Price Center for Entrepreneurship at UCLA Anderson. I am indebted to my dissertation advisors Hugo Hopenhayn, Nico Voigtländer, Andy Atkeson, Ed Leamer, Stavros Panageas and Romain Wacziarg for their invaluable advice and guidance. I benefited from thoughtful conference discussions by Michael Gofman, Naveen Khanna, Isabelle Mejean and Jane Olmstead-Rumsey. I also received valuable comments from John Asker, David Baqaee, Alberto Bisin, Ariel Burstein, Michela Giorcelli, Thomas Groesbeck, Kurt Mitman, Oleg Itskhoki, Nicola Persico, Sebastian Ottinger, Tommaso Valletti, Maria Lucia Yanguas, workshop participants at the the NBER Mega-Firms Conference, the WFA, the NBER Summer Institute (Income Distribution and Macroeconomics), the EEA and many other conferences and seminars. Finally, I acknowledge outstanding data collection assistance by Michele Schincaglia.

1. Introduction

Industry concentration, markups, and corporate profit rates have all increased in the United States during the past two decades (Grullon et al., 2019; De Loecker et al., 2020). This fact has spurred important public debates over whether these trends reflect an oligopolization of U.S. industries and whether a revised antitrust paradigm is necessary (Khan, 2018). While standard price theory arguments suggest that the welfare implications of these trends might be significant, interpreting these trends presents an imposing methodological challenge. The study of market power has traditionally resided within the literature on Empirical Industrial Organization (EIO). Yet, there is a consensus that these trends are macroeconomic in nature: standard EIO methodologies are unfeasible, as they require data that is not available for more than a handful of industries (Syverson, 2019).

This paper investigates the welfare consequences of the increased concentration of U.S. industries. I address the existing methodological challenges by introducing a novel general equilibrium model with two types of firms: a finite set of *granular* firms that behave as oligopolists and a continuum of atomistic producers that behave competitively and can enter and exit. To model product market competition among the oligopolists, I propose a Generalized Hedonic-Linear (GHL) demand system, which I estimate using a dataset recently developed by Hoberg and Phillips (2016). This dataset provides measures of product similarity for all pairs of publicly-traded corporations in the US. The empirical implementation of the model allows me address the following question: how have consumer surplus and the welfare costs of oligopoly evolved as a consequence of industry consolidation during this period?

Using my novel theoretical framework, I document that the increased concentration of US industries over the past twenty years was accompanied by an increase in oligopoly power, measured as: (1) an increase in the deadweight losses induced by oligopolistic behavior; (2) a decline in the share of total surplus that accrues to consumers. My methodology also allows to me associate these trends with another well-known stylized fact: the dramatic rise in takeovers of startups that began in the mid 1990s, and which coincided with the well-known secular decline in Initial Public Offerings (IPOs) (Kahle and Stulz, 2017).

Economists have long been concerned with market power. Since the 1980s, the EIO literature has been developing a conceptual “toolkit” that researchers and antitrust enforcement practitioners have used to analyze market power within industries (Einav and Levin, 2010). The EIO approach requires the researcher to first understand the structure of product market rivalries in an industry: a firm’s ability to price above marginal cost depends critically on the intensity of competition from firms that produce similar products. As a consequence, this literature has shown that oligopoly power is inextricably linked to the notion of product differentiation: to measure a firm’s market power in an industry with n firms, the economist effectively needs to first estimate n^2 cross-price demand elasticities—one for each pair of rivals. This is because variations in the supply of any product cause the residual demand curve of every competing product to shift, and the shift is larger for products that are closer substitutes.

In industry studies, demand estimation is nowadays largely based on hedonic models (Berry, Levinsohn and Pakes, 1995). The current resurgence in market power and antitrust research, however, has a distinctive macroeconomic angle. Because we do not observe output volume, prices, or product characteristics for a sufficiently large cross-section of industries, this EIO approach cannot be directly applied in a macroeconomic context. This challenge is compounded by the problem that, at the macro level, product-market rivalry is not well approximated by industry classifications. Industry classifications (such as NAICS) tend to be based on similarities in the production process, not on the degree of product substitutability. In other words, they are appropriate for estimating production functions, but they are unreliable when it comes to estimating cross-price demand elasticities. In addition, the very concepts of industry or sector are more fluid

than macroeconomists tend to assume. While industry classifications are static, larger companies (those more likely to have market power) move frequently from one industry to another, and have been shown to strategically manipulate their industry classification—a phenomenon that has been dubbed *industry window dressing* (Chen et al., 2016).

Despite these challenges, the macroeconomics literature has made significant progress in incorporating market power in general equilibrium models: Baqaee and Farhi (2020, henceforth BF) have recently shown how to approximate the welfare costs of imperfect competition, under minimal assumptions, using the cross-sectional distribution of markups. Their approach is (by design) very general, but it is also agnostic about the origin of markups, which are assumed to be exogenous and observable. Because this approach is silent about how the observed cross-section of markups arises in the first place, a separate theory of markup formation is needed in order to simulate the welfare consequences of a change in market structure.

This study breaks new ground by providing a theory of firm size and profitability that generalizes the Cournot oligopoly model to differentiated products and hedonic demand, and embeds it in a general equilibrium model. The objective of my model is not to capture all sources of variation in markups, but rather to isolate the variation that can be reliably attributed to product market rivalry. Through this approach, I can quantify the contribution of each individual producer to aggregate welfare, and I can study the general equilibrium effects of events that are relevant to antitrust policy, such as mergers or the entry of additional firms.

To achieve this, my theoretical model dispenses with the notions of industry and sector altogether, building instead on the tradition of hedonic demand (Lancaster, 1966; Rosen, 1974). Thus, I can link the cross-price elasticity of demand between all firms in the economy to the fundamental attributes of each firm’s product portfolio. Each firm’s output is modeled as a bundle of characteristics that are individually valued by the representative consumer. The cross-price elasticity of demand between two firms depends on the characteristics embedded in their output. If the product portfolios of two companies contain similar characteristics, the cross-price elasticity of demand between their products is high. The result is a rather different picture of the product market: not a collection of sectors, but a network, in which the distance between nodes reflects product similarity and strategic interaction between firms.

The key assumptions of my model are: (1) the representative consumer’s preferences are described by a utility function that is quadratic-in-characteristics ; (2) firms compete à la Cournot¹; (3) the marginal cost function is linear in output. Based on these assumptions, the firms in my model play a linear-quadratic game over a weighted network, a type of potential game that has been extensively studied in the micro theory literature (see Ballester, Calvó-Armengol and Zenou, 2006; Ushchev and Zenou, 2018).

This is the first paper to show how to derive the network Cournot model starting from a hedonic utility specification, to embed the game in a general equilibrium framework and to take the model to the data in a structural way.

To estimate my model, I use a recently-developed data set (Hoberg and Phillips, 2016, henceforth HP) that provides continuous measures of product similarity for every pair of publicly traded firms in the United States. These product-similarity scores are based on a computational-linguistics analysis of product descriptions obtained from SEC filings, and are mapped by my model into an $n \times n$ matrix of cross-price demand elasticities. Moreover, because HP’s similarity scores are time-varying (yearly observations since 1997), my model is unique in that the degree of product substitution between individual firms is allowed to change over time.

Crucially, the empirical implementation of my model does not require any proprietary or confidential

¹I also study the Bertrand case in the Appendix.

data, and is computationally tractable. Two datasets are required: Compustat and HP’s cosine similarity data, which the authors have made publicly-accessible through an online repository.²

I use my model to compute the deadweight loss from oligopoly and to simulate changes in total surplus and consumer surplus for a number of counterfactuals. I find that the welfare costs of oligopoly are sizable. By moving to an allocation in which firms price at marginal cost (that is, in which they behave as if they were atomistic players in a perfectly-competitive market), total surplus would rise by approximately 11 percentage points; consumer surplus would double, partly due to total surplus being reallocated from producers to consumers. By computing a separate counterfactual that keeps the aggregate labor supply fixed (markups are equalized, rather than eliminated), I can determine that a significant share of the welfare loss from oligopoly—about 7.7 percentage points of the aforementioned 11—occurs by way of factor misallocation. In other words, the deadweight loss is driven not only by an underutilization of inputs, but also by a suboptimal mix of goods being produced. I also simulate a counterfactual in which all firms in the economy are owned by a single producer that implements a collusive equilibrium. Under this scenario, total surplus would drop by about one-tenth: with some degree of abstraction, we can think of this estimate as an upper bound to the welfare benefits of antitrust. Also, in this monopolistic/collusive equilibrium consumer surplus would decrease by about 38%, due partly to surplus being reallocated from consumers to producers.

By mapping my model to firm-level data for a period of 21 consecutive years, I investigate the welfare consequences of the rise in concentration and markups between 1997 and 2017. I find that the share of surplus appropriated by companies in the form of oligopoly profits has increased from about 50% (in 1997) to 55.5% (in 2017). When I subtract fixed costs (such as capital and overhead) from profits and total surplus, this increase becomes significantly steeper: from 11% in 1997 to 22% in 2017. This result is robust to different measurements of fixed costs and intangible capital, and suggests that the increase in the profit share of surplus is not justified by larger fixed costs.

I also use the model to decompose the increase in markups among US public firms into two components: one driven by (quality-adjusted) productivity and one driven by the firms’ centrality in the product network. Both these components contribute to the rise of markups, albeit through different margins: the productivity component increases by way of reallocation of market shares towards firms that differentiate vertically, while the increase of the increase in the centrality component, which reflects the availability of substitutable products, drives the within-firm growth of markups.

The welfare costs of oligopoly have also increased over this period. In terms of total surplus, the gap between the oligopolistic equilibrium and perfect competition (the deadweight loss) has increased from 8.5% (in 1997) to 11% (in 2017). Consumer surplus is adversely affected via two channels: less surplus is produced overall (as a percentage of the surplus that *could* be produced), and less of the diminished surplus is allocated to the consumer in equilibrium. Thus, an important contribution of this paper is to investigate the distributional implications of oligopoly. Overall, my empirical findings suggest that rising concentration and markups did indeed translate into a measurable increase in the welfare loss from oligopoly, and affected how surplus is shared between producers and consumers.

Finally, I use the counterfactual-building capabilities of the model to better understand the drivers of these trends. In particular, I study the impact of a dramatic secular shift that has occurred among venture capital (VC) startups over the past 20 years. In the early 1990s, most VC-backed startups (80%–90%), if successful, would exit³ through an Initial Public Offering (IPO). Today, the near entirety (about 94%) of

²See hobergphillips.tuck.dartmouth.edu

³In the entrepreneurial finance literature, an “exit” is the termination of a VC investment and should not be confused with a business termination. If the VC investor exits with an IPO, that event marks the entry of that firm in the universe of public firms, not an enterprise death.

successful VC-backed startups exit by getting acquired by an incumbent. I find that this shift accounts not only for the secular decline in the number of public corporations in the US (from about 7,500 in 1997 to about 3,500 in 2017) but also for a large share of the increase in the welfare costs of oligopoly, and of the rise in the profit share of surplus.

This paper aims to connect the new EIO literature (Einav and Levin, 2010) to two recent and growing branches of macroeconomics that use micro-data.

The first is the literature on networks (Acemoglu, Carvalho, Ozdaglar and Tahbaz-Salehi, 2012; Carvalho, 2014; Acemoglu, Ozdaglar and Tahbaz-Salehi, 2017; Carvalho and Tahbaz-Salehi, 2019; Baqaee and Farhi, 2020; Carvalho, Nirei, Saito and Tahbaz-Salehi, 2020). I contribute to and expand this literature, which has mostly focused on input-output networks, by considering a different type of network: that of product market rivalries.⁴

The second is the literature on markups and industry concentration (De Loecker, Eeckhout and Unger, 2020; Autor, Dorn, Katz, Patterson and Van Reenen, 2020; Edmond, Midrigan and Xu, 2018; Covarrubias, Gutiérrez and Philippon, 2020; Syverson, 2019). This paper builds on and adds to this body of work by incorporating hedonic demand as well as new data. These features allow me to go beyond markups and concentration, and to create a rich, high-dimensional representation of the competitive environment. In my model, firms differ not only by their productivity, but also by their products' characteristics; as a consequence, each firm has a distinct set of competitors that changes over time, as firms update their product's description in their SEC filings.

This paper also connects the recent literature on market power to the secular decline of public companies (Kahle and Stulz, 2017) and IPOs (Bowen, Frésard and Hoberg, 2018; Gao, Ritter and Zhu, 2013). My model allows to quantify the effects of these phenomena on the intensity of product market competition.

The rest of the paper is organized as follows. In Section 2, I present my theoretical model. In Section 3, I present the data used in the empirical part of the paper and show how it is mapped to the model. In Section 4, I present my empirical results. In Section 5, I discuss a number of extensions and robustness checks. In Section 6, I present my conclusions and discuss how my findings can inform the current debate on market power and antitrust policy.

⁴Outside the macro-networks literature, Bloom, Schankerman and Van Reenen (2013) have studied rivalry networks in a seminal empirical study of R&D spillovers.

2. A Theory of Imperfect, Networked Competition

In this section, I present a general equilibrium model in which firms produce differentiated products and compete à la Cournot. For expositional purposes, I start by laying out the basic model that only includes granular oligopolistic firms. After characterizing the equilibrium of this model economy and outlining a series of counterfactuals of interest, I extend the model (in Subsection 2.6) by adding a continuum of perfectly-competitive atomistic firms.

2.1. Basic Setup: the Generalized Hedonic-Linear (GHL) Demand System

There are n firms, indexed by $i \in \{1, 2, \dots, n\}$ that produce differentiated products. Following the tradition of hedonic demand in differentiated product markets (Lancaster, 1966; Rosen, 1974), I assume that consumers value each product as a bundle of characteristics. The number of characteristics is $(n + J)$.

There are two types of characteristics. The first J characteristics are common across all goods and are indexed by $j \in \{1, 2, \dots, J\}$, while the remaining n characteristics are idiosyncratic (that is, they are product-specific and cannot be imitated by other products) and therefore have the same index i as the corresponding product. The scalar a_{ji} is the number of units of common characteristic j provided by product i . Each product is described by a k -dimensional column vector \mathbf{a}_i , which I assume (without loss of generality) to be of unit length – formally:

$$\mathbf{a}_i = \begin{bmatrix} a_{1i} & a_{2i} & \dots & a_{Ji} \end{bmatrix}' \quad (2.1)$$

$$\text{such that } \sum_{j=1}^J a_{ji}^2 = 1 \quad \forall i \in \{1, 2, \dots, n\} \quad (2.2)$$

The vector \mathbf{a}_i therefore provides firm i 's coordinates in the space of common characteristics. We can stack all the coordinate vectors \mathbf{a}_i inside a $n \times J$ matrix that we call \mathbf{A} :

$$\mathbf{A} = \begin{bmatrix} \mathbf{a}_1 & \mathbf{a}_2 & \dots & \mathbf{a}_n \end{bmatrix} = \begin{bmatrix} a_{11} & a_{12} & \dots & a_{1n} \\ a_{21} & a_{22} & \dots & a_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ a_{k1} & a_{k2} & \dots & a_{kn} \end{bmatrix} \quad (2.3)$$

Let q_i be the number of units produced by firm i and consumed by the representative agent, which we write inside the n -dimensional vector \mathbf{q} :

$$\mathbf{q} = \begin{bmatrix} q_1 & q_2 & \dots & q_n \end{bmatrix}' \quad (2.4)$$

Definition 1. A vector \mathbf{q} that specifies, for every firm, the number of units produced is called an *allocation*.

I assume that there exists a representative agent. Consistent with the hedonic demand literature, the consumer combines linearly the characteristics of different products, and their preferences are defined in terms of these characteristics. Letting x_j be the total units of common characteristic j , we have:

$$x_j = \sum_{i=1}^n a_{ji} q_i \quad (2.5)$$

Hence, intuitively, the matrix \mathbf{A} transforms units of goods into units of common characteristics:

$$\mathbf{x} = \mathbf{A}\mathbf{q} \quad (2.6)$$

With regard to the n idiosyncratic characteristics, I assume that each unit of good i provides exactly one unit of its corresponding idiosyncratic characteristic. Hence, we can just write q_i in place of the units of idiosyncratic characteristic i .

The representative agent's preferences are described by a utility function that is quadratic in both common characteristics (\mathbf{x}) and idiosyncratic characteristics (\mathbf{q}). The agent's preferences also incorporate a linear disutility for the total number of hours of work supplied (H):

$$U(\mathbf{x}, \mathbf{q}, H) \stackrel{\text{def}}{=} \alpha \cdot \sum_{j=1}^k \left(b_j^x x_j - \frac{1}{2} x_j^2 \right) + (1 - \alpha) \sum_{i=1}^n \left(b_i^q q_i - \frac{1}{2} q_i^2 \right) - H \quad (2.7)$$

where b_j^x and b_i^q are characteristic-specific preference shifters. In linear algebra notation:

$$U(\mathbf{x}, \mathbf{q}, H) \stackrel{\text{def}}{=} \alpha \left(\mathbf{x}' \mathbf{b}^x - \frac{1}{2} \cdot \mathbf{x}' \mathbf{x} \right) + (1 - \alpha) \left(\mathbf{q}' \mathbf{b}^q - \frac{1}{2} \cdot \mathbf{q}' \mathbf{q} \right) - H \quad (2.8)$$

$\alpha \in [0, 1]$ is the utility weight that is assigned to common characteristics. Hence, it governs the degree of *horizontal differentiation* among products. This utility specification is a generalization of the preferences used by Epple (1987). In addition to introducing idiosyncratic characteristics, I make leisure the outside good: that allows me to close the model and make it general equilibrium.

I denote by h_i the labor input acquired by every firm, so that the labor market clearing condition is:

$$H = \sum_i h_i \quad (2.9)$$

I assume (without loss of generality) that labor is the numéraire of this economy (the price of one unit of labor is 1\$), therefore h_i is also the total variable cost incurred by firm i . Firm i produces output q_i using a quasi-Cobb Douglas production function:

$$q_i = k_i^\theta \cdot \ell(h_i) \quad (2.10)$$

where k_i is the capital input (fixed) and the function $\ell(\cdot)$ is such that firm i 's technology can be described by the following quadratic total variable cost function:

$$h_i = f_i + c_i q_i + \frac{\delta_i}{2} q_i^2 \quad (2.11)$$

Due to equation (2.10), we can see that c_i and δ_i must depend on k_i in the following way:⁵

$$c_i = \frac{\tilde{c}_i}{k_i^\theta}; \quad \delta_i = \frac{\tilde{\delta}_i}{k_i^{2\theta}}; \quad (2.12)$$

where \tilde{c}_i and $\tilde{\delta}_i$ are positive constants. MC and AVC denote, respectively, the marginal cost and the average variable cost:

$$\text{MC}_i = c_i + \delta_i q_i; \quad \text{AVC}_i = c_i + \frac{\delta_i}{2} q_i \quad (2.13)$$

⁵This in turn implies that $\ell(h_i) = \sqrt{\frac{h_i - f_i}{\delta_i} + \frac{\tilde{c}_i^2}{\delta_i^2}} - \frac{\tilde{c}_i}{\delta_i}$.

The representative consumer buys the goods bundle \mathbf{q} taking \mathbf{p} (the vector of prices) as given. Moreover, I assume that the representative consumer is endowed with the shares of all the companies in the economy. As a consequence, the aggregate profits are paid back to them. Their consumption basket, defined in terms of the unit purchased \mathbf{q} , has to respect the following budget constraint:

$$H + \Pi = \sum_{i=1}^k p_i q_i \quad (2.14)$$

Notice that for now we have defined aggregate economic profits Π to include all non-labor compensation (which equates to assuming that f_i is sunk). We will later consider a narrower metric of profits from which fixed costs ($F \stackrel{\text{def}}{=} \sum_i f_i$) are netted out.

2.2. Equilibrium

To streamline notation, let us define:

$$b_i \stackrel{\text{def}}{=} \alpha \sum_j b_{ji}^x x_j + (1 - \alpha) b_i^q \quad (2.15)$$

or, in linear algebra notation:

$$\mathbf{b} \stackrel{\text{def}}{=} \alpha \mathbf{A}' \mathbf{b}^x + (1 - \alpha) \mathbf{b}^q \quad (2.16)$$

Then, plugging equation (2.6) and (2.16) inside equation (2.8), we obtain the following Lagrangian for the representative consumer:

$$\mathcal{L}(\mathbf{q}, H) = \mathbf{q}' \mathbf{b} - \frac{1}{2} \mathbf{q}' [\mathbf{I} + \alpha (\mathbf{A}' \mathbf{A} - \mathbf{I})] \mathbf{q} - H - \lambda (\mathbf{q}' \mathbf{p} - H - \Pi) \quad (2.17)$$

The choice of labor hours as the numéraire immediately pins down the Lagrange multiplier $\lambda = 1$. Then, the consumer chooses a demand function $\mathbf{q}(\mathbf{p})$ to maximize the following consumer surplus function:

$$S(\mathbf{q}) \stackrel{\text{def}}{=} \mathbf{q}' (\mathbf{b} - \mathbf{p}) - \frac{1}{2} \mathbf{q}' [\mathbf{I} + \alpha (\mathbf{A}' \mathbf{A} - \mathbf{I})] \mathbf{q} \quad (2.18)$$

Let us now define the concept of *cosine similarity*.

Definition 2. We call the dot product $\mathbf{a}'_i \mathbf{a}_j$ the *cosine similarity* between i and j .

The rationale for this nomenclature is that – geometrically – $\mathbf{a}'_i \mathbf{a}_j$ measures the cosine of the angle between vectors \mathbf{a}_i and \mathbf{a}_j in the space of common characteristics \mathbb{R}^k . Hence, the cosine similarity ranges from zero to one. Because, by definition:

$$(\mathbf{A}' \mathbf{A})_{ij} = \mathbf{a}'_i \mathbf{a}_j \quad (2.19)$$

the matrix $\mathbf{A}' \mathbf{A}$ contains the *cosine similarities* between all firm pairs. A higher cosine similarity implies that two products provide a more overlapping mix of characteristics, and this reflects in patterns of product substitution: if $\mathbf{a}'_i \mathbf{a}_j > \mathbf{a}'_i \mathbf{a}_k$, an increase in the supply of product i leads to a larger decline in the marginal utility of product j than it does on the marginal utility of product k .

Figure 1 helps visualize this setup for the simple case of two firms—1 and 2—competing in the space of two common characteristics A and B. As can be seen in the figure, both firms exist as vectors on the unit circle (with more than three characteristics, it would be a hypersphere instead). The cosine similarity $\mathbf{a}'_i \mathbf{a}_j$

FIGURE 1: EXAMPLE PRODUCT SPACE: TWO FIRMS, TWO CHARACTERISTICS




FIGURE NOTES: The following diagram exemplifies the hedonic demand model, for the simple case where there are only two product characteristics (A and B) and only two competitors (1 and 2). Each firm exists as a vector on the unit hypersphere of product characteristics (in this example, we have a circle). The dot product $\mathbf{a}_i' \mathbf{a}_j$ equals the cosine of the angle θ . The tighter the angle, the higher the cosine similarity, and the larger (in absolute value) the inverse cross-price elasticity of demand.

captures the width of the angle θ . An increase in the cosine of the angle θ implies a lower angular distance, and therefore a more overlapping set of common characteristics.

The assumption that \mathbf{a}_i has unit length is a normalization assumption on volumetric units (kilograms, pounds, gallons, etc.). The normalization consists in picking, for each good i , the volume unit so that i is geometrically represented by a point on the k -dimensional hypersphere.

We can streamline the notation further by defining:

$$\Sigma \stackrel{\text{def}}{=} \alpha (\mathbf{A}' \mathbf{A} - \mathbf{I}) \tag{2.20}$$

then the demand and inverse demand functions are given by:

$$\text{Aggregate demand : } \mathbf{q} = (\mathbf{I} + \Sigma)^{-1} (\mathbf{b} - \mathbf{p}) \tag{2.21}$$

$$\text{Inverse demand : } \mathbf{p} = \mathbf{b} - (\mathbf{I} + \Sigma) \mathbf{q} \tag{2.22}$$

Notice that the quantity sold by each firm may affect the price of the output sold by every other firm in the economy (unless the matrix Σ is null). The derivative $\partial p_i / \partial q_j$ is proportional to $\mathbf{a}'_i \mathbf{a}_j$, the product similarity between i and j . The closer these two firms are in the product characteristics space, the larger is this derivative in absolute value. Because $\mathbf{A}'\mathbf{A}$ is symmetric, we have $\partial q_i / \partial p_j = \partial q_j / \partial p_i$ by construction. My rationale for using a linear demand is discussed at length in Appendix G.

In terms of elasticities, we have:

$$\text{Inverse cross - price elasticity of demand : } \frac{\partial \log p_i}{\partial \log q_j} = -\frac{q_j}{p_i} \cdot \sigma_{ij} \quad \forall i \neq j \quad (2.23)$$

$$\text{Cross - price elasticity of demand : } \frac{\partial \log q_i}{\partial \log p_j} = -\frac{p_j}{q_i} \cdot (\mathbf{I} + \Sigma)_{ij}^{-1} \quad (2.24)$$

It is worth stopping to inspect equation (2.24) more closely. The first thing to notice is that the cross-price demand elasticities depend on the inverse $(\mathbf{I} + \Sigma)^{-1}$. This implies that, while cosine similarities are positive by construction, it is entirely possible for goods to be complements. This property of the model is discussed at length in Section 5.

Next, let us consider the case $i = j$, where (2.24) simply becomes the own residual demand elasticity. The first major difference between the GHL demand system and CES is that, while in CES the own demand elasticity is equal to a constant, here the own demand elasticity is an equilibrium object (as it depends on \mathbf{q}) and will generally differ among firm pairs. This implies that, unlike CES, this demand system produces heterogenous markups. In fact, we can see that two forces drive cross-sectional differences in market power across firms. The more familiar one is the incomplete passthrough from marginal cost to prices: that is, larger firms (high q_i) charge higher markups. The second force, which is instead a feature of hedonic demand models, is asymmetric product differentiation. That is, firms that produce “unique” products, as measured by the term $(\mathbf{I} + \Sigma)_{ii}^{-1}$, face a less elastic residual demand.

Next, I define the economic profits π_i as follows:

$$\begin{aligned} \pi_i(\mathbf{q}) &\stackrel{\text{def}}{=} p_i(\mathbf{q}) \cdot q_i - h_i \\ &= q_i (b_i - c_i) - \left(1 + \frac{\delta_i}{2}\right) q_i^2 - \sum_{j \neq i} \sigma_{ij} q_i q_j \end{aligned}$$

Firms compete à la Cournot: each firm i strategically chooses its output volume q_i by taking as given the output of all other firms. By taking the profit vector as a payoff function and the vector of quantities produced \mathbf{q} as a strategy profile, I have implicitly defined a *linear-quadratic network game* (Ballester, Calvó-Armengol and Zenou, 2006, henceforth BCZ). The reason is that the matrix Σ can be conceptualized as the adjacency matrix of a weighted network: in this specific instance, it is the network of product market rivalry relationships that exists among the firms, based on the substitutability of their products.

Linear-quadratic network games belong to a larger class of games known as “potential games” (Monderer and Shapley, 1996): the key feature of potential games is that they can be described by a scalar function $\Phi(\mathbf{q})$, which we call the game’s *potential*. The potential function can be thought of, intuitively, as the objective function of the *pseudo-planner* problem that is solved by the Nash equilibrium allocation. The potential function is shown below, together with the aggregate profit function $\Pi(\mathbf{q})$ and the aggregate welfare function $W(\mathbf{q})$:

$$\begin{aligned}
\text{Aggregate Profit : } \Pi(\mathbf{q}) &= \mathbf{q}'(\mathbf{b} - \mathbf{c}) - \mathbf{q}'\left(\mathbf{I} + \frac{1}{2}\mathbf{\Delta} + \mathbf{\Sigma}\right)\mathbf{q} \\
\text{Cournot Potential : } \Phi(\mathbf{q}) &= \mathbf{q}'(\mathbf{b} - \mathbf{c}) - \mathbf{q}'\left(\mathbf{I} + \frac{1}{2}\mathbf{\Delta} + \frac{1}{2}\mathbf{\Sigma}\right)\mathbf{q} \\
\text{Total Surplus : } W(\mathbf{q}) &= \mathbf{q}'(\mathbf{b} - \mathbf{c}) - \frac{1}{2} \cdot \mathbf{q}'(\mathbf{I} + \mathbf{\Delta} + \mathbf{\Sigma})\mathbf{q}
\end{aligned} \tag{2.25}$$

$$\text{where } \mathbf{\Delta} \stackrel{\text{def}}{=} \begin{bmatrix} \delta_1 & 0 & \cdots & 0 \\ 0 & \delta_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \delta_n \end{bmatrix} \tag{2.26}$$

The three functions in equation (2.25) are visually similar to each other; they only differ by the scalar weight applied to the quadratic terms. The Cournot potential Φ is somewhat of a hybrid between the aggregate profit Π and the total surplus W : the diagonal elements of the quadratic term are the same as the aggregate profit function, while the off-diagonal terms are the same as the aggregate surplus function. By maximizing the potential $\Phi(\mathbf{q})$, we find the Cournot-Nash equilibrium. I shall assume all these three functions are concave. Because the oligopolists in this model will be actual firms in the data (who produce positive output by definition) we can look directly at the unique internal solution.

Proposition 1. *The Cournot-Nash equilibrium of the game described above is \mathbf{q}^Φ – the maximizer of the potential function $\Phi(\cdot)$:*

$$\mathbf{q}^\Phi \stackrel{\text{def}}{=} \arg \max_{\mathbf{q}} \Phi(\mathbf{q}) = (2\mathbf{I} + \mathbf{\Delta} + \mathbf{\Sigma})^{-1}(\mathbf{b} - \mathbf{c}) \tag{2.27}$$

Proof. The derivation of the potential function, as well as the proof that its maximizer \mathbf{q}^Φ is the genuine Nash equilibrium, appear in Appendix A. \square

Equation (2.27), which characterizes the Cournot-Nash equilibrium, tells us which factors determine the size of each firm in equilibrium. The diagonal matrix $\mathbf{\Delta}$, which contains the slopes of the marginal cost functions, captures economies of scale. $\mathbf{\Sigma}$ is the adjacency matrix of the network of product rivalries. \mathbf{b} and \mathbf{c} are, respectively, the demand and supply function intercepts. Hence, $(b_i - c_i)$ is simply the marginal surplus of the very first unit produced by firm i ; also, b_i can be interpreted as a measure of vertical product differentiation (quality).

BCZ show that another way to interpret equation (2.27) is as a measure of *network centrality* – specifically, that developed by Katz (1953) and Bonacich (1987). The intuition is that firms that are more “isolated” in the network of product similarities face less product market competition and behave more like monopolists. Centrality measures are a recurring feature of the literature on networks in macroeconomics (see Carvalho and Tahbaz-Salehi, 2019). In Appendix B, I discuss in further detail the link between Nash equilibrium and network centrality.

The discrepancy between the potential function and the total-surplus function implies that the network Cournot game delivers an equilibrium allocation that is not socially-optimal. A benevolent social planner can theoretically improve on the market outcome for two reasons. First, they can coordinate output choices across firms; second, they can internalize consumer surplus.

2.3. Consumer Surplus, the Shapley Value and Surplus Appropriation

How does a firm's position in the product network affect its ability to appropriate surplus? To answer this question, we must first answer the question of how to attribute consumer surplus to specific firms. Luckily, there is a natural, economically-meaningful way to break down consumer surplus and attribute, to every firm, a fraction of it: is the Shapley Value.

To see why this measure makes economic sense, we start by computing the marginal consumer surplus generated by q_i units of good i when all other firms supply $\hat{\mathbf{q}}_{-i}$. We know from basic price theory that this quantity can be computed by integrating the vertical distance between the residual demand and the purchase price p_i ⁶:

$$\begin{aligned} \text{CS}_i(q_i, p_i, \hat{\mathbf{q}}_{-i}) &= \int_0^{q_i} \left(b_i - \hat{q}_i - \sum_{j \neq i} \sigma_{ij} \hat{q}_j - p_i \right) d\hat{q}_i \\ &= q_i (b_i - p_i) - \frac{1}{2} q_i^2 + \sum_{j \neq i} \sigma_{ij} q_i \hat{q}_j \end{aligned} \quad (2.28)$$

We can then compute s_i , the Shapley Value in terms of consumer surplus, by taking the average of (2.28) across all possible allocations where each firm $j \neq i$ produces either $\hat{q}_j = q_j$ or $\hat{q}_j = 0$:

$$s_i \stackrel{\text{def}}{=} \frac{1}{2^{(n-1)}} \sum_{\iota_1=0}^1 \sum_{\iota_2=0}^1 \cdots \sum_{\iota_n=0}^1 \left[q_i (b_i - p_i) - \frac{1}{2} q_i^2 + \sum_{j \neq i} \iota_j \sigma_{ij} q_i q_j \right] \quad (2.29)$$

because we know that each firm j produces q_j in exactly half of the allocations considered, the expression above simplifies to:

$$s_i = q_i (b_i - p_i) - \frac{1}{2} \left(q_i^2 + \sum_{j \neq i} \sigma_{ij} q_i q_j \right) \quad (2.30)$$

by plugging p_i (as a function of q_i), this expression further simplifies to:

$$s_i \stackrel{\text{def}}{=} \frac{1}{2} \left(q_i^2 + \sum_{j \neq i} \sigma_{ij} q_i q_j \right) \quad (2.31)$$

Because it is a Shapley Value, one of s_i 's desirable properties is that it will always aggregate to total consumer surplus, that is:

$$S = \sum_i s_i \quad (2.32)$$

We shall therefore call s_i , going forward, the consumer surplus generated by firm i . By plugging the equilibrium expression for

We can similarly define a firm-level *total* surplus function, which specifies for every firm i a certain share w_i of the total surplus $W(\mathbf{q})$:

⁶When integrating consumer surplus, we must remember to treat p_i as a constant, since consumers are price-takers.

$$\begin{aligned}
w_i &\stackrel{\text{def}}{=} \pi_i + s_i \\
&= q_i (b_i - c_i) - \frac{1}{2} \left[(1 + \delta_i) q_i^2 + \sum_{j \neq i} \sigma_{ij} q_i q_j \right]
\end{aligned} \tag{2.33}$$

Next, we define a suitable measure of market share for our network oligopoly model, and derive an equation that links this quantity to firm's ability to appropriate surplus in the form of monopoly profits.

Definition 3. I define γ_i , the *weighted market share* of firm i , as follows:

$$\gamma_i \stackrel{\text{def}}{=} \frac{q_i}{q_i + \sum_j \sigma_{ij} q_j} \tag{2.34}$$

Notice that, under homogenous products ($\sigma_{ij} = 1 \forall i, j$) this is simply the market share of firm i . It is possible to show that the ratio of firm profits π_i to consumer surplus (measured using the Shapley Value s_i) is proportional to the weighted market share γ_i .

Proposition 2. *In the Cournot-Nash equilibrium allocation, the ratio of profits to consumer surplus for firm i is proportional to its weighted market share - specifically:*

$$\frac{\pi_i}{s_i} = (2 + \delta_i) \gamma_i \tag{2.35}$$

Proof. See Appendix K. □

The similarity-weighted market share γ_i plays, in the network Cournot model, a similar role as the Herfindahl index (HHI) in the canonical Cournot model. Proposition 2 reflects the fact that, in my model, there are no clearly-defined industry boundaries. This is also the case in the real world: if we consider antitrust lawsuits for example, a major object of litigation is the market's definition. Defendants (alleged monopolies) have an incentive to define the relevant market broadly, while plaintiffs have an incentive to define the relevant market narrowly. In my model, firms exist in a continuous space of product characteristics. Hence, there is no uniquely-defined peer group that we can compare each firm to. To understand how dominant firm i is, we need to compare its market share vis-à-vis every other firm in the economy, weighting each of them by their distance in the space of product characteristics.

The HHI can be seen as a special case of the weighted market share: as a measure of surplus appropriation, it is only valid in the special case where similarity scores are dichotomous (implying sharp industry boundaries) and firms are exchangeable ($b_i - c_i$ is constant across firms).

2.4. Markups and Product Market Centrality

In this model, we can think of market power being shaped by two forces. Firms more productive firms become larger in the Nash-Cournot equilibrium: because of imperfect passthrough these firms end up charging higher markups. Second, firms that are more central in the network of product similarities face increased competition (all else being equal), and thus a more elastic residual demand: these firms charge lower markups.

We next show how to decompose a firm's markup into two statistics that summarize these forces. We start by defining product market centrality. To define centrality, we start from the residual demand for firm

i , which is given by:

$$q_i = b_i - \left(\sum_{j=1}^n \sigma_{ij} q_j \right) - p_i \quad (2.36)$$

Equation (2.36) shows that, from the point of view of firm i , an increase in the similarity-weighted supply of competitors induces a parallel downward shift in the residual demand for firm i 's product. Because i takes the term in brackets as given, we can re-write equation (2.36) at no loss as:

$$q_i = \frac{b_i}{\chi_i} - p_i \quad (2.37)$$

where χ_i , which we call the *Product Market Centrality*, by definition equates (2.36) and (2.37).

Definition 4. We define χ_i , the Product Market Centrality of firm i , as the parallel downward shift in firm i 's residual demand induced by the combined supply of all potential competitors – formally:

$$\chi_i \stackrel{\text{def}}{=} \frac{b_i}{b_i - \sum_{j=1}^n \sigma_{ij} q_j} \quad (2.38)$$

We can then re-write equation (2.27) as:

$$q_i = \frac{1}{2} \left(\frac{b_i}{\chi_i} - \text{MC}_i \right) \quad (2.39)$$

We call χ_i a measure of centrality for two reasons. Firstly, the vector $(\chi_1, \chi_2, \dots, \chi_n)$ “summarizes” the entire matrix of cross-price derivatives into an n -dimensional vector; the idea is that, in order to solve its own profit maximization problem, firm i does not need to know all of the cross-price derivatives (σ_{ij}) , it only needs to know its own residual demand. Second, the summation term in the denominator of equation (2.38) is just the output-weighted degree centrality of firm i . Hence, the product market centrality is a monotonic transform of a measure of degree centrality, normalized by b_i .

Definition 5. We define ω_i –the “quality-adjusted productivity” of firm i – as the ratio between the marginal utility of the very first unit produced (b_i) and the marginal cost – formally:

$$\omega_i \stackrel{\text{def}}{=} \frac{b_i}{\text{MC}_i} \quad (2.40)$$

We can now express firm i 's markup μ_i (the price-marginal cost ratio) in terms of the ratio ω_i/χ_i .

Proposition 3. *The equilibrium markup μ_i is equal to the average of one and the ratio of the quality-adjusted productivity to the product market centrality:*

$$\mu_i \stackrel{\text{def}}{=} \frac{p_i}{\text{MC}_i} = \frac{1}{2} \left(1 + \frac{\omega_i}{\chi_i} \right) \quad (2.41)$$

Proof. Plug $(q_i^\Phi = p_i - \text{MC}_i)$ inside equation (2.39), divide both sides by MC_i , substitute (2.40) and rearrange. \square

2.5. Market Structure Counterfactuals

A key application of my theoretical model is to study how welfare statistics - such as total surplus - respond to changes in market structure. What that means is that, having made the required assumption that firms

compete by a well-defined set of rules (thus far we have assumed Cournot oligopoly), we can then consider counterfactuals in which the same firms play by a different set of rules. In this Subsection, I define four of these counterfactuals: each of these counterfactuals corresponds to the solution of a specific maximization problem.⁷

The first counterfactual that I consider is *perfect competition*: firms act as atomistic producers, and price all units sold at marginal cost.

Definition 6. The *Perfect Competition* allocation \mathbf{q}^W is defined as the maximizer of the aggregate total surplus function $W(\mathbf{q})$:

$$\mathbf{q}^W \stackrel{\text{def}}{=} \arg \max_{\mathbf{q}} W(\mathbf{q}) = (\mathbf{I} + \mathbf{\Delta} + \mathbf{\Sigma})^{-1} (\mathbf{b} - \mathbf{c}) \quad (2.42)$$

The second counterfactual that I consider is called *Monopoly*: it represents a situation in which one agent (that does not internalize consumer surplus) has control over all the firms in the economy and maximizes aggregate profits.

Definition 7. The *Monopoly* allocation is defined as the maximizer of the aggregate profit function $\Pi(\mathbf{q})$:

$$\mathbf{q}^\Pi \stackrel{\text{def}}{=} \arg \max_{\mathbf{q}} \Pi(\mathbf{q}) = (2\mathbf{I} + \mathbf{\Delta} + 2\mathbf{\Sigma})^{-1} (\mathbf{b} - \mathbf{c}) \quad (2.43)$$

This allocation can be alternatively conceptualized as an economy with no antitrust policy, where firms have unlimited ability to coordinate their supply choices.

While the *Monopoly* counterfactual is an interesting limit case, using the model we can also study the welfare impact of mergers and collusion among specific firms.

When it comes to modeling mergers and collusions, the I.O. literature has used multiple approaches. Following Baker and Bresnahan (1985), I model mergers and collusion interchangeably as *coordinated pricing*. That is, I assume that the merger or the collusion does not affect the product range offered by the merging/colluding enterprises; instead, a single agent determines the output of the merging firms to maximize the joint profits.⁸

Lemma 1. Consider, without loss of generality, a merger or collusion between companies $\{1, 2, \dots, m\}$; then, partition the matrix $\mathbf{\Sigma}$ by separating the first m rows and columns as follows:

$$\mathbf{\Sigma} = \begin{bmatrix} \mathbf{\Sigma}_{11} & \mathbf{\Sigma}_{12} \\ \mathbf{\Sigma}_{21} & \mathbf{\Sigma}_{22} \end{bmatrix} \quad (2.44)$$

The post-merger equilibrium allocation maximizes the following modified potential function:

$$\Phi(\mathbf{q}) = \mathbf{q}' (\mathbf{b} - \mathbf{c}) - \mathbf{q}' \left(\mathbf{I} + \frac{1}{2} \mathbf{\Delta} \right) \mathbf{q} - \frac{1}{2} \cdot \mathbf{q}' \begin{bmatrix} 2\mathbf{\Sigma}_{11} & \mathbf{\Sigma}_{12} \\ \mathbf{\Sigma}_{21} & \mathbf{\Sigma}_{22} \end{bmatrix} \mathbf{q} \quad (2.45)$$

Proof. See Appendix K. □

⁷The closed-form expressions for the output vector \mathbf{q} which I provide below assume an internal solution. For my empirical analysis, I also compute a numerical solution that is subject to a non-negativity constraint on \mathbf{q} and I verify it is approximately equal to the unconstrained solution (error < 0.1% for the total surplus function in Perfect Competition). The non-negativity constraint binds for very few firms.

⁸This approach is particularly attractive, from a tractability standpoint, in a setting like this, where products are differentiated. The alternative would be to make some heroic assumptions about the nature and characteristics of the product supplied by the combined entity.

The maximizer of the re-defined $\Phi(\mathbf{q})$, which corresponds to the post-merger equilibrium allocation, is:

$$\mathbf{q}^\Phi = \left(2\mathbf{I} + \mathbf{\Delta} + \begin{bmatrix} 2\mathbf{\Sigma}_{11} & \mathbf{\Sigma}_{12} \\ \mathbf{\Sigma}_{21} & \mathbf{\Sigma}_{22} \end{bmatrix} \right)^{-1} (\mathbf{b} - \mathbf{c}) \quad (2.46)$$

That is, to simulate the new equilibrium following a merger or a collusion among existing firms, one only needs to amend the potential function by doubling the off-diagonal quadratic terms corresponding to the merging firms. It is easily verified that when all firms are merged, $\Phi(\mathbf{q})$ simply becomes the aggregate profit function $\Pi(\mathbf{q})$, and the equilibrium allocation converges to the Monopoly counterfactual (equation 2.43).

Another interesting counterfactual is one in which resources are allocated efficiently but the labor supply is fixed. That is, the social planner maximizes the aggregate surplus function subject to the constraint of using no more labor than in the observed Cournot equilibrium.

Definition 8. I define the resource-efficient counterfactual \mathbf{q}^H as the solution to the following constrained maximization problem:

$$\mathbf{q}^H \stackrel{\text{def}}{=} \arg \max_{\mathbf{q}} W(\mathbf{q}) \quad \text{s.t.} \quad H(\mathbf{q}) = H(\mathbf{q}^\Phi) \quad (2.47)$$

Setting up the Lagrangian and using $(1 - \mu)$ as the Lagrange multiplier, we find that the resource-efficient counterfactual takes the form:

$$\mathbf{q}^H = (\mathbf{I} + \mu\mathbf{\Delta} + \mathbf{\Sigma})^{-1} (\mathbf{b} - \mu\mathbf{c}) \quad (2.48)$$

where μ solves:

$$H(\mathbf{q}^H(\mu)) = H(\mathbf{q}^\Phi) \quad (2.49)$$

The Lagrange multiplier term μ turns out to be the common markup charged by all firms in the resource-efficient counterfactual.

Lemma 2. *The Resource-efficient counterfactual \mathbf{q}^H equalizes markups across firms.*

Proof. Let all firms price at a constant markup μ over marginal cost:

$$p_i = \mu \cdot \text{MC}_i \quad (2.50)$$

expanding the expression for the marginal cost and the equilibrium price we have:

$$\mathbf{b} - (\mathbf{I} + \mathbf{\Sigma})\mathbf{q} = \mu(\mathbf{c} + \mathbf{\Delta}\mathbf{q}) \quad (2.51)$$

rearranging the equation above we obtain (2.48). □

Because this counterfactual uses the same amount of labor as the observed equilibrium, by comparing welfare in this allocation to the first-best we can effectively break down the deadweight loss into two components – one linked to misallocation, the other linked to labor suppression. We can also interpret this counterfactual as the deadweight loss in an alternative model where the supply of labor is completely inelastic. Notice that when this allocation is not constrained by the labor supply (the Lagrange multiplier $1 - \mu$ is zero), the common markup is one (firms price at marginal cost) and the resource-efficient allocation coincides with perfect competition.

The counterfactuals considered thus far do not account for how a firm's incentives to participate in the market are affected by the intensity of competition. When the market moves from Cournot competition to

(say) Bertrand ⁹ or perfect competition, the resulting lower profits might insufficient to cover fixed costs, and therefore too low to justify a firm’s continued existence. If this is the case, perfect competition may not be a realistic benchmark in the long-run: this is the classical criticism of static welfare analysis.

Next, I construct an “efficient” allocation that takes into account (to the extent possible in a static model) these dynamic incentives. The starting point is again a benevolent social planner, to which we are adding a constraint, in the form of a participation condition on the firms’ side: firms have to be able (on average) to recover their fixed costs (F) at the optimum.¹⁰

Definition 9. The Second-Best Allocation \mathbf{q}^{2nd} is defined as the solution to the following constrained maximization problem:

$$\mathbf{q}^{2nd} \stackrel{\text{def}}{=} \arg \max_{\mathbf{q}} W(\mathbf{q}) \quad \text{s.t.} \quad \Pi(\mathbf{q}) \geq F \quad (2.52)$$

where

$$F \stackrel{\text{def}}{=} \sum_i f_i \quad (2.53)$$

Setting up the Lagrangian of this problem and imposing λ as the Lagrangian multiplier, we find that the resource-efficient counterfactual takes the form:

$$\mathbf{q}^{2nd} = \left[\frac{1 + 2\lambda}{1 + \lambda} \cdot (\mathbf{I} + \mathbf{\Sigma}) + \mathbf{\Delta} \right]^{-1} (\mathbf{b} - \mathbf{c}) \quad (2.54)$$

Assuming that the constraint binds at the optimum, the Lagrange multiplier λ solves:

$$\Pi(\mathbf{q}^{2nd}(\lambda)) = F \quad (2.55)$$

As the constraint is relaxed ($\lambda \rightarrow 0$), this counterfactual allocation converges to the first-best. When the constraint becomes arbitrarily tight ($\lambda \rightarrow \infty$), it converges to the *Monopoly* allocation.

In addition to the counterfactuals considered above, which admit closed-form solutions, we can simulate the introduction or the removal of granular firms. The latter can be trivially implemented by computing an allocation where a firm’s output is constrained to be zero. In order to simulate instead the introduction of new firms, we require additional assumptions or data. Namely, in order to simulate the introduction of an additional firm (let us label it firm zero), we would need to know the value of $(b_0 - c_0)$, as well as the firm’s similarity to every other firm in the economy (\mathbf{a}_{i0}). One such counterfactual is considered in Section (4).

2.6. Adding a Continuum of firms with Endogenous Entry

Next, I show how to expand the model to include a continuum of atomistic firms that behave competitively and can enter and exit endogenously. This extension of the model allows me to accomplish two things: 1) incorporate firms for whom we do not observe product similarity data – that is, foreign and private firms; 2) it allows to incorporate entry and exit in an otherwise static model. The idea is that we can model unobserved companies as atomistic firms.

The key to tractably integrating these atomistic firms in the model is an aggregation result. I describe

⁹The Bertrand model is covered in Appendix F.

¹⁰There are two reasons why I consider a constraint on aggregate profits rather than individual profits ($\pi_i \geq f_i$). The first is that such individual constraint is already violated by many firms in the observed (Cournot) equilibrium. The second reason is that adding individual constraints for each firms would make the optimization problem numerically intractable, since we would need to solve for thousands of constraints (one for each firm in the model).

these atomistic firms through a productivity distribution: the set of active atomistic firms is then characterized by a productivity cut-off value, in the style of Hopenhayn (1992).

Next, I show that these atomistic companies can be aggregated into a representative firm: variations in the size of the representative firm reflect the intensive margin of production as well as the extensive margin (the entry/exit of the atomistic firms). I index this representative firm $i = n + 1$, effectively adding a row and a column to the matrices $\mathbf{A}'\mathbf{A}$ and $\mathbf{\Delta}$ and adding one dimension to the vector \mathbf{b} .

Proposition 4. *Assume that there is a mass one of potential entrants that are indexed by a productivity parameter $z \in (\underline{z}, \infty)$ and that produce a homogeneous good using the following quadratic cost function:*

$$h(z) = \frac{\delta(z)}{2} \cdot q^2(z) \quad (2.56)$$

with $\underline{z} > 0$ and

$$\delta(z) = \frac{1}{z} \quad (2.57)$$

Assume also that the firms face cost of entry equal to one unit of labor and that the probability density of type- z potential entrants is given by

$$f(z) = \frac{\beta - 1}{z^{\beta+1}} \quad (2.58)$$

implying that z follows a Pareto distribution with shape parameter β and scale parameter $\underline{z} \stackrel{\text{def}}{=} [(\beta - 1) / \beta]^{\frac{1}{\beta}}$.¹¹ Then, as the parameter β converges down to 1, the cost function of the corresponding aggregate representative firm is approximated by

$$h_{n+1} = \frac{q_{n+1}^2}{2} \quad (2.59)$$

where and h_{n+1} and q_{n+1} are, respectively, the labor input and the output of the representative firm, and the productivity cutoff for entry converges to $z_{\min} = \frac{1}{q_{n+1}}$.

Proof. See Appendix K. □

Because employment and revenues are proportional to z , it follows that, if the assumptions above are respected, both the revenue and employment distribution of firms also approximate a Pareto distribution with shape parameter $\beta = 1$, sometimes called a Zipf Law.

Although this might look like a knife-edge assumption, it is not. It is a well-documented empirical regularity that the size distribution of firms closely approximates a Pareto distribution with shape parameter $\beta = 1$. This stylized fact was confirmed to hold for both the employment *and* the revenue distribution of US firms by Axtell (2001), using Census micro-data.

Because the representative firm behaves competitively, its first order condition will differ from that of granular firms $\{1, 2, \dots, n\}$. The latter maximize individual profits:

$$\pi'_i(q_i) = 0 \quad \text{for } i = 1, 2, \dots, n \quad (2.60)$$

The representative firm, on the other hand, prices at marginal cost, and therefore maximizes total surplus:

$$W'(q_i) = 0 \quad \text{for } i = n + 1 \quad (2.61)$$

¹¹While the revenue and employment distribution of US firms approximates a Pareto Distribution with scale parameter equal to one (a Zipf Law), this distribution has the undesirable property that its mean (and therefore q_{n+1} and h_{n+1}) grows unboundedly as $\beta \rightarrow 1^+$. This particular choice of the scale parameter ensures that q_{n+1} and h_{n+1} integrate to a finite number as $\beta \rightarrow 1^+$.

We can write the full system of first order conditions in linear algebra notation as:

$$0 = \begin{bmatrix} \mathbf{b}^{(n)} - \mathbf{c}^{(n)} \\ b_{n+1} - c_{n+1} \end{bmatrix} + \left(\begin{bmatrix} 2\mathbf{I} & \mathbf{0} \\ \mathbf{0} & 1 \end{bmatrix} + \mathbf{\Sigma} + \mathbf{\Delta} \right) \begin{bmatrix} \mathbf{q}^{(n)} \\ q_{n+1} \end{bmatrix} \quad (2.62)$$

where $c_{n+1} = 0$, $\delta_{n+1} = 1$ and the superscript (n) identifies the sub-vector corresponding to the granular firms. A simpler way to rewrite this set of equations is

$$0 = \mathbf{b} - \mathbf{c} - (\mathbf{I} + \mathbf{G} + \mathbf{\Sigma} + \mathbf{\Delta}) \mathbf{q} \quad (2.63)$$

where \mathbf{G} is a diagonal matrix that identifies granular firms – that is, whose diagonal elements equal 1 for firms 1 to n and to 0 for firm $n + 1$:

$$\mathbf{G} = \begin{bmatrix} 1 & 0 & \cdots & 0 & 0 \\ 0 & 1 & \cdots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & 1 & 0 \\ 0 & 0 & \cdots & 0 & 0 \end{bmatrix} \quad (2.64)$$

The potential function for the model that includes the representative firm is:

$$\Phi(\mathbf{q}) = \mathbf{q}'(\mathbf{b} - \mathbf{c}) - \frac{1}{2} \mathbf{q}'(\mathbf{I} + \mathbf{G} + \mathbf{\Sigma} + \mathbf{\Delta}) \mathbf{q} \quad (2.65)$$

and the equilibrium quantity vector is:

$$\mathbf{q}^\Phi = (\mathbf{I} + \mathbf{G} + \mathbf{\Delta} + \mathbf{\Sigma})^{-1} (\mathbf{b} - \mathbf{c}) \quad (2.66)$$

3. Data and Identification/Calibration

In this section, I outline the data used to estimate the model in Section 2. Additional details are provided in Appendix D, which also contains Table 2, where model mapping and identification are summarized.

3.1. Firm Financials

My data source for firm financials is the Compustat database, which I access via the Wharton Research Data Services (WRDS) platform. From this database, I extract information on firm revenues, Costs of Goods Sold (COGS), Selling General and Administrative (SGA) costs, R&D expenditures and Property Plant and Equipment (PPE).

I follow (De Loecker, Eeckhout and Unger, 2020, henceforth DEU) in mapping accounting revenues to model revenues, COGS to variable costs, and in computing an estimate of fixed costs costs (f_i):

$$f_i \leftarrow \text{SGA}_i + \text{Property Plant \& Equipment}_i \times \text{User Cost of Capital} \quad (3.1)$$

3.2. Text-Based Product Similarity

The key data ingredient that we need, in order to estimate my model, is the matrix of product similarities $\mathbf{A}'\mathbf{A}$. The empirical counterpart of this object is provided by Hoberg and Phillips (2016, henceforth HP).

HP created a publicly-available database that provides product cosine similarities for the universe of public corporations in the United States. These cosine similarities originate from natural language processing (NLP) of 10-K filings, and are time-varying. A complete matrix of similarities is provided for every year, beginning in 1997.

The 10-K is a mandatory form that is filed by American public corporations with the U.S. Securities and Exchange Commission on a yearly basis. Item 1 of the 10-K is a long and detailed description of the product or service sold by the company. HP's product cosine similarities are constructed by comparing these textual product descriptions.

I briefly outline the construction of this dataset. HP start by building a vocabulary of 61,146 words that firms use to describe the characteristics of their products.¹² Based on this vocabulary, HP produce, for each firm i , a vector of word frequencies \mathbf{o}_i . Each of component of this vector corresponds to a word in HP's vocabulary, and is equal to the number of times that word appears in firm i 's 10-K product description:

$$\mathbf{o}_i = \begin{bmatrix} o_{i,1} \\ o_{i,2} \\ \vdots \\ o_{i,61146} \end{bmatrix} \quad (3.2)$$

Similar to the model in Section 2, this vector is then normalized (divided by the Euclidean norm). We have thus obtained the empirical counterpart of \mathbf{a}_i :

$$\mathbf{a}_i = \frac{\mathbf{o}_i}{\|\mathbf{o}_i\|} \quad (3.3)$$

finally, all \mathbf{a}_i vectors are dot-multiplied to obtain $\mathbf{A}'\mathbf{A}$:

$$\mathbf{A}'\mathbf{A} = \begin{bmatrix} \mathbf{a}'_1\mathbf{a}_1 & \mathbf{a}'_1\mathbf{a}_2 & \cdots & \mathbf{a}'_1\mathbf{a}_n \\ \mathbf{a}'_2\mathbf{a}_1 & \mathbf{a}'_2\mathbf{a}_2 & \cdots & \mathbf{a}'_2\mathbf{a}_n \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{a}'_n\mathbf{a}_1 & \mathbf{a}'_n\mathbf{a}_2 & \cdots & \mathbf{a}'_n\mathbf{a}_n \end{bmatrix} \quad (3.4)$$

Next, we state explicitly the key identifying assumption that allows us to map HP's cosine similarity dataset to $\mathbf{A}'\mathbf{A}$.

Assumption (Identification). *the normalized frequencies of the words from Hoberg and Phillips' vocabulary in 10-K business descriptions approximate the product loadings on common characteristics.*

This assumption is obviously is strong, and will need to be validated empirically before proceeding to perform welfare measurements using the model. However, the implications of this assumption for identification

¹²I report here verbatim the methodology description from the original paper by Hoberg and Phillips (2016):“[...] In our main specification, we limit attention to nouns (defined by Webster.com) and proper nouns that appear in no more than 25 percent of all product descriptions in order to avoid common words. We define proper nouns as words that appear with the first letter capitalized at least 90 percent of the time in our sample of 10-Ks. We also omit common words that are used by more than 25 percent of all firms, and we omit geographical words including country and state names, as well as the names of the top 50 cities in the United States and in the world. [...]”

are powerful: to the extent that this assumption is satisfied, HP’s matrix of cosine similarities is not simply “resembling” the object $\mathbf{A}'\mathbf{A}$: it is that exact object, there is literally no difference between the theoretical object and the data counterpart.

The fact that all publicly-traded firms in the United States are required to file a 10-K form makes the this data set unique, in that it covers the near entirety (97.8%) of the Compustat universe. HP use these cosine similarities to produce a dynamic industry classification, called TNIC, which they extensively validate: one way they validate their data (in the paper that presents their methodology) is by using another dataset called CapitalIQ. This dataset provides dummy variables for a sub-set of Compustat firm pairs which identify product market rivalry relationships; they are based on corporate filings as well as other sources (no time variation is available in this dataset). HP show that TNIC outperforms SIC and NAICS in predicting competitor pairs in CapitalIQ.

Since their introduction in 2011, HP’s industry classifications have become standard in the empirical corporate finance literature, where they have replaced NAICS and SIC for a variety of applications. A major reason for this methodological shift is that HP’s dataset addressed an important limitation of traditional industry classifications. While these have often been used (for lack of better alternatives) to capture product market competition¹³, it is well-known that they are based on the concept of *production process* similarity, not product similarity¹⁴. This is also one reason why, in the I.O. and Antitrust literature, NAICS and SIC are generally only used to estimate production functions¹⁵.

There are other factors that differentiate HP’s database from traditional industry classifications. While NAICS and SIC are binary (firms are either in the same industry or different industries), HP’s database also provides continuous similarity scores ranging from zero to one, thus accommodating the inherent fuzziness of product market rivalries. While NAICS and SIC are seldom updated, HP’s similarity scores are updated yearly. While NAICS and SIC are arbitrarily assigned (Chen et al., 2016 show that firms strategically manipulate their industry classifications), HP’s similarity scores are rule-driven and incentive-compatible: executives face legal liability for misrepresenting company information in SEC filings.

I begin my empirical analysis by visualizing HP’s dataset. To do so, I have to reduce the dimensionality of the dataset from 61,146 (the number of words in the HP’s vocabulary) to two. I do so using the algorithm of Fruchterman and Reingold (1991, henceforth FR), which is widely used in network science to visualize weighted networks¹⁶.

The result of this exercise is Figure 2: every dot in the graph is a publicly traded firm as of 2004. Firm pairs that have a high cosine similarity appear closer, and are joined by a thicker line. Conversely, firms that are more dissimilar are not joined, and are more distant. From the graph, we can see that the distribution of firms over the space of product characteristics is manifestly uneven: some areas are significantly more densely populated with firms than others. Also, the network displays a pronounced community structure: large groups of firms tend to cluster in certain areas of the network.

In Appendix C, I show that this visualization is not an artifact of dimensionality reduction or measurement error: notwithstanding the dimensionality reduction, a remarkable degree of overlap exists between

¹³Before HP’s data was published, Bloom, Schankerman and Van Reenen (2013) constructed cosine similarities to estimate R&D spillovers. They used Compustat Segments data, which is based on NAICS/SIC industries. This dataset’s coverage of Compustat is insufficient to estimate my model (not enough firms/years available).

¹⁴See the following [Bureau of Labor Statistics Guide](#).

¹⁵For example: DEU’s method to compute markups uses production function estimates for NAICS industries.

¹⁶The algorithm models the network nodes as particles, letting them dynamically arrange themselves on a bidimensional surface as if they were subject to attractive and repulsive forces. One known shortcoming of this algorithm is that it is sensitive to the initial configurations of the nodes, and it can have a hard time uncovering the cluster structure of large networks. To mitigate this problem, and to make sure that the cluster structure of the network is properly displayed, I pre-arrange the nodes using the OpenOrd algorithm (which was developed for this purpose) before running FR.

FIGURE 2: NETWORK VISUALIZATION OF THE HOBERG-PHILLIPS DATASET




FIGURE NOTES: The following diagram is a two-dimensional representation of the network of product similarities computed by Hoberg and Phillips (2016), which is used in the estimation of the model presented in Section 2. The data covers the universe of Compustat firms in 2004. Firm pairs that have thicker links are closer in the product market space. These distances are computed in a space that has approximately 61,000 dimensions. To plot this high-dimensional object over a plane, I applied the gravity algorithm of Fruchterman and Reingold (1991), which is standard in social network analysis.

the macro-clusters of this network and broad economic sectors. In addition, this exercise allows me to independently validate HP's product similarity data.

3.3. Identification of Output, Prices and Cost Intercept

All of the unobserved variables in the model are identified subject to two parameters: (α) which controls the degree of horizontal differentiation between goods; and the diagonal matrix (Δ), which controls returns to scale. I will first show how to identify the remaining variables conditional on these two parameters and then I will illustrate my calibration procedure for α and Δ .

I measure the matrix $\mathbf{A}'\mathbf{A}$ using Hoberg and Phillips (2016)'s cosine similarity data. Given α , this provides an estimate of Σ . We can then identify real output q_i from revenues and total variable cost data:

$$q_i = \sqrt{\frac{\pi_i}{1 + \delta_i/2}} \quad \text{if } i \leq n \quad (3.5)$$

If the model includes a representative competitive firm $n + 1$, the identification of q_i for this firm will be different. Specifically, the marginal cost pricing condition ($p_{n+1} = \text{MC}_{n+1}$) implies that:

$$q_i = \sqrt{\pi_i + h_i} \quad \text{if } i = n + 1 \quad (3.6)$$

where $(\pi_{n+1} + h_{n+1})$ is measured as the Gross Value Added of private and foreign firms, which I compute using the OECD Trade in Value Added (TiVA) Dataset.

Having identified q_i , we can then pin down the vector of prices and the cost function intercepts:

$$p_i = \frac{p_i q_i}{q_i} \quad c_i = \frac{h_i}{q_i} - \frac{\delta_i}{2} q_i \quad (3.7)$$

Finally, I identify the demand intercept b_i using equation (2.27):

$$\mathbf{b} = (2\mathbf{I} + \mathbf{\Delta} + \mathbf{\Sigma}) \mathbf{q} + \mathbf{c} \quad (3.8)$$

or, in the presence of an representative competitive firm:

$$\mathbf{b} = (\mathbf{I} + \mathbf{G} + \mathbf{\Delta} + \mathbf{\Sigma}) \mathbf{q} + \mathbf{c} \quad (3.9)$$

3.4. Calibration of α and $\mathbf{\Delta}$

The last step required to take the model to the data is to calibrate the scalar α and the diagonal matrix $\mathbf{\Delta}$. Let us start from the latter. To calibrate each diagonal element δ_i , we use the fact that the markup (price-marginal cost ratio) of firm i can be written as a function of observables (revenues, total variable costs) and δ_i – i.e. the markup μ_i is identified conditional δ_i .

Lemma 3. *The markup μ_i is equal to:*

$$\mu_i \stackrel{\text{def}}{=} \frac{p_i}{\text{MC}_i} = \frac{(2 + \delta_i) \cdot p_i q_i}{2 \cdot h_i + \delta_i \cdot p_i q_i} \quad (3.10)$$

Proof. See Appendix (K). □

DEU compute the revenue-weighted average markup for the same universe of companies (Compustat). My strategy for calibrating $\mathbf{\Delta}$ is to target their estimate. The detailed methodology for calibrating $\mathbf{\Delta}$ is outlined in detail in Appendix D.

To calibrate α , we rewrite equation (2.21) as:

$$\left| \frac{\partial \log p_i}{\partial \log q_j} \right| = \alpha \cdot \mathbf{a}'_i \mathbf{a}_j \frac{q_j}{p_i} \quad \forall i \neq j \quad (3.11)$$

By calibrating $\mathbf{\Delta}$, we have already pinned q_i and p_i . The matrix of cosine similarities of Hoberg and Phillips (2016) provides the empirical counterpart to $\mathbf{A}'\mathbf{A}$. Hence, the matrix of equilibrium cross-price demand elasticities is identified given α .

My strategy for calibrating α is to target target microeconomic estimates from the Industrial Organization literature. I obtain, for a number of firm pairs, estimates of the cross-price demand elasticity from

empirical IO studies that estimate the demand function econometrically. These estimates of the cross-price demand elasticity are then manually matched to the corresponding firm pair in Compustat. Finally, for each firm pair, I can obtain an estimate of α by rearranging equation (3.11):

$$\hat{\alpha}_{ij} = \left| \frac{\partial \log p_i}{\partial \log q_j} \right| \bigg/ \left(\mathbf{a}'_i \mathbf{a}_j \frac{q_i}{p_i} \right) \quad (3.12)$$

In the absence of mis-specification and measurement error, all these estimates $\hat{\alpha}_{ij}$ would return the same estimate. What I obtain in this case, instead, is a range of estimates. I calibrate α to the median value among these estimates, which is 0.05. The full methodology is presented in Appendix D, where I also discuss how the model fits non-targeted moments in the data.

4. Empirical Findings

In this section, I present the results of the estimation of my model. My baseline estimates reflect the model implementation that only includes granular firms (Compustat). In the next section, I discuss the robustness of my estimates to the inclusion of private and foreign firms as a continuum of atomistic firms that enter endogenously.

4.1. Welfare Statics

My first empirical exercise is to compute total surplus and to break it down into profits and consumer surplus. This is done for both the observed equilibrium (which is assumed to be a Nash-Cournot equilibrium) and the counterfactuals considered in Section 2. These estimates are all shown in Table 1.

I estimate that the publicly-traded firms earn an aggregate economic profit of \$5 trillion and produce an estimated total surplus of \$9.1 trillion. Consumer surplus is therefore estimated to be about \$4 trillion. About 55% of the total surplus produced is appropriated by the companies in the form of oligopoly profits. For context, the GDP of U.S. corporations in the same year (2017) is \$11 trillion¹⁷.

The first counterfactual I consider, *Perfect Competition*, appears in the second column. By comparing this counterfactual with the Cournot-Nash allocation we can see that the welfare costs of oligopoly are significant. Under perfect competition, aggregate surplus is significantly higher – \$10.2 trillion – hence, the deadweight loss amounts to about 11% of the total surplus.

While the effects of oligopoly on Pareto efficiency are significant, perhaps even more significant are the *distributional* effects. When firms price at marginal cost, a much larger share of the surplus goes to the consumer: \$8.2 trillion, more than double than in the Cournot allocation. This amounts to 80% of the total surplus.

The next counterfactual I analyze, the *Monopoly* counterfactual, appears in the third column: it represents a scenario in which all firms are controlled by a single decision-maker that coordinates supply choices. In this allocation, aggregate surplus is significantly lower than in the Network Cournot equilibrium allocation: \$8.2 trillion. Despite the decrease in aggregate welfare, profits are markedly higher: \$5.7 trillion. Consequently, consumer surplus is reduced to just \$2.5 trillion, a mere 33% of the total. One interpretation

¹⁷The difference between GDP and total surplus is that total surplus does not include the value of labor input but it does include the value of inframarginal consumption. GDP, on the other hand, includes the value of labor input but not the inframarginal value of consumption. In this model each unit of labor is paid exactly its marginal disutility, hence there is no inframarginal value of leisure.

TABLE 1: WELFARE STATISTICS (2017)

Welfare Statistic	Variable	Scenario				
		(1) Cournot-Nash	(2) Perfect Competition	(3) Monopoly	(4) Resource-Efficient	(5) Second-Best
	\mathbf{q}^Φ	\mathbf{q}^W	\mathbf{q}^Π	\mathbf{q}^H	\mathbf{q}^{2nd}	
Total Surplus (US\$ trillions)	$W(\mathbf{q})$	9.086	10.208	8.183	9.869	10.017
Aggregate Profits (US\$ trillions)	$\Pi(\mathbf{q})$	5.043	1.995	5.673	3.135	3.910
Consumer Surplus (US\$ trillions)	$S(\mathbf{q})$	4.043	8.213	2.510	6.735	6.106
Total Surplus / Perfect Competition	$\frac{W(\mathbf{q})}{W(\mathbf{q}^W)}$	0.890	1.000	0.802	0.967	0.981
Aggregate Profit / Total Surplus	$\frac{\Pi(\mathbf{q})}{W(\mathbf{q})}$	0.555	0.195	0.693	0.318	0.390
Consumer Surplus / Total Surplus	$\frac{S(\mathbf{q})}{W(\mathbf{q})}$	0.445	0.805	0.307	0.682	0.610

TABLE NOTES: The following table shows my estimates of aggregate profits, consumer surplus and total surplus in each of the counterfactual scenarios presented in Section 4.

of this exercise is that policies that prevent coordination between firms (antitrust) have a large positive impact on welfare.

Next, I consider the *Resource Efficient* counterfactual, in which the social planner maximizes total surplus subject to the constraint of not using more labor than the Cournot equilibrium.

In this scenario, markups across firms have been equalized, but not eliminated. By removing all dispersion in markups, this counterfactual targets the malallocative effects of oligopoly.

The total surplus produced in this counterfactual is \$9.9 trillion, about 3.3 percentage points lower than in perfect competition, and \$800 billion higher than the observed Cournot-Nash equilibrium. Most of the surplus produced – \$6.9 trillion, or 68% of the total – goes to the consumer; profits are reduced to \$3.1 trillion. Because labor is fixed, all the welfare gains with respect the Cournot equilibrium come from the *reallocation* of labor. Hence, an important take-away from this counterfactual is that a large share of the inefficiencies from oligopoly are driven by resource misallocation. A different way to say this is that the dispersion in markups (caused by oligopolistic competition with differentiated products) matters at least as much as the level of markups in determining the overall deadweight loss.

The last counterfactual that I consider is the *Second-Best*, in which a benevolent social planner maximizes aggregate surplus subject to an aggregate participation constraint (profits must cover fixed costs on average). In this counterfactual, total surplus is very close to the level achieved by the perfectly-competitive outcome: \$10 trillion. The main difference is the surplus split: the consumer receives \$6.1 trillion (two trillion less than under perfect competition but two more than under Cournot), while total profits amount to \$3.9 trillion.

In sum, my measurements suggest that the oligopoly power of U.S. public firms has significant consequences for aggregate welfare, and that it impacts consumer welfare through two channels: it increases the dispersion of markups, generating resource misallocation which raises the deadweight loss; it also increases the level of markups, which in turn affects how surplus is shared between producers and consumers.

4.2. Markup Decomposition

While there’s evidence in support of the fact that the markups of US corporations have increased in the past decades (De Loecker, Eeckhout and Unger, 2020), the reasons behind this increase are still very much up for debate: are higher markups the result of a right tail of highly-productive superstar firms pulling away from the rest the competition, or are they the result of a softening of product market competition (i.e. fewer substitutes)?

Equation (2.41) allows us to study this question by decomposing markups into two components reflecting these forces: by implementing the decomposition within our model, can investigate whether the rise of markup is driven by firms with a low product market centrality (χ_i) or whether it’s driven by higher quality-adjusted productivity (ω_i).

Defining the operator $\mathbb{E}(\cdot)$ as the (weighted) cross-sectional average of a variable, and letting $\text{cov}(\cdot)$ be the corresponding cross-sectional (weighted) covariance, we can write the following formula for the revenue-weighted markup, based on equation (2.41):

$$\mathbb{E}(\mu_{it}) = \frac{1}{2} [1 + \mathbb{E}(\omega_{it}) \cdot \mathbb{E}(\chi_{it}^{-1}) + \text{cov}(\omega_{it}, \chi_{it}^{-1})] \quad (4.1)$$

where t represents the year of observation. Differentiating with respect to time, we obtain the following

FIGURE 3: MARKUPS GROWTH DECOMPOSITION (1995-2019)

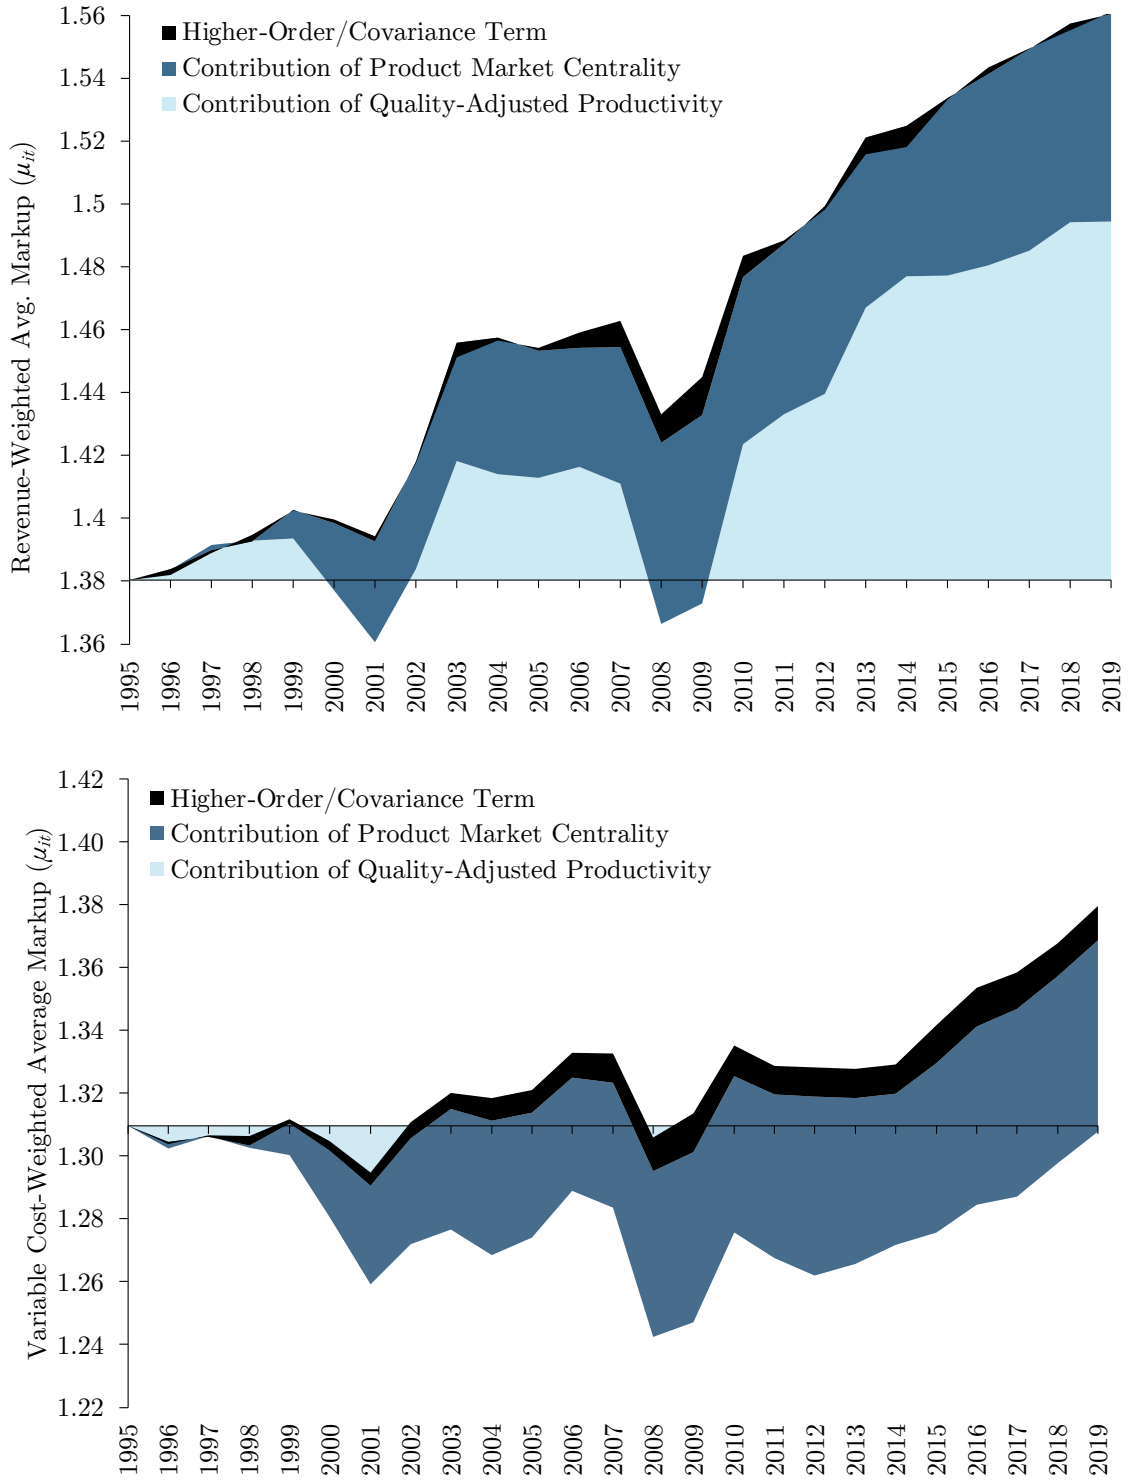


FIGURE NOTES: The figure decomposes the growth rate of the average markup (μ_i) using the formula in equation (4.2). The upper panel displays the decomposition of the revenue-weighted average markup, while the lower panel displays the decomposition for the variable cost-weighted average markup.

decomposition for the change in the revenue-weighted average markup:

$$\Delta \mathbb{E}(\mu_{it}) \approx \underbrace{\frac{1}{2} [\mathbb{E}(\omega_{it}) \cdot \mathbb{E}(\chi_{it}^{-1})] \Delta \log \mathbb{E}(\omega_{it})}_{\text{Contribution of Quality-Adjusted Productivity}} + \underbrace{\frac{1}{2} [\mathbb{E}(\omega_{it}) \cdot \mathbb{E}(\chi_{it}^{-1})] \Delta \log \mathbb{E}(\chi_{it}^{-1})}_{\text{Contribution of Centrality}} \quad (4.2)$$

where the discrepancy is determined by a higher-order/covariance term.

Figure 3 illustrates the implementation of this decomposition for the period 1995-2019. The upper panel presents the decomposition for the revenue-weighted average markup, while the lower panel presents the decomposition of the average markup weighted by variable costs. The difference between these two statistics consists in the fact that a reallocation of market shares towards higher markups will affect the first, but not the latter.

Over this period, we can see that the revenue-weighted average markup increased from 1.38 to 1.56. Relative to the minimum theoretical value of one, this value represents an increase of 47%. The decomposition shows that approximately three-fifths of this increase (+0.11) can be attributed to changes in quality-adjusted productivity (ω_i), with the remaining share of the increase (+0.06) being driven by the contribution of product market centrality (χ_i).

When we consider cost-weighted markups instead, the increase in markup is more muted, from 1.31 to 1.38 (+23%). The effect of quality-adjusted productivity is completely absent, while the contribution of product market centrality is broadly unchanged.

The interpretation of these findings is that the fall in product market centrality occurs *within* firms, and reflects a lesser availability of substitutes for the typical product. The increase in (revenue-weighted) quality-adjusted productivity, on the other hand, is driven by a reallocation of revenues towards more productive firms. Because market shares are driven by the difference ($b_i - MC_i$), it must be the case that this reallocation is driven by vertical differentiation: high-productivity firms (which charge higher markups to begin with) experience an increase in product quality (b_i), combined by a more-or-less proportional increase in marginal costs (MC_i). This leaves their ratio (ω_i) unaffected but increases their difference ($b_i - MC_i$), thus triggering a “superstar” effect: a reallocation of market shares towards high-productivity firms. Because demand is linear these firms also charge higher markups: this results in an overall increase in sales-weighted markups.

4.3. Time Trends in Total Surplus and Consumer Surplus

HP’s cosine similarity data is available starting from 1997. By mapping my model to Compustat data year by year, I can produce annual estimates of the welfare metrics previously presented. This allows me to study the welfare implications of the rising concentration of US industries. Most importantly, because my model leverages HP’s time-varying product similarity data, these estimates account for how the product offering of US public firms changed over time. This is another contribution of this study.

In Figure 4, I plot aggregate consumer surplus S (the dark area) and profits Π (the light area) for every year between 1997 and 2017. The combined area represents total surplus W . I also plot, on the right axis (dotted black line), profits as a share of total surplus Π/W . All these statistics refer to the (observed) Cournot equilibrium.

The graph shows that the total surplus produced by US public corporations has nearly doubled between 1997 and 2017 from \$4.6 trillion to \$9.1 trillion between 1997 and 2017. Profits have increased more-than-proportionally with respect to consumer surplus – from about \$2.2 trillion to about \$5 trillion. Consumer surplus increased instead from \$2.2 trillion in 1997 to about \$4 trillion in 2017. As a consequence, the profit share of surplus has increased from about 50% of total surplus to nearly 56%. The consumer appears to

FIGURE 4: TOTAL SURPLUS OF US PUBLIC FIRMS (1997-2017)




FIGURE NOTES: The figure above plots the evolution, between 1997 and 2017, of aggregate (economic) profits $\Pi(\mathbf{q})$, aggregate consumer surplus $S(\mathbf{q})$ and total surplus $W(\mathbf{q})$, as defined in the model from Section 2. Profits as a percentage of total surplus (Π/W , black dotted line) are shown on the right axis. These statistics are estimated over the universe of the US publicly-listed corporations. These surplus measures are gross of fixed costs. Appendix E replicates this graph using surplus measures that are net of fixed costs.

capture a decreasing share of the surplus generated by public companies.

In Figure 5, I plot, over the same period, the percentage gain in total surplus from moving from the competitive equilibrium \mathbf{q}^Φ to the first best \mathbf{q}^W . This is the deadweight loss from oligopolistic behavior, and is plotted as the darker line. Its trend that mimic that of profit share of surplus: it increased from 8.5% (in 1997) to the current level of 11% (in 2017). This suggests that the impact of oligopoly on surplus creation has increased over time.

To investigate the impact of fixed costs on these results, I plot, in the same figure (light line), the percentage difference in total surplus between the Cournot equilibrium and the *Second Best*: it has increased from 6.4 percentage points (in 1997) to 9.3 percentage points (in 2017). In other words, when fixed costs are taken into account, the welfare loss caused by oligopoly starts from a lower level (as is to be expected, mathematically) but increases more sharply over time (+45% over the period)

Overall, my findings are consistent with the interpretation that U.S. public firms have more oligopoly power than they had in 1997, and that this increase in oligopoly power has had a significant impact on both allocative efficiency and consumer welfare.

4.4. The Role of Entry Costs

The increasing concentration of US industries is reflected in the stark decline of the number of public companies: their number has dropped from about 7,500 in 1997 to about 3,500 in 2017 (Kahle and Stulz, 2017). This is largely the result of a collapse in the rate of new Initial Public Offerings (IPOs). This lower

FIGURE 5: DEADWEIGHT LOSS FROM OLIGOPOLY (1997-2017)

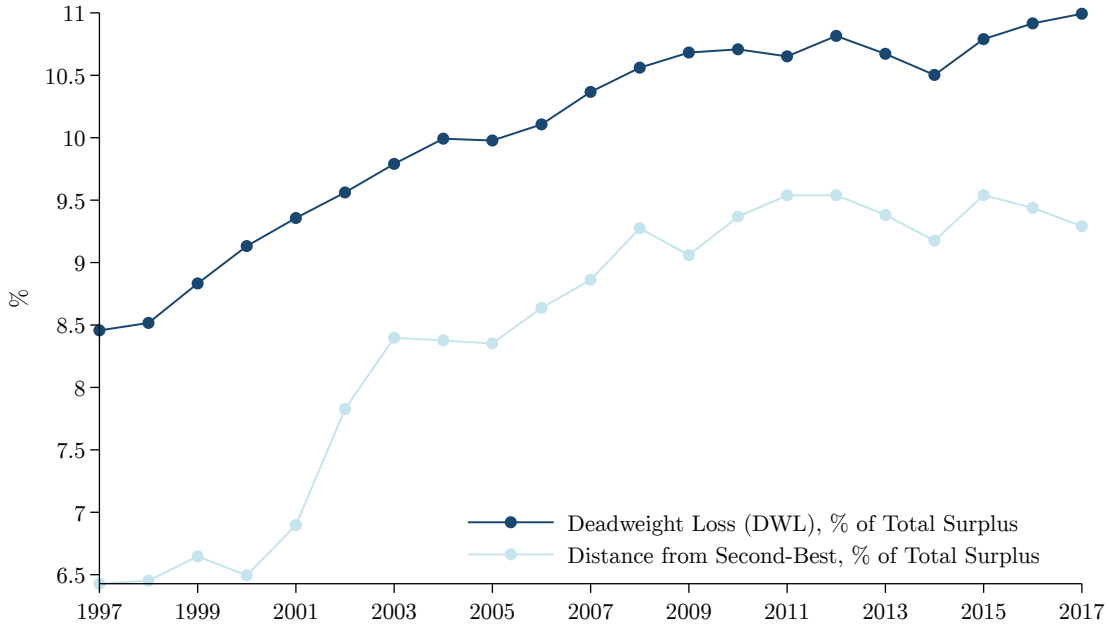


FIGURE NOTES: The following figure plots the estimated deadweight loss (DWL) from oligopoly, between 1997 and 2017. The darker line is the traditionally-defined DWL - the % difference in total surplus between the Cournot equilibrium and the First-Best scenario, while the lighter line is the % difference between the Cournot equilibrium and the Second-Best scenario as defined in Section 2. These surplus measures are gross of fixed costs. Appendix E replicates this graph using surplus measures that are net of fixed costs.

rate of entry among public firms, as well as the higher profits and deadweight losses, could be the result of higher entry costs. Obviously, this is hard to know for sure, since entry costs are not directly observable.

Next, I use my model to obtain an estimate of how costs of entry have changed over time. The intuition behind this exercise is to ask how high would the cost of entry have to be, in order to deter a marginal startup from entering the market. Thus far, I have kept the number of granular firms fixed. Now, I consider the entry problem of a marginal startup indexed by $i = 0$ that has quality-adjusted productivity ($b_0 - c_0$) and fixed costs (f_0) that are equal to the median in the population of incumbents.

We can use the model to compute firm 0's economic profits conditional on entry (π_0). Let us then assume that firm 0 faces an entry cost of e_0 , and that it enters the market if $\pi_0 - f_0 \geq e_0$. Given these assumptions, the net profits conditional on entry ($\pi_0 - f_0$) provide a lower bound for the entry cost e_0 .

The challenge in computing π_0 is that, in my model, firms exist in a space of product characteristics. Therefore, to estimate π_0 , I also need to measure 0's cosine similarity vis-à-vis every other firm in the dataset ($\mathbf{a}'_0 \mathbf{a}_i$). To measure $\mathbf{a}'_0 \mathbf{a}_i$ (for all i), I obtain another set of cosine similarities from Hoberg, Phillips and Prabhala (2014). These cosine similarities are computed, for Compustat firms, against a generic venture capital-backed startup. To construct them, the authors use startup product descriptions from the VenturExpert database. This data is available for 1997-2008.¹⁸

Figure 6 presents the resulting estimate of the implied entry cost, normalized at the level of 1997. As of 2008, it has increased by approximately 60%. What this exercise tells us is that incentives to enter have

¹⁸I thank the authors for retrieving and sharing this data, which is not part of the public HP repository.

FIGURE 6: IMPLIED COST OF ENTRY FOR A VC-BACKED STARTUP




FIGURE NOTES: the figure above plots the implied cost of entry for a VC-backed startup with median quality-adjusted productivity ($b_i - c_i$) and fixed cost (f_i), as implied by the model in Section 2. The similarity scores for the marginal entrant are computed by Hoberg, Phillips and Prabhala (2014), using VenturExpert product descriptions.

increased; hence, from a marginal analysis perspective, the model suggests that constraints or disincentives to enter should also have increased over this period. This occurs at approximately the same time as the number of new IPOs collapses.

This exercise is, obviously, not informative about the nature of these unobserved entry costs. In the next Subsection, I further investigate the collapse of the rate of entry among public firms.

4.5. Startup Takeovers and the Secular Decline of IPOs

One interesting aspect of the decline of IPOs is that it appears to be unrelated to the decline of the startup rate that has been measured in the broader economy (Decker, Haltiwanger, Jarmin and Miranda, 2014). Far from declining, the number of startups that are backed by Venture Capital (VC), which make up the majority of startups that eventually become public companies, has boomed over this period.

Figure 7 plots the number of Venture Capital exits in the United States by year and type, for the period 1985-2017¹⁹. In the diagram, I separate IPOs from acquisitions. We can see from this figure that, at the beginning of the 1990s, the vast majority of VC exits were IPOs. However, since the mid-90s, there has been a dramatic shift toward acquisitions. One implication of this fact is that the decline of IPOs was not driven by a dearth of startups. Instead, the reason why IPOs have decreased is that they have been largely replaced by acquisitions.

The next question I ask is to what extent can the shift from IPOs to acquisitions account for the rising

¹⁹This data is sourced from the National Venture Capital Association (NVCA).

FIGURE 7: VENTURE CAPITAL STARTUP EXITS BY TYPE




FIGURE NOTES: the figure above plots the number of successful venture capital exits in the United States by year and type (Initial Public Offering v/s Acquisition). The data is sourced from the National Venture Capital Association (NVCA).

profit share of surplus, and the increasing deadweight loss?

To investigate this question, I construct a counterfactual in which I add granular firms to the model, with the objective of simulating the ratio of IPOs to acquisitions remaining constant after 1997. For each firm i entering Compustat after 1997, I add $(N_i - 1)$ firms to the model. These firms are “similar” to i in the sense that they share the same value of $(b_i - c_i)$ as well as the same coordinates in the space of common characteristics (\mathbf{a}_i); they also exit the sample in whichever year firm i exits the dataset. Yet, they are not perfect substitutes to i , due to the presence of idiosyncratic characteristics.

N_i is determined so that, in this counterfactual, the ratio of IPOs to acquisitions remains constant after 1997. Specifically, let IR_t be ratio of IPOs to total VC exits at time t . I set N_i as follows:

$$N_i = \begin{cases} \frac{IR_{1997}}{IR_t} & \text{if } i \text{ went public at time } t \\ 1 & \text{otherwise} \end{cases} \quad (4.3)$$

We can see the result of this counterfactual analysis in Figure 8: it shows the percent difference in consumer surplus between the Cournot equilibrium and the *Perfect Competition* counterfactual, under two alternative scenarios. The lighter line shows the baseline case: consistent with the findings of Subsection 4.3, the percentage gap in consumer surplus increases from 42% to 51%, reflecting the larger deadweight loss as well as the larger share of total surplus accruing to producers.

The darker line shows the counterfactual where the ratio of IPOs to acquisitions stays constant after 1997. Under this alternative scenario, the increase in the consumer surplus “gap” is significantly less pronounced, leveling at 43.9% in 2017. This reflects a more muted increase in the deadweight loss, as well as a slight decrease in the profit share of surplus.

FIGURE 8: CONSUMER SURPLUS, % DIFFERENCE FROM PERFECT COMPETITION

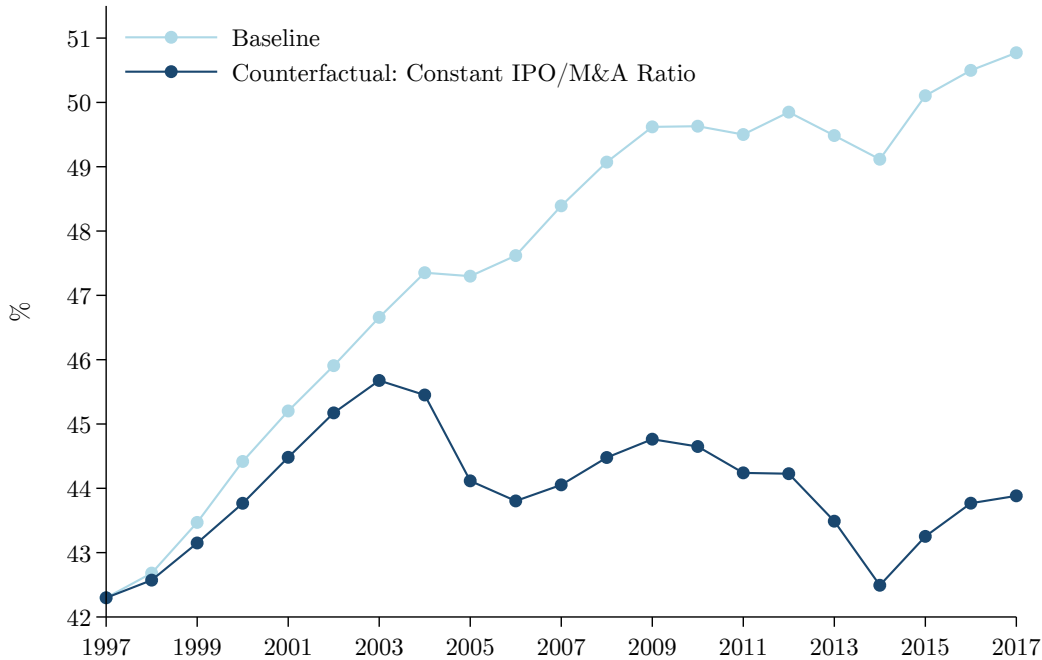


FIGURE NOTES: The following figure plots the percentage difference in consumer surplus between oligopolistic competition and perfect competition, between 1997 and 2017. The observed equilibrium value (light line) is plotted against a counterfactual scenario (darker line) in which the ratio of IPOs to startup acquisitions remains constant after 1997.

This counterfactual analysis suggests that the shift from IPOs to takeovers may have contributed at least to some degree to the measured increase in oligopoly power. There are some caveats, of course. First, this is just one of many factors that may have contributed – and it only offers a *proximate* explanation for the increasing oligopoly power: my model is silent on the reasons why this shift from IPOs to acquisitions has occurred in the first place²⁰. It is also important to note that this counterfactual analysis should not be interpreted as causal evidence.

That being said, these results do seem to align with a recent empirical literature that has focused on the anti-competitive effects of startup acquisitions. Wollmann (2019), for example, argues that startup acquisitions may have been used by large corporations to engage in what he calls *stealth consolidation*. The idea is that the majority of startup acquisitions fall under the size threshold for mandatory merger review: as a consequence, startup acquisitions rarely undergo merger review. Large companies may then be able to use startup acquisitions to engage in monopolization with little risk of attracting antitrust scrutiny.

In another study, Cunningham, Ederer and Ma (2018) provide evidence from the pharmaceutical industry of so-called *killer acquisitions*: specifically, they show that a significant share of startup acquisitions by drug-makers results in the arrested development of startups’ drugs that might compete with the acquirers’ own.

²⁰See Bowen, Frésard and Hoberg (2018) and Ewens and Farre-Mensa (2020) for a discussion of the potential causes.

5. Robustness and Extensions

In this section, I investigate the robustness of my empirical analyses to a variety of assumptions; I propose extensions to the model; and I discuss the limitations of this paper and potential follow-up work.

5.1. Private and Foreign Firms, Endogenous Entry

I verify that my baseline results are not sensitive to the inclusion of private and foreign firms. Based on the aggregation result derived in Subsection 2.6, I am able to include non-Compustat companies in the model by adding a representative firm that acts competitively and whose size reflects the endogenous entry and exit of atomistic players. This representative firm is used to capture the economic activity of private and foreign firms.

In order to implement this version of the model, I need to assume a cosine similarity between this representative firm, which I label $n + 1$, and every other firm $i = 1, 2, \dots, n$. I assume that the cosine similarity between i and $n + 1$ is simply equal to the average cosine similarity between i and every other firm – formally:

$$\mathbf{a}'_i \mathbf{a}_{n+1} = \frac{1}{n-1} \sum_{j \neq i} \mathbf{a}'_i \mathbf{a}_j \quad (5.1)$$

My empirical results only change slightly as a consequence of this modification: corporate profits, as a percentage of the surplus produced by public firms, increase from 49.5% in 1997 to 55% in 2017. The deadweight loss increases is virtually unchanged.

To ensure that the approach of modeling non-Compustat firms as atomistic is valid, I investigate whether Compustat becomes a much larger or smaller share of US GDP over time. I compute an estimate of the value added by Compustat companies, and investigate how it changes over time as a percentage of corporate business GDP (obtained from BLS NIPA tables). Reassuringly, I find that this percentage does not trend (either positively or negatively) over the period considered: it is 46.8 in 1997 and 47.4 in 2017, with a standard deviation of 5.9 percentage points over the period.

5.2. Fixed Costs and Capital

Thus far, I have defined aggregate profits (Π) and total surplus (W) *gross* of fixed costs (F). Next, I study how Figures 4 and 5 change if we subtract F from Π and from W . In other words, I want to investigate whether the higher economic profits are somehow justified by higher fixed costs.

In Appendix E, I reproduce both these figures, after redefining Π and W to be computed net of fixed costs. By comparing Figures 4 and 5 with Figures 14 and 15, we can see that my core empirical results are largely unaffected by how fixed costs are incorporated in the analysis. The most significant change in the findings is that the profit share of surplus increases much more dramatically when fixed costs are subtracted: from 11% in 1997 to 22% in 2017.

With regard to the estimation of fixed costs, there has been some debate in the literature about how Selling, General & Administrative (SGA) costs should be treated from an accounting standpoint. This item, as presented in Compustat, includes miscellaneous costs that are not directly linked to production (see Traina, 2018). It also includes R&D expenditures.

While it is generally understood that these are not variable costs, it is also not entirely clear that this cost item is simply overhead. Eisfeldt and Papanikolaou (2013) have argued that SGA partly embeds investments in *intangible capital*, and therefore should not be treated as overhead but capitalized.

Based on this argument, Peters and Taylor (2017) have developed a measure of intangible capital for Compustat: they treat R&D expenditures, plus 30% of the remaining portion of SGA, as investment in intangible capital. They then computed the firm-level intangible capital stock by applying a perpetual inventory model. If we choose to capitalize, rather than expense, these putative investments in intangible capital, we then obtain the following alternative measure of fixed costs:

$$f_i = (\text{SGA}_i - \text{R\&D}_i) \times 0.7 + (\text{PPE}_i + \text{Intan. Capital}_i) \times \text{User CoC} \quad (5.2)$$

Changing the definition of fixed costs (by definition) does not affect my measurements in Figure 4, nor the deadweight loss from Figure 5, because fixed costs do not enter these measures. It does, however, affect the distance from the second-best (which is shown in Figure 5) as well as the analysis from Appendix E, which I have discussed above. In the same Appendix, I show that my results are largely unaffected if intangible investments are capitalized, rather than expensed.

5.3. Labor Supply Elasticity

Next, I investigate how my empirical results change if I assume a completely-inelastic labor supply function.

By definition, profit as a share of total surplus would be unchanged. The deadweight loss would instead become the total surplus difference between the Cournot equilibrium and the *Resource-Efficient* counterfactual, which we previously defined in Subsection 2.5. As can be seen in Table 1, this welfare difference is smaller than the deadweight loss. Intuitively, this is because the labor supply (by definition) cannot respond to the removal of the oligopolistic distortions.

I compute this alternative measure of the deadweight loss (the percentage difference in total surplus between Cournot and *Resource Efficient*) over the period 1997-2017. I find that my core empirical results carry through: the level of this “alternative” deadweight loss is 5.2% in 1997, and it increases to 7.9% by 2017. In other words, the level of the deadweight loss is lower if we assume a fixed labor supply (as should be expected), but it increases more sharply (by half) over the 20-year period.

5.4. Multi-product Firms (Diversification vs. Differentiation)

Like most of the macroeconomics literature, this paper does not use product-level data on characteristics. Hence, while the model does reasonably well in capturing measured variation in markups across firms (see Appendix D), it cannot speak to how the market power of a firm varies across individual product markets (say, how Apple’s market power in the smartphone industry may be higher or lower than in the personal computer market). Nevertheless, given the extensive presence in Compustat of multi-product firms, it is worth investigating what additional assumptions are required for my model to apply to multi-product firms.

Suppose that there are still n firms and k characteristics, but now the n firms produce a total of $m \geq n$ products. The same product might be produced by multiple firms and the same firm may produce more than one product. The vector of units produced for each good is now the m -dimensional vector \mathbf{y} . Similarly to matrix \mathbf{A} in Section 2, matrix \mathbf{A}_1 transforms units of products into units of characteristics:

$$\mathbf{x} = \mathbf{A}_1 \mathbf{y} \quad (5.3)$$

Because firms are diversified, each firm now produces a basket of goods: instead of representing the number of units produced of each product, the vector \mathbf{q} now represents the number of baskets produced by each firm.

The matrix \mathbf{A}_2 transforms the vector of quantity indices into units of products supplied:

$$\mathbf{y} = \mathbf{A}_2 \mathbf{q} \quad (5.4)$$

Now I put together the previous two equations. Letting $\mathbf{A} = \mathbf{A}_1 \mathbf{A}_2$, I have

$$\mathbf{x} = \mathbf{A}_1 \mathbf{y} = \mathbf{A}_1 \mathbf{A}_2 \mathbf{q} = \mathbf{A} \mathbf{q} \quad (5.5)$$

The identity above demonstrates how the linear-hedonic structure makes the model easily generalizable to multi-product firms. The intuition is that, if the output of a firm i is not a single product, but rather a basket of products, one can project the quantity index q_i on the characteristics space either in two steps (by projecting it first on goods and then on characteristics), or in one single step (using the composite projection matrix \mathbf{A}).

The limitation of this multi-product interpretation of the model is that, while firms can change their supply q_i , the vector \mathbf{a}_i must stay fixed. What this means is that while firms may produce more than one product and can scale up or down the quantity of baskets produced, they must keep producing the products in constant quantity ratios.

A different way of saying this is that the limitation of firm-level data is that it does not allow to study the reallocation of resources *within* firms, but only *across* firms. This implies that my estimates of the deadweight loss may be conservative: if firms have different degree of market power in different markets, this will generate *within-firm* variation in markups. My model does not capture the additional welfare gains that would be realized if the social planner could remove *within-firm* dispersion in markups.

It is important to emphasize that this limitation is not specific to this paper, but it is endemic in the literature. While my linear demand specification does not fully address it, I claim that the GHL demand system handles multi-product firms better than CES preferences.

Additionally, suppose that we had product or plant-level data (including similarity scores for individual products), so that we may relax the assumption that \mathbf{A}_2 is fixed (firms can change output ratios among products). The model would then need to account for the fact that different business units within the same firm coordinate their supply decisions, similar to what merging/colluding firms do in the counterfactual from equation (2.45). Hence, one argument for using firm-level data is that it is a tractable way of modeling coordination of plants within the same firm.

Finally, it useful to ask how does diversification affect the measured intensity of competition across firms. If, for example, firms became more diversified, would the model erroneously interpret the resulting change in cosine similarity as an increase in market power? The answer to this question is *it depends* – i.e. the effect of diversification of measured competition is ambiguous. It is possible to construct examples where diversification leads to higher, lower, or unchanged cosine similarity. In Appendix I, I construct three such examples to illustrate this argument.

5.5. Complements

Because the matrix Σ is non-negative by construction, the marginal utility from one unit of product j is always non-increasing in q_i - formally:

$$\frac{\partial^2 S}{\partial q_i \partial q_j} = -\sigma_{ij} \leq 0 \quad \forall i \neq j \quad (5.6)$$

In light of equation (5.6), it is tempting to jump to the conclusion that all products are by construction

substitutes and that no pair of products are complements. That conclusion is, however, incorrect.

To understand why, we need to recall the textbook definition of substitution and complementarity. Two goods (i, j) are²¹:

$$\text{Complements if } \frac{\partial q_i}{\partial p_j} < 0 \quad \text{Substitutes if } \frac{\partial q_i}{\partial p_j} > 0 \quad (5.7)$$

We intuitively expect this derivative to have the opposite sign of that in equation (5.6). In the case of CES, this intuition is correct. In the case of my model, however, this intuition fails. This is a consequence of the fact that the cross-price demand elasticity depends on the inverted matrix $(\mathbf{I} + \mathbf{\Sigma})^{-1}$, not on $\mathbf{\Sigma}$ itself. If $\mathbf{\Sigma}$ is not symmetric (here it is not) the off-diagonal elements of $-(\mathbf{I} + \mathbf{\Sigma})^{-1}$ will generally include positive as well as negative elements.²² This implies that, in the empirical implementation of the model, many producer pairs are strategic complements.

For example, if we compute the vector of cross-price derivatives for car manufacturer General Motors in 2017, we will find that it includes several negative elements (i.e. complements), mostly corresponding to energy and consumer finance companies. This makes sense: intuitively, we expect higher oil prices, loan rates or insurance premia to adversely affect the residual demand for cars.

Hence, despite the property of the model described by equation (5.6), my model does produce strategic complementarity. Indeed, I argue that one of the strengths of the network Cournot model is its ability to produce a rich competitive environment that includes complement goods.

5.6. Limitations and Future Work

This model (like every other model) has limitations and leaves out certain aspects of market power that might be relevant to the current debate on antitrust policy.

For example, one important assumption that I make in my model (in order to preserve tractability) is that all firms are final goods firms. In other words, input-output linkages between individual firms are not part of the model. This might lead to underestimating the welfare costs of oligopoly, if input-output linkages result in double marginalization. One interesting extension of the model would be to modify the firm's production function to allow firms to use other firm's output as inputs. In order to take this extended model to the data, granular input-output data on firm-to-firm relationships would be required.

Another important restriction of my model is that it treats the firms' position in the product characteristics space as fixed. The assumption that product characteristics are exogenous is standard in the demand estimation literature (Berry et al., 1995; Nevo, 2001); however, it makes policy simulations less robust in the long-run. This is because, given enough time, firms may be able to endogenously change their product portfolios. Some recent work in the IO literature (see Fan, 2013; Wollmann, 2018) has considered endogenous product characteristics. If product characteristics data can be obtained (HP only provide cosine similarities), another interesting direction for future research is to generalize my model by endogenizing the firms' position in the characteristics space.

Finally, another force that is left out of this paper is labor market power. Concentration may lead not only to oligopoly power in output markets: it may also lead to oligopsony in input markets. Because the model that I presented does not speak to this channel, my estimates of income distribution and aggregate

²¹Usually, these derivatives refer to the *Hicksian* demand, to ensure that $\text{sign}(\partial q_i / \partial p_j) = \text{sign}(\partial q_j / \partial p_i)$ (see Kreps, 2012, section 11.6). This distinction can be ignored for my model, since the sign of this derivative is the same for the Marshallian and the Hicksian demand.

²²It is fairly easy to come up with examples. Consider for instance three goods (1,2,3) and three common characteristics (A,B,C). If goods 1 and 3 load entirely on characteristic A and C respectively, and good 2 loads equally on all three characteristics, then it can be verified that goods 1 and 3 are strategic complements.

efficiency do not include the effects of the labor market power, which might also have increased over this period.

With additional labor market data, it is definitely possible to apply the methods developed in this paper to a labor market setting. Cosine similarities can easily be constructed for textual descriptions of job vacancies: if we see labor input as a differentiated good, we can then develop and estimate a labor market version of the model presented in Section 2. In Appendix J, I propose two ways of reframing my model to study labor market power: the first considers *workers'* monopoly power (workers with unique characteristics command higher wages); the latter looks at oligopsony (firms that utilize unique inputs are able to charge higher markups).

6. Conclusions

In this study, I have presented a new general equilibrium model of oligopolistic competition with hedonic demand and differentiated products, with the objective of measuring the welfare consequences of rising oligopoly in the United States from 1997 to 2017. To estimate my model, I used a data set (recently developed by Hoberg and Phillips, 2016) of product cosine similarities that covers all public firms in the United States on a yearly basis. Through the lens of my model, these similarity scores are used to retrieve the cross-price elasticity of demand for every pair of publicly traded firms.

My measurements suggest that oligopoly has a considerable and growing effect on aggregate welfare. In particular, I estimate that, if all publicly traded firms were to behave as atomistic competitors, the total surplus produced by this set of companies would increase by 11 percentage points. Consumer welfare would increase even more dramatically—it would more than double—as surplus would be largely reallocated from producers to consumers. I find that a large share of the deadweight loss caused by oligopoly (7.7 percentage points) can be attributed to resource misallocation—that is, a significant share of the deadweight losses could theoretically be recovered by a benevolent social planner, even if we assumed labor to be inelastically supplied.

I also find evidence that the deadweight losses increases significantly in the presence of collusion: consolidating firm ownership in the hands of one owner, who induces firms to coordinate, would depress aggregate surplus by an additional 10 percentage points. Consumer surplus would drop even more, with a projected decrease of about 38 percentage points. Overall, my analysis of firm-level data suggests that there is evidence of sizable welfare distortions due to oligopoly power.

By mapping my model to firm-level data for every year between 1997 and 2017, I find that, while both the profits earned by U.S. public corporations and the corresponding consumer surplus have increased over this period, profits have increased at a significantly faster pace: consequently, the share of surplus appropriated by firms in the form of profits has increased substantially (from 50% to 55.5%). Consistent with this finding, I estimate that the welfare costs of oligopoly, computed as the percentage increase in surplus that is obtained by moving to the competitive outcome, have increased (from 8.5% to 11%). Overall, my estimates are consistent with the hypothesis that the observed secular trends in markups and concentration have resulted in an increased welfare loss, particularly at the expense of the consumer.

The model also disentangle the role of two factors that have contributed to secular rise of markups (previously documented by De Loecker, Eeckhout and Unger, 2020): productivity (adjusted for quality) and product market centrality (which captures the availability of substitute products). The first drives the reallocation of market shares towards high-markup firms, while the latter drives the within-firm growth component. This finding provides a proximate explanation for the rise of markups, and imposes significant constraints to theories that seek to provide a deep explanation: a credible theoretical account for the secular

rise of markups should accommodate both of these forces.

The model allows me to compute a number of novel counterfactuals that are relevant to competition policy, and to shed additional light on these ongoing trends. I have shown that a potential contributor to the increased oligopoly power might lie in the secular decline in IPOs and the surge in takeovers of VC-backed startups. Through the lens of my model, this shift can account for a large share of the increase in the deadweight loss and in the producer share of surplus.

This paper contributes—both methodologically and empirically—to a growing literature in macroeconomics and finance that is devoted to incorporating heterogeneity, imperfect competition and Industrial Organization methods in general equilibrium models. In particular, it shows that combining firm financials with measures of similarity based on natural-language processing of regulatory filings offers a promising avenue to model product differentiation and imperfect substitutability at the macroeconomic level: it affords the opportunity to impose a less arbitrary structure on the degree of substitution across sectors and firms.

In addition to the theory contribution, this paper provides measurements that align with a growing body of empirical work on rising market power (De Loecker, Eeckhout and Unger, 2020) and the anti-competitive effects of startup acquisitions (Cunningham et al., 2018; Wollmann, 2019). One potential policy implication of my findings is that startup takeovers may have significant anti-competitive effects at the aggregate level – even if one takes the view that individual acquisitions have negligible impact. Because competition authorities are much less likely to challenge these acquisitions than they are to oppose a merger between large incumbents, these findings strengthen the case for increasing the antitrust oversight of these transactions.

References

- ACEMOGLU, D., V. M. CARVALHO, A. OZDAGLAR, AND A. TAHBAZ-SALEHI (2012): “The network origins of aggregate fluctuations,” *Econometrica*, 80, 1977–2016.
- ACEMOGLU, D., A. OZDAGLAR, AND A. TAHBAZ-SALEHI (2017): “Microeconomic origins of macroeconomic tail risks,” *American Economic Review*, 107, 54–108.
- ASPLUND, M., AND V. NOCKE (2006): “Firm turnover in imperfectly competitive markets,” *The Review of Economic Studies*, 73, 295–327.
- AUTOR, D., D. DORN, L. F. KATZ, C. PATTERSON, AND J. VAN REENEN (2020): “The fall of the labor share and the rise of superstar firms,” *The Quarterly Journal of Economics*, 135, 645–709.
- AXTELL, R. L. (2001): “Zipf Distribution of US Firm Sizes,” *Science*, 1818–1820.
- BAKER, J. B., AND T. F. BRESNAHAN (1985): “The gains from merger or collusion in product-differentiated industries,” *The Journal of Industrial Economics*, 427–444.
- BALLESTER, C., A. CALVÓ-ARMENGOL, AND Y. ZENOU (2006): “Who’s who in networks. Wanted: The key player,” *Econometrica*, 74, 1403–1417.
- BAQAEE, D. R., AND E. FARHI (2020): “Productivity and misallocation in general equilibrium,” *The Quarterly Journal of Economics*, 135, 105–163.
- BERRY, S., J. LEVINSOHN, AND A. PAKES (1995): “Automobile prices in market equilibrium,” *Econometrica: Journal of the Econometric Society*, 841–890.
- BLOOM, N., M. SCHANKERMAN, AND J. VAN REENEN (2013): “Identifying technology spillovers and product market rivalry,” *Econometrica*, 81, 1347–1393.
- BONACICH, P. (1987): “Power and centrality: A family of measures,” *American journal of sociology*, 92, 1170–1182.

- BOWEN, D. E. I., L. FRÉSARD, AND G. HOBERG (2018): “Technological disruptiveness and the evolution of IPOs and sell-outs.”
- CARVALHO, V. M. (2014): “From micro to macro via production networks,” *Journal of Economic Perspectives*, 28, 23–48.
- CARVALHO, V. M., M. NIREI, Y. SAITO, AND A. TAHBAZ-SALEHI (2020): “Supply chain disruptions: Evidence from the great east japan earthquake,” *The Quarterly Journal of Economics*, forthcoming.
- CARVALHO, V. M., AND A. TAHBAZ-SALEHI (2019): “Production networks: A primer,” *Annual Review of Economics*, 11, 635–663.
- CHEN, H., L. COHEN, AND D. LOU (2016): “Industry window dressing,” *The Review of Financial Studies*, 29, 3354–3393.
- CHINTAGUNTA, P. K. (2002): “Investigating category pricing behavior at a retail chain,” *Journal of Marketing Research*, 39, 141–154.
- COVARRUBIAS, M., G. GUTIÉRREZ, AND T. PHILIPPON (2020): “From Good to Bad Concentration? US Industries over the past 30 years,” *NBER Macroeconomics Annual*, 34, 1–46.
- CUNNINGHAM, C., F. EDERER, AND S. MA (2018): “Killer acquisitions,” *Available at SSRN 3241707*.
- DE LOECKER, J., J. EECKHOUT, AND G. UNGER (2020): “The rise of market power and the macroeconomic implications,” *The Quarterly Journal of Economics*.
- DE LOECKER, J., AND P. T. SCOTT (2016): “Estimating market power: Evidence from the US Brewing Industry,” working paper, National Bureau of Economic Research.
- DECKER, R., J. HALTIWANGER, R. JARMIN, AND J. MIRANDA (2014): “The role of entrepreneurship in US job creation and economic dynamism,” *Journal of Economic Perspectives*, 28, 3–24.
- DHINGRA, S., AND J. MORROW (2019): “Monopolistic competition and optimum product diversity under firm heterogeneity,” *Journal of Political Economy*, 127, 196–232.
- EDMOND, C., V. MIDRIGAN, AND D. Y. XU (2018): “How costly are markups?” working paper, National Bureau of Economic Research.
- EINAV, L., AND J. LEVIN (2010): “Empirical industrial organization: A progress report,” *Journal of Economic Perspectives*, 24, 145–62.
- EISFELDT, A. L., AND D. PAPANIKOLAOU (2013): “Organization capital and the cross-section of expected returns,” *The Journal of Finance*, 68, 1365–1406.
- EPPLE, D. (1987): “Hedonic prices and implicit markets: estimating demand and supply functions for differentiated products,” *Journal of political economy*, 95, 59–80.
- EWENS, M., AND J. FARRE-MENSA (2020): “The Deregulation of the Private Equity Markets and the Decline in IPOs,” *The Review of Financial Studies*, 10.1093/rfs/hhaa053, hhaa053.
- FAN, Y. (2013): “Ownership consolidation and product characteristics: A study of the US daily newspaper market,” *American Economic Review*, 103, 1598–1628.
- FOSTER, L., J. HALTIWANGER, AND C. SYVERSON (2008): “Reallocation, firm turnover, and efficiency: Selection on productivity or profitability?” *American Economic Review*, 98, 394–425.
- FRUCHTERMAN, T. M., AND E. M. REINGOLD (1991): “Graph drawing by force-directed placement,” *Software: Practice and experience*, 21, 1129–1164.
- GAO, X., J. R. RITTER, AND Z. ZHU (2013): “Where have all the IPOs gone?” *Journal of Financial and Quantitative Analysis*, 48, 1663–1692.

- GOEREE, M. S. (2008): “Limited information and advertising in the US personal computer industry,” *Econometrica*, 76, 1017–1074.
- GRULLON, G., Y. LARKIN, AND R. MICHAELY (2019): “Are US Industries Becoming More Concentrated?” *Review of Finance*, 23, 697–743.
- HOBERG, G., AND G. PHILLIPS (2016): “Text-based network industries and endogenous product differentiation,” *Journal of Political Economy*, 124, 1423–1465.
- HOBERG, G., G. PHILLIPS, AND N. PRABHALA (2014): “Product market threats, payouts, and financial flexibility,” *The Journal of Finance*, 69, 293–324.
- HOPENHAYN, H. A. (1992): “Entry, exit, and firm dynamics in long run equilibrium,” *Econometrica: Journal of the Econometric Society*, 1127–1150.
- KAHLE, K. M., AND R. M. STULZ (2017): “Is the US public corporation in trouble?” *Journal of Economic Perspectives*, 31, 67–88.
- KATZ, L. (1953): “A new status index derived from sociometric analysis,” *Psychometrika*, 18, 39–43.
- KHAN, L. (2018): “The new Brandeis movement: America’s antimonopoly debate.”
- KREPS, D. M. (2012): *Microeconomic Foundations I: Choice and Competitive Markets*: Princeton University Press.
- LANCASTER, K. J. (1966): “A New Approach to Consumer Theory,” *Journal of Political Economy*, 74, 132–157.
- MELITZ, M. J., AND G. I. OTTAVIANO (2008): “Market size, trade, and productivity,” *The review of economic studies*, 75, 295–316.
- MONDERER, D., AND L. S. SHAPLEY (1996): “Potential games,” *Games and economic behavior*, 14, 124–143.
- NEVO, A. (2001): “Measuring market power in the ready-to-eat cereal industry,” *Econometrica*, 69, 307–342.
- PETERS, R. H., AND L. A. TAYLOR (2017): “Intangible capital and the investment-q relation,” *Journal of Financial Economics*, 123, 251–272.
- ROSEN, S. (1974): “Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition,” *The Journal of Political Economy*, 82, 34–55.
- SYVERSON, C. (2019): “Macroeconomics and Market Power: Context, Implications, and Open Questions,” *Journal of Economic Perspectives*, 33, 23–43.
- TRAINA, J. (2018): “Is aggregate market power increasing? production trends using financial statements,” working paper, Chicago Booth Stigler Center.
- USHCHEV, P., AND Y. ZENOU (2018): “Price competition in product variety networks,” *Games and Economic Behavior*, 110, 226–247.
- WOLLMANN, T. G. (2018): “Trucks without bailouts: Equilibrium product characteristics for commercial vehicles,” *American Economic Review*, 108, 1364–1406.
- (2019): “Stealth consolidation: Evidence from an amendment to the Hart-Scott-Rodino Act,” *American Economic Review: Insights*, 1, 77–94.

Online Appendices

PRODUCT DIFFERENTIATION AND OLIGOPOLY: A NETWORK APPROACH

Bruno Pellegrino, UMD

A. Derivation of the Cournot Potential

In the Network Cournot model, each firm i chooses its own output level to maximize its own profit by taking as given the output of every other firm:

$$q_i^* = \arg \max_{q_i} \pi(q_i; \bar{\mathbf{q}}_{-i}) \quad (\text{A.1})$$

where $\bar{\mathbf{q}}_{-i}$ is the vector of output for every firm except i . The upper bar sign $\bar{\cdot}$ indicates that i treats the quantity supplied by other firms as fixed. The system of first order conditions for this problem is

$$0 = b_i - c_i - (2 + \delta_i) q_i - \alpha \sum_{j \neq i} (\mathbf{a}'_i \mathbf{a}_j) \bar{q}_j \quad (\text{A.2})$$

which can be expressed, in vector form, as:

$$0 = \mathbf{b} - \mathbf{c} - (2\mathbf{I} + \mathbf{\Delta}) \mathbf{q} - \mathbf{\Sigma} \bar{\mathbf{q}} \quad (\text{A.3})$$

This system of reaction functions defines a vector field $\mathbf{q}(\bar{\mathbf{q}})$ which represents the firms' best response as a function of every other firms' strategy. To find the Cournot-Nash Equilibrium, we look for the fixed point \mathbf{q}^* such that $\mathbf{q} = \bar{\mathbf{q}} = \mathbf{q}^*$. Plugging this inside the equation above yields the first order condition that is needed to maximize the potential function $\Phi(\mathbf{q})$:

$$0 = (\mathbf{b} - \mathbf{c}) - (2\mathbf{I} + \mathbf{\Delta} + \mathbf{\Sigma}) \mathbf{q}^* \quad (\text{A.4})$$


which clarifies why the maximizer of the potential function solves the Network Cournot game. The potential $\Phi(\mathbf{q})$ is then obtained as the solution to the following system of partial differential equations

$$\nabla \Phi(\mathbf{q}) = (\mathbf{b} - \mathbf{c}) - (2\mathbf{I} + \mathbf{\Delta} + \mathbf{\Sigma}) \mathbf{q} \quad (\text{A.5})$$

which equates the gradient of the potential function to the linear system of Cournot reaction functions.

The relationship between the potential and the Cournot-Nash equilibrium is represented graphically, for the two-firm case, in Figure 9. The arrows represent the vector field defined by the firms' reaction functions. The potential function is defined to be the scalar-valued function whose gradient coincides with this vector field. A game is a potential game if the vector field defined by the players' reaction functions is a conservative field - that is, if it is the gradient of some scalar function. We call that function the game's potential.

FIGURE 9: GRAPHING THE COURNOT POTENTIAL FOR THE TWO-FIRM CASE



B. Nash-Cournot Equilibrium and Network Centrality

In this appendix, I provide additional details on the relationship between equation (2.27), which describes the equilibrium size of firms in my model, and the measures of network centrality developed by Katz (1953) and Bonacich (1987), which are widely used in the social networks literature.

The game played by the firms from Section 2 is a linear quadratic game played over a weighted network. Ballester, Calvó-Armengol and Zenou (2006, henceforth BCZ) show that players' equilibrium actions and payoffs in this class of games depends on their centrality in the network.

In the game played the firms that populate by model, the adjacency matrix of the network over which the game is played, is given by the matrix $(-\mathbf{\Sigma})$. This matrix appears in the quadratic term of all the welfare functions (profits, total surplus and the Cournot potential).

Before discussing how the linkage extends to my model, I am going to formally define the metric of centrality.

Definition 10 (Katz-Bonacich Centrality). For a weighted network with adjacency matrix \mathbf{G} , we define the vector of centralities \mathbf{f} , with parameters (λ, \mathbf{z}) :

$$\begin{aligned} \mathbf{f}(\mathbf{G}; \lambda, \mathbf{z}) : \quad \mathbf{f} &= \lambda \mathbf{G} \mathbf{f} + \mathbf{z} \\ &= (\mathbf{I} - \lambda \mathbf{G})^{-1} \mathbf{z} \end{aligned} \tag{B.1}$$

The Katz-Bonacich centrality is defined recursively: a node receives a higher centrality score the higher is the centrality of the nodes it is connected to.

The Nash-Bonacich linkage extends to my model. Suppose there are constant returns to scale ($\mathbf{\Delta} = \mathbf{0}$): then, the Cournot-Nash equilibrium allocation of the model presented in Section 2 (equation 2.27) can be easily verified to coincide with the vector of Katz-Bonacich centralities, with parametrization $(\frac{1}{2}, \frac{1}{2}(\mathbf{b} - \mathbf{c}))$:



$$\mathbf{q}^\Phi \equiv \mathbf{f} \left(-\mathbf{\Sigma}; \frac{1}{2}, \frac{1}{2}(\mathbf{b} - \mathbf{c}) \right) \tag{B.2}$$

The peculiarity of the Cournot game played by the firms in my model is that it is played over a *negatively-weighted* network (the adjacency matrix is $-\mathbf{\Sigma}$). The consequence is that the interpretation of \mathbf{q}^Φ as a measure of centrality is reversed: a higher centrality actually reflects a more peripheral position in the network with positive weights $\mathbf{\Sigma}$.

C. Independent Validation of the Hoberg & Phillips Dataset

In this appendix, I validate independently the text-based product similarity measures of Hoberg and Phillips (2016). In the figure below I produce a graph similar to that of Figure 2, while coloring different nodes according to the respective firm's GIC economic sector. The figure shows that there is significant overlap between the macro clusters of the network of product similarity and the broad GIC sectors. To produce this visualization, the dimensionality of data has been reduced from 61,000 to 2; yet, the overlap is nonetheless very clearly visible. The GIC sectors were *not* targeted in producing this graph.

FIGURE 10: VISUALIZATION OF THE PRODUCT SPACE (ALTERNATE COLORING)



D. Data and Calibration

D.1. Mapping to Data and Identification

I present in Table 2 the correspondence between model variables and actual data, with sources. In constructing the dataset, I follow De Loecker, Eeckhout and Unger (2020, henceforth DEU) in excluding firms with negative revenues or costs of goods sold, or negative gross margin (revenues less COGS).

I also follow DEU in the computation of the user cost of capital, which is equal to the federal funds rate (FEDFUNDS from FRED), minus capital goods inflation (PIRIC from FRED), plus a combined depreciation rate and risk premium set at 12%.

D.2. Calibrating the Cost Function

The firms' variable cost function (in linear algebra notation) is:

$$h_i = c_i q_i + \frac{\delta_i}{2} q_i^2 \quad (\text{D.1})$$

We have seen in Section (3) how to recover the intercepts c_i given the slopes δ_i . Calibrating the cost function means calibrating the slopes δ_i . Recall that firms produce using the following quasi-Cobb-Douglas production function:

$$q_i = k_i^\theta \cdot \ell(h_i) \quad (\text{D.2})$$

where k_i is the capital input. Following De Loecker, Eeckhout and Unger (2020, henceforth DEU), who also use Compustat data, I measure capital in Compustat as Property, Plant and Equipment). DEU estimate the production function econometrically. I use DEU's estimates to calibrate $\theta = 0.15$.

Then it must be the case that $\ell(\cdot)$ respects:

$$h_i = \ell^{-1}\left(\frac{q_i}{k_i^\theta}\right) \stackrel{\text{def}}{=} \tilde{c}_i \left(\frac{q_i}{k_i^\theta}\right) + \frac{\tilde{\delta}}{2} \left(\frac{q_i}{k_i^\theta}\right)^2 \quad (\text{D.3})$$

where, by definition:

$$c_i = \frac{\tilde{c}_i}{k_i^\theta} \quad \text{and} \quad \delta_i = \frac{\tilde{\delta}}{k_i^{2\theta}} \quad (\text{D.4})$$

I make the assumption that $\tilde{\delta}$ is constant across firms and over time. Through this assumption, I am reducing the dimensionality of the cost parameter vector to one ($\tilde{\delta}$). We can then rewrite equation (3.10) – the markup of firm i – as a function of $\tilde{\delta}$ and observables:

$$\mu_i = \frac{\left(2k_i^{2\theta} + \tilde{\delta}\right) \cdot p_i q_i}{2k_i^{2\theta} \cdot h_i + \tilde{\delta} \cdot p_i q_i} \quad (\text{D.5})$$

where $(p_i q_i, h_i, k_i)$ are all obtained from Compustat.

TABLE 2: VARIABLE DEFINITIONS AND MAPPING TO COMPUSTAT


Panel A: Observed Variables

Notation	Concept	Measurement
$p_i q_i$	Revenues	Revenues (source: Compustat)
h_i	Total Variable Costs	Costs of Goods Sold (source: Compustat)
f_i	Fixed Costs <i>alternative measure:</i>	$SGA_i + \text{Property Plant \& Equipment}_i \times \text{User Cost of Capital}$ $(SGA_i - R\&D_i) \times 0.7 + (\text{Property Plant \& Equipment}_i + \text{Intangible Capital}_i) \times \text{UCC}$ (source: Compustat, Federal Reserve Economic Data)
$\mathbf{a}'_i \mathbf{a}_j$	Product Cosine Similarity	Word cosine similarity in 10-K product description (source: Hoberg and Phillips, 2016)

Panel B: Identified Variables

Notation	Derived Variable	Computation/Identification
q_i	Output	$= \sqrt{\frac{\pi_i}{1 + \delta_i/2}}$
c_i	Marginal Cost Intercept	$= h_i/q_i - \frac{1}{2}\delta_i q_i$
\mathbf{b}	Demand Intercept	$= (2\mathbf{I} + \mathbf{\Delta} + \mathbf{\Sigma}) \mathbf{q} + \mathbf{c}$

FIGURE 11: REVENUE-WEIGHTED AVERAGE MARKUP



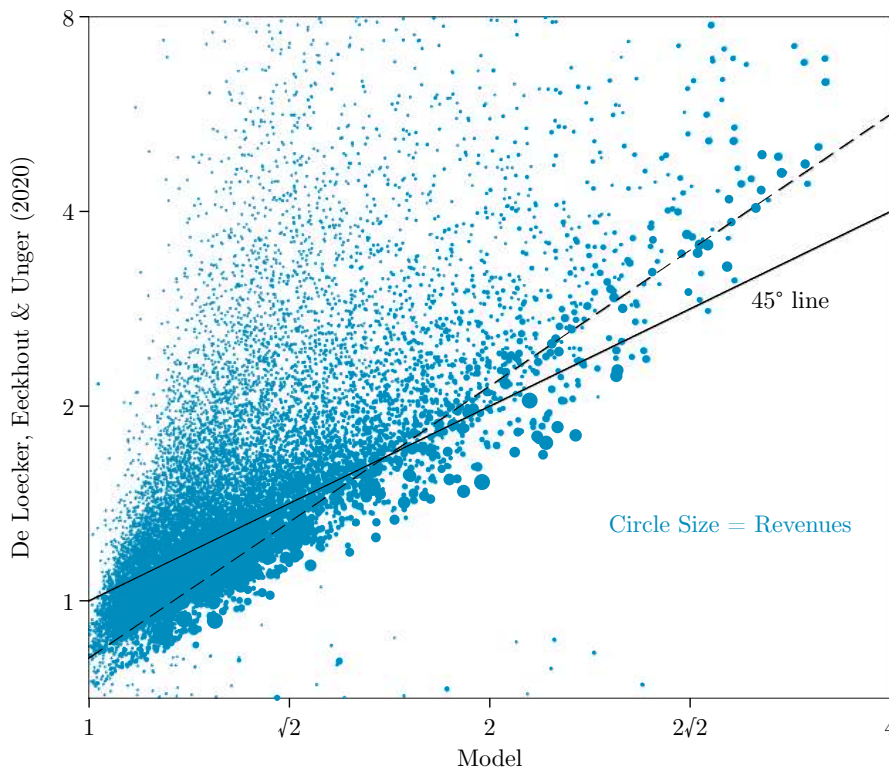
To calibrate $\tilde{\delta}$, I target the markup estimates of De Loecker, Eeckhout and Unger (2020, DEU). Specifically, I calibrate $\tilde{\delta}$ so that the model-inferred average markup matches exactly that estimated by DEU:

$$\tilde{\delta} : \mathbb{E} \left[\frac{(2k_i^{2\theta} + \tilde{\delta}) \cdot p_i q_i}{2k_i^{2\theta} \cdot h_i + \tilde{\delta} \cdot p_i q_i} \right] = \mathbb{E}(\mu_i^{\text{DEU}}) \quad (\text{D.6})$$

Following DEU, I weight markups by revenues in order to compute this average. While I can only estimate my full model for the period 1997-2017, I can actually compute markups as far back as the 1950s, since the identification of markups does not require HP’s product similarity data.

I use GMM to solve the moment equation (D.6). I pool data from the years preceding the rise in markups that was documented by DEU (1957-1982). This implies that only the long-term average markup before 1982 is targeted (not the time trend). The rationale for this approach is that I can then test the assumption that $\tilde{\delta}$ is stable over time by evaluating whether the $\tilde{\delta}$ calibrated before 1982 allows to capture the rise of markups that is observed after 1982. I obtain $\tilde{\delta} = 6.3$.

FIGURE 12: MARKUPS: MODEL VS. DEU ESTIMATES



D.3. Non-Targeted Moments (Model Fit)

For the model to realistically capture the deadweight losses from oligopoly it should generate meaningful time-series and cross-sectional variation in markups. Moreover, we would expect the model-implied markups to correlate with existing econometric estimates (for the Compustat universe, the benchmark would be DEU).

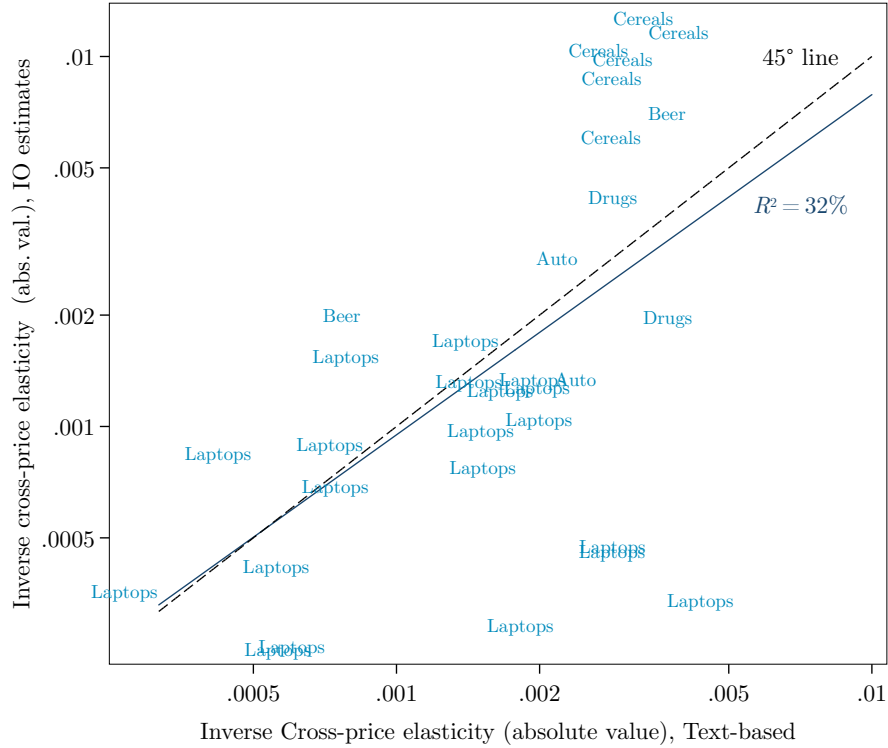
I look at non-targeted moments of the distribution of markups to evaluate the empirical performance of the model. The first non-targeted moment is the time-series correlation between the revenue-weighted average markup (μ_i) from my model and the corresponding estimate by DEU, outside of the calibration window (1983-2017). Figure 11 plots the two series. The navy line is the average markup based on my model. The light blue line is the series computed by DEU.

While only the average markup over 1957-1982 was targeted to calibrate $\tilde{\delta}$, the model-implied average markup tracks closely the time trend of the average markup of DEU after 1982.

Next, I examine how much cross-sectional variation in markups the model is able to generate. One major limitation of current general equilibrium models with market power is that the cross-sectional variation in markups that they can endogenously generate is nowhere close to that estimated by DEU.

In Figure 12, I plot the cross-section of markups obtained from the model against DEU's estimates. Both axes use a log scale. I include data from 1997, 2007 and 2017. As in DEU, I weight observations by revenue (weights are represented by circle sizes). On the one hand, it is fairly evident from the graph that there is more variation in the markups estimated by DEU than in the model-based estimates: the range is wider and the fitted regression slope (the dotted black line) is larger than one. On the other hand, I find that my model-based estimates capture a remarkable 77% of the cross-sectional variation in markups (as measured

FIGURE 13: CALIBRATED ELASTICITIES VS. MICROECONOMETRIC ESTIMATES



Industry	Citation	Journal/Series
Auto	Berry, Levinsohn and Pakes (1995)	Econometrica
Beer	De Loecker and Scott (2016)	NBER Working Papers
Cereals	Nevo (2001)	Econometrica
Drugs	Chintagunta (2002)	Journal of Marketing Research
Laptops	Goeree (2008)	Econometrica

by the R -squared of a regression where the slope is constrained to one²³).

Another important question from the point of view of model fit is how well my model-based estimates of the cross-price elasticity of demand correlate with the corresponding microeconomic estimates that were used to calibrate the parameter α . To this end, notice that equation (3.12) can be rearranged as:

$$\log \left| \frac{\partial \log p_i}{\partial \log q_j} \right| \approx \log \alpha + \log \left(\mathbf{a}'_i \mathbf{a}_j \frac{q_j}{p_i} \right) \quad (\text{D.7})$$

²³If we don't constrain the coefficient to one, the R^2 increases to 83%.

where we use the symbol \approx to denote the fact that measurement error contaminates what should be a linear-in-logs relationship.

By calibrating α , I was only able to choose the intercept of this linear relationship, not the slope nor the correlation. I will use the (untargeted) slope and the goodness-of-fit of equation (D.7) to evaluate the model fit.

In Figure 13, I plot the relationship between the microeconomic estimates (the left side of equation D.7) and the model-based estimates that use the calibrated value of α , together with the 45 degree line (the predicted slope) and a linear regression line. Each observation is a firm pair (i, j) and the labels reflect the industry study from whom the microeconomic estimates are sourced.

Despite the fact that we cannot affect the slope of the relationship by calibrating α , the slope of the linear fit between the two variables is statistically indistinguishable from one, and there is a strong positive correlation between the two series. The R^2 is 32%, which is particularly high if we consider that the microeconomic estimates were obtained by the respective authors using different assumptions about the underlying demand system, as well as different econometric methodologies.

E. Accounting for Fixed Costs and Intangible Capital

FIGURE 14: REPLICATION OF FIGURE 4, USING SURPLUS NET OF FIXED COSTS





FIGURE 15: REPLICATION OF FIGURE 5, USING SURPLUS NET OF FIXED COSTS



In this appendix, I replicate Figures 4 and 5 using a different definition of aggregate profits, which is computed net of fixed costs. All results described in the main body of the paper are strengthened if we detract fixed costs. This suggests that fixed costs are not driving the measured increase in oligopoly power.

FIGURE 16: REPLICATION OF FIGURE 4, USING SURPLUS NET OF FIXED COSTS (ALTERNATE MEASURE)


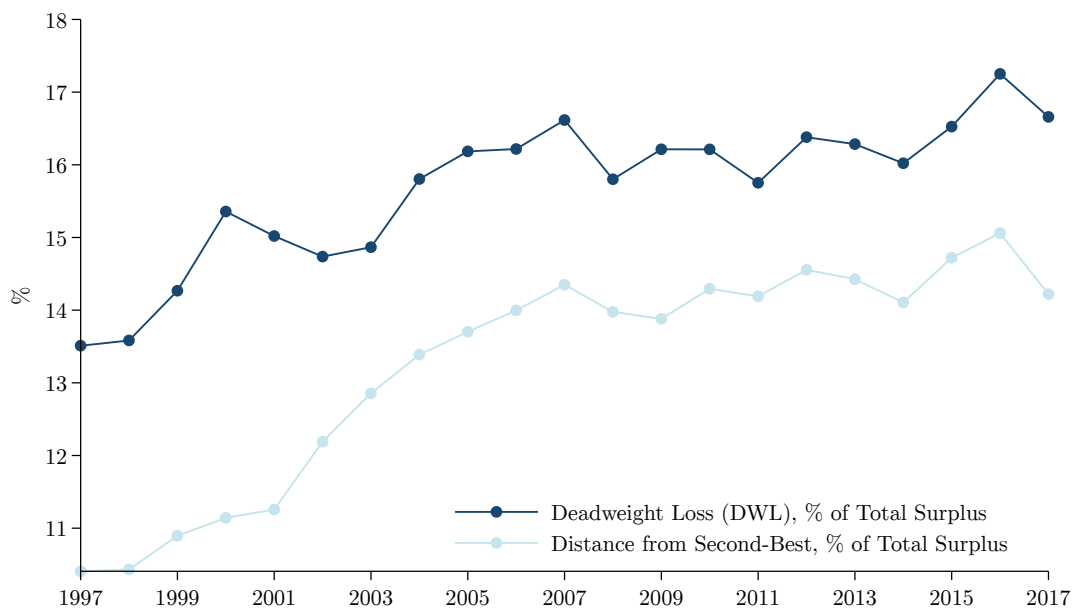


FIGURE 17: REPLICATION OF FIGURE 5, USING SURPLUS NET OF FIXED COSTS (ALTERNATE MEASURE)



Figures 16 and 17 repeat this exercise using the alternate measure of fixed costs described in equation (5.2) (in which intangible investments are capitalized rather than expensed).

F. Bertrand Competition

In this Appendix I derive the equilibrium of the Bertrand counterpart of the Network Cournot game played by the oligopolists in the Model from Section 2. I focus on the special case of flat marginal cost – that is $\Delta = \text{diag}(\mathbf{0})$ (the case with non-flat marginal cost is less tractable). I start by writing the vector of economic profits in terms of the price vector \mathbf{p} :

$$\pi = \text{diag}(\mathbf{p} - \mathbf{c})(\mathbf{I} + \Sigma)^{-1}(\mathbf{b} - \mathbf{p}) \quad (\text{F.1})$$

Now let \mathbb{D} and \mathbb{O} be, respectively, the matrices containing the diagonal and off-diagonal elements of $(\mathbf{I} + \Sigma)^{-1}$ so that:

$$(\mathbf{I} + \Sigma)^{-1} = \mathbb{D} + \mathbb{O} \quad (\text{F.2})$$

Then we can write:

$$\pi = \text{diag}(\mathbf{p} - \mathbf{c})[\mathbb{D}(\mathbf{b} - \mathbf{p}) + \mathbb{O}(\mathbf{b} - \bar{\mathbf{p}})] \quad (\text{F.3})$$

and take the first order condition firm-by-firm by taking the price vector of other firms (denoted by the upper bar as in Appendix A) as given:

$$0 = \mathbb{D}(\mathbf{b} - 2\mathbf{p}) + \mathbb{O}(\mathbf{b} - \bar{\mathbf{p}}) - \mathbb{D}\mathbf{c} \quad (\text{F.4})$$

which we can re-write in terms of \mathbf{q} as:

$$0 = \mathbf{b} - \mathbf{c} - (\mathbb{D}^{-1} + \mathbf{I} + \Sigma)\mathbf{q} \quad (\text{F.5})$$

the corresponding Bertrand potential is

$$\Phi^B = \mathbf{q}'(\mathbf{b} - \mathbf{c}) - \frac{1}{2}\mathbf{q}'(\mathbf{I} + \mathbb{D}^{-1} + \Sigma)\mathbf{q} \quad (\text{F.6})$$

and the Bertrand equilibrium is:

$$\mathbf{q}^B = (\mathbf{I} + \mathbb{D}^{-1} + \Sigma)^{-1}(\mathbf{b} - \mathbf{c}) \quad (\text{F.7})$$

Compare this with the Cournot Equilibrium (flat marginal cost):

$$\mathbf{q}^\Phi = (2\mathbf{I} + \Sigma)^{-1}(\mathbf{b} - \mathbf{c}) \quad (\text{F.8})$$

and perfect competition:

$$\mathbf{q}^W = (\mathbf{I} + \Sigma)^{-1}(\mathbf{b} - \mathbf{c}) \quad (\text{F.9})$$

because \mathbb{D}^{-1} is a diagonal matrix whose diagonal components are between zero and one, we can see how Bertrand is a more “intense” form of competition than Cournot.

G. Further discussion of the GHL demand system

In this Appendix, I discuss more in depth how the GHL demand system introduced in this paper relates to other demand systems already encountered in the literature.

At a very basic level, GHL can be thought of as the hedonic, discrete counterpart of a type of preference aggregator that is already widely used in macroeconomics, trade, and industrial organization (see for example Asplund and Nocke 2006; Foster, Haltiwanger and Syverson 2008; Melitz and Ottaviano 2008):

$$U(\mathbf{q}) = q_0 + \alpha \int_{i \in \mathcal{I}} q_i \, di - \frac{\eta}{2} \left(\int_{i \in \mathcal{I}} q_i \, di \right)^2 - \frac{\gamma}{2} \int_{i \in \mathcal{I}} q_i^2 \, di \quad (\text{G.1})$$

These linear-quadratic aggregators are frequently encountered in models with variable markups. My hedonic demand specification adds to this setup the following three features: 1) granular firms; 2) asymmetry in the degree of substitutability between different firm pairs; 3) the ability to interpret the cross derivatives of the demand system in terms of product similarity.

Hedonic demand is something that I adopt from empirical industrial organization (see for example Berry, Levinsohn and Pakes, 1995). There are two main points of departure between GHL and standard IO demand models: the first is that the typical IO demand system is derived from a discrete choice model at the level of the individual consumer: it is obvious why, in a macroeconomics setting, with multiple products, it would be inappropriate to start from a discrete choice model (the consumer would have to choose *either* cars or laptops). The second difference is the functional form, which is linear instead of logit.

A key motivation for using linear demand as opposed to the widely-used CES aggregator is the ability of linear demand to produce heterogeneous markups. The reason why this is a desirable feature is that homogeneous markups, a feature of CES preferences, imply allocative efficiency (Dhingra and Morrow, 2019). While CES preferences have a number of desirable properties (such as the ability to aggregate the output of firms with Pareto-distributed productivity), they become hard to rationalize within the context of this paper, not just empirically but also theoretically. If we try, for example, to write the isoelastic counterpart to the demand system in equation (2.21)

$$\log \mathbf{p} = \mathbf{b} - (\mathbf{I} + \mathbf{\Sigma}) \log \mathbf{q} \quad (\text{G.2})$$

and derive the utility specification that generates it, we find that it will generally fail to satisfy the conditions for integrability, unless we impose the assumption that $\sigma_{ij} = 0 \, \forall ij$, which would make the model trivial. A more intuitive way to rephrase this mathematical fact is that in order to model the cross-price derivatives of the demand system in terms of product cosine similarities, we *need* a model with variable markups.

H. Additional Details on the ‘‘Constant IPO Ratio’’ Counterfactual

In this appendix, I provide additional details on how I implement the ‘‘constant IPO ratio’’ counterfactual. Without loss of generality, consider ‘‘spawning’’ $(N_1 - 1)$ additional firms that are identical to firm 1, except for the idiosyncratic characteristic. The profit and total surplus functions for each of these firms is:

$$\pi_1 = (b_1 - c_1) q_1 - \left(1 + \frac{\delta_1}{2}\right) q_1^2 - \alpha (Q_1 - q_1) q_1 - \alpha \sum_{j \neq 1} (\mathbf{a}'_1 \mathbf{a}_j) q_1 q_j \quad (\text{H.1})$$

$$w_1 = (b_1 - c_1) q_1 - \frac{1}{2} (1 + \delta_1) q_1^2 - \frac{\alpha}{2} (Q_1 - q_1) q_1 - \frac{\alpha}{2} \sum_{j \neq 1} (\mathbf{a}'_1 \mathbf{a}_j) q_1 q_j \quad (\text{H.2})$$

The first order condition is:

$$0 = b_1 - c_1 - (2 + \delta_1) q_1 - \alpha (Q_1 - q_1) - \alpha \sum_{j \neq 1} (\mathbf{a}'_1 \mathbf{a}_j) q_j \quad (\text{H.3})$$

where Q_1 is the joint output of all the N_1 resulting entities, and q_1 is the output of the individual firm. Because the child companies are all identical, in equilibrium they must produce the same quantity, therefore $q_1 = \frac{1}{N} Q_1$. We can then re-arrange:

$$0 = b_1 - c_1 - \frac{2 + \delta_1 + \alpha (N_1 - 1)}{N_1} Q_1 - \sum_{j \neq 1} \sigma_{1j} q_j \quad (\text{H.4})$$

re-labelling Q_1 as q_1 we obtain this new set of first order conditions

$$0 = \mathbf{b} - \mathbf{c} - \left[\begin{array}{c} \frac{2 + \delta_1 + \alpha (N_1 - 1)}{N} q_1 \\ (2\mathbf{I} + \mathbf{\Delta}_2) \mathbf{q}_2 \end{array} \right] - \mathbf{\Sigma} \mathbf{q} \quad (\text{H.5})$$

where q_1 is no longer the output of the individual company of type 1 but the joint output of all the type-1 companies. The new Cournot equilibrium allocation maximizes the following modified potential:

$$\Phi(\mathbf{q}) = \mathbf{q}' (\mathbf{b} - \mathbf{c}) - \frac{1}{2} \mathbf{q}' \left[\begin{array}{cc} \frac{2 + \delta_1 + \alpha (N_1 - 1)}{N} & 0 \\ 0 & (2\mathbf{I} + \mathbf{\Delta}_2) \end{array} \right] \mathbf{q} - \frac{1}{2} \mathbf{q}' \mathbf{\Sigma} \mathbf{q} \quad (\text{H.6})$$

The maximizer of $\Phi(\mathbf{q})$, which corresponds to the post-breakup equilibrium allocation, is

$$\mathbf{q}^\Phi = \left(\left[\begin{array}{cc} \frac{2 + \delta_1 + \alpha (N_1 - 1)}{N} & 0 \\ 0 & (2\mathbf{I} + \mathbf{\Delta}_2) \end{array} \right] \mathbf{q} + \mathbf{\Sigma} \right)^{-1} (\mathbf{b} - \mathbf{c}) \quad (\text{H.7})$$

The total surplus and aggregate profit function have to be modified accordingly:

$$\Pi(\mathbf{q}) = \mathbf{q}' (\mathbf{b} - \mathbf{c}) - \mathbf{q}' \left[\begin{array}{cc} \frac{1 + \frac{1}{2} \delta_1 + \alpha (N_1 - 1)}{N} & 0 \\ 0 & \mathbf{I} + \frac{1}{2} \mathbf{\Delta}_2 \end{array} \right] \mathbf{q} - \mathbf{q}' \mathbf{\Sigma} \mathbf{q} \quad (\text{H.8})$$

$$W(\mathbf{q}) = \mathbf{q}' (\mathbf{b} - \mathbf{c}) - \frac{1}{2} \mathbf{q}' \left[\begin{array}{cc} \frac{1 + \delta_1 + \alpha (N_1 - 1)}{N} & 0 \\ 0 & (\mathbf{I} + \mathbf{\Delta}_2) \end{array} \right] \mathbf{q} - \frac{1}{2} \mathbf{q}' \mathbf{\Sigma} \mathbf{q} \quad (\text{H.9})$$

I. Differentiation and diversification: three examples

In this Appendix, I provide three simple numerical examples that illustrate how diversification could be associated, in the model, with decreasing, unchanged, or increasing competition. The direction of the effect clearly depends on whether cosine similarity increases, decreases or remains unchanged as a consequence of diversification.

I.1. Example 1: Similarity unaffected by diversification

Let us consider the example of two companies – say Apple and Samsung – that start by producing perfectly-substitutable computers (hence they compete fiercely) and then later diversify into producing mobile phones that are also perfectly-substitutable (this implies that there are no idiosyncratic characteristics, or $\alpha \rightarrow 0$). Quantitatively

$$\mathbf{A}_2 = \begin{array}{cc} & \begin{array}{cc} \text{Apple} & \text{Samsung} \\ \downarrow & \downarrow \end{array} \\ \begin{bmatrix} 1 & 1 \\ 0 & 0 \end{bmatrix} & \begin{array}{l} \leftarrow \text{Computers} \\ \leftarrow \text{Mobile Phones} \end{array} \end{array} \quad (\text{I.1})$$

after diversification occurs, this matrix becomes:

$$\tilde{\mathbf{A}}_2 = \begin{bmatrix} \sqrt{.5} & \sqrt{.5} \\ \sqrt{.5} & \sqrt{.5} \end{bmatrix} \quad (\text{I.2})$$

It can be easily verified that, for any matrix \mathbf{A}_1 such that the columns of $\mathbf{A} = \mathbf{A}_1\mathbf{A}_2$ have unit length (that is, regardless of how substitutable computers and mobile phones are):

$$\tilde{\mathbf{A}}_2'\mathbf{A}'_1\mathbf{A}_1\tilde{\mathbf{A}}_2 = \mathbf{A}'_2\mathbf{A}'_1\mathbf{A}_1\mathbf{A}_2 = \mathbf{A}'\mathbf{A} = \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix} \quad (\text{I.3})$$

by computing the $(\mathbf{I} + \boldsymbol{\Sigma})^{-1}$ as $\alpha \rightarrow 0$ (no common characteristics) we find that the derivatives:

$$\lim_{\alpha \rightarrow 0} \left(-\frac{\partial q_i}{\partial p_i} \right) = \lim_{\alpha \rightarrow 0} \left(\frac{\partial q_i}{\partial p_{j \neq i}} \right) = \infty \quad (\text{I.4})$$

hence, as long as Apple and Samsung diversify symmetrically (as measured by the change from \mathbf{A}_2 to $\tilde{\mathbf{A}}_2$) the model captures the fact that Apple and Samsung behave as perfect substitutes both before and after the change.

I.2. Example 2: Diversification reduces similarity

Next, consider a scenario in which only Samsung diversifies into mobile phones:

$$\tilde{\mathbf{A}}_2 = \begin{bmatrix} 1 & \sqrt{.5} \\ 0 & \sqrt{.5} \end{bmatrix} \quad (\text{I.5})$$

and suppose, for simplicity, that \mathbf{A}_1 is equal to an identity matrix (hence there is no substitution between computers and phones). The derivative of q_{Samsung} with respect to p_{Apple} is 1.41, indicating that Apple and Samsung are no longer perfect substitutes from the point of view of the model.

This is not an accurate representation of the elasticities at the product level. However, we must be mindful that what the model takes as inputs are not product-level revenues and costs, but firm-level revenues and costs. When Apple raises computer prices, only the computer part of Samsung’s product portfolio reacts to that price change.

Clearly, the ideal scenario would be to estimate the model at the product level. Lacking product-level data, I argue that it is appropriate to treat products Apple and Samsung as imperfect substitutes. If we treated Apple and Samsung as perfect substitutes after this diversification has occurred, we would be making the implicit assumption that Apple *computers* are perfect substitutes to Samsung *phones*.

In sum, in the example considered above, treating differentiated firms as imperfect substitutes is less-than-ideal. It is, however, still preferable to the alternative, which is to treat them as perfect substitutes while only being able to observe firm-level data.

I.3. Example 3: Diversification increases similarity

Finally, consider a scenario where, initially, Apple produces “mostly” computers and Samsung produces “mostly” phones. We continue to assume that all phones are perfect substitutes and that the same is true for computers. Also, phones do not interact strategically with computers. The two companies then become more diversified, with equal weight being applied to computers and phones:

$$\mathbf{A}_2 = \begin{bmatrix} \sqrt{8} & \sqrt{.2} \\ \sqrt{.2} & \sqrt{.8} \end{bmatrix} \quad \tilde{\mathbf{A}}_2 = \begin{bmatrix} \sqrt{.5} & \sqrt{.5} \\ \sqrt{.5} & \sqrt{.5} \end{bmatrix} \quad (\text{I.6})$$

In this case, the model initially sees the two companies as imperfect substitutes, although their products are, in each market, perfectly substitutable. After the product loadings are equalized, the cosine similarity between Apple and Samsung increases, and the two companies become eventually perfect substitutes in the model.

J. Labor Market Applications/Extensions

In this Appendix, I explore two labor-market applications of the linear-quadratic framework. The first is a model of skill worker monopoly: workers that have a unique set of skills (that are unlike other workers) are able to command higher wages. The second is a model of monopsony: firms that use inputs or labor markets not used by other firms are able to charge larger unit margins.

J.1. Worker Skill Monopoly

There are k types of workers. Each type is endowed with a unit of labor. There are n firms indexed by i that produce a homogeneous good (whose price is normalized to one) and act as price takers. They produce output q_i using a quadratic production function:

$$q_i = \mathbf{s}'_i \mathbf{b}_i - \frac{1}{2} \mathbf{s}'_i \mathbf{s}_i \quad (\text{J.1})$$

where \mathbf{s}_i is the vector of skill units that depend on the vector of labor units \mathbf{h}_i hired:

$$\mathbf{s}_i = \mathbf{A} \mathbf{h}_i \quad (\text{J.2})$$

As in the model from Section 2, the matrix \mathbf{A} transforms labor units into skill units, and $\|\mathbf{a}_i\| = 1$. Hence

$$q_i = \mathbf{h}'_i \mathbf{A}' \mathbf{b}_i - \frac{1}{2} \mathbf{h}'_i \mathbf{A}' \mathbf{A} \mathbf{h}_i \quad (\text{J.3})$$

The optimality condition for firm i is that the marginal revenue product of each type of labor is equal to the corresponding wage

$$\text{MRPL}_i = \mathbf{A}' \mathbf{b}_i - \mathbf{A}' \mathbf{A} \mathbf{h}_i = \mathbf{w} \quad (\text{J.4})$$

The labor demand for firm i is:

$$\mathbf{h}_i = (\mathbf{A}' \mathbf{A})^{-1} (\mathbf{A}' \mathbf{b}_i - \mathbf{w}) \quad (\text{J.5})$$

Summing across firms and defining $\mathbf{b} = \sum_i \mathbf{A}' \mathbf{b}_i$

$$\mathbf{1} = (\mathbf{A}' \mathbf{A})^{-1} (\mathbf{b} - n \mathbf{w}) \quad (\text{J.6})$$

Hence j 's equilibrium wage is equal to the marginal revenue product of the first hour (b_j) less the average cosine similarity of j to every other worker:

$$w_j = b_j - \frac{1}{n} \sum_{m=1}^k \mathbf{a}'_j \mathbf{a}_m \quad (\text{J.7})$$

J.2. Monopsony

There are k types of workers and n firms indexed by i that act as price-takers in goods markets and face output price vector \mathbf{p} . The firms produce using a Leontief production function, so that the following vector of labor inputs \mathbf{h} is required in order to produce output vector \mathbf{q} :

$$\mathbf{h} = \mathbf{T}\mathbf{q} \quad (\text{J.8})$$

Each dimension of \mathbf{h} is a worker and each dimension of \mathbf{q} is a firm. Once again, we normalize $\|\mathbf{t}_i\| = 1$. Assume that the labor supply function is linear, so that i 's wage is:

$$w_i = \lambda_i h_i \quad (\text{J.9})$$

The total cost of firm i is equal to

$$\text{TC}_i = \sum_j w_j x_{ij} = q_i \sum_j w_j t_{ji} = \lambda_i q_i \sum_j \sum_i q_j t_{ij} t_{ji} \quad (\text{J.10})$$

We can then write the vector of profit as:

$$\boldsymbol{\pi} = \text{diag}(\mathbf{q})\mathbf{p} - \boldsymbol{\Lambda} \cdot \text{diag}(\mathbf{q})\mathbf{T}'\mathbf{T}\mathbf{q} \quad (\text{J.11})$$

where

$$\boldsymbol{\Lambda} = \begin{bmatrix} \lambda_1 & 0 & \cdots & 0 \\ 0 & \lambda_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \lambda_n \end{bmatrix} \quad (\text{J.12})$$

Firms compete on quantity, taking the output of all other firms as given. The potential function is:

$$\Phi(\mathbf{q}) = \mathbf{q}'\mathbf{p} - \boldsymbol{\Lambda} \cdot \mathbf{q}(\mathbf{I} + \mathbf{T}'\mathbf{T})\mathbf{q} \quad (\text{J.13})$$

Taking the first order condition we find the Nash equilibrium. It is once again a measure of centrality - this time in the input/technological space:

$$\mathbf{q}^\Phi = (\mathbf{I} + \boldsymbol{\Lambda}\mathbf{T}'\mathbf{T})^{-1} \mathbf{p} \quad (\text{J.14})$$

Because q_i is proportional to the unit margin, one implication of the equation above is that firms that use labor not used by other firms have monopsony power in those labor markets.

K. Proofs and Derivations

Proof to Proposition 2. To prove the proposition we use the fact that, at the Cournot equilibrium, the following two relationships hold:

$$\pi_i = \left(1 + \frac{\delta_i}{2}\right) q_i^2 \quad \text{and} \quad s_i = \frac{1}{2} \left(q_i^2 + \sum_{j \neq i} \sigma_{ij} q_i q_j \right) \quad (\text{K.1})$$

then:

$$\frac{\pi_i}{s_i} = \frac{\left(1 + \frac{\delta_i}{2}\right) \cdot q_i^2}{\frac{1}{2} \cdot q_i \left(q_i + \sum_{j \neq i} \sigma_{ij} q_j \right)} = (2 + \delta_i) \cdot \frac{q_i}{q_i + \sum_{j \neq i} \sigma_{ij} q_j} \equiv (2 + \delta_i) \cdot \omega_i \quad (\text{K.2})$$

□

Proof to Lemma 1. To simulate a merger (or a collusion), we sum the first rows of the profit function that correspond to the merging firms:

$$\begin{bmatrix} \Pi_1 \\ \pi_2 \end{bmatrix} = \begin{bmatrix} \mathbf{q}'_1 \\ \mathbf{Q}_2 \end{bmatrix} \begin{bmatrix} \mathbf{b} - \mathbf{c}_1 \\ \mathbf{b} - \mathbf{c}_2 \end{bmatrix} - \begin{bmatrix} \mathbf{q}_1 \\ \mathbf{Q}_2 \end{bmatrix}' \begin{bmatrix} (\mathbf{I} + \frac{1}{2} \mathbf{\Delta}_1) \mathbf{q}_1 \\ (\mathbf{I} + \frac{1}{2} \mathbf{\Delta}_2) \mathbf{q}_2 \end{bmatrix} - \begin{bmatrix} \mathbf{q}_1 \\ \mathbf{Q}_2 \end{bmatrix}' \begin{bmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} \end{bmatrix} \begin{bmatrix} \mathbf{q}_1 \\ \mathbf{q}_2 \end{bmatrix} \quad (\text{K.3})$$

where

$$\mathbf{Q}_2 \stackrel{\text{def}}{=} \text{diag}(\mathbf{q}_2) \quad (\text{K.4})$$

I have partitioned the profits vector into a scalar Π_1 , which collects the joint profits of the new entity, and vector π_2 , in which I stack the profits of all the other companies that are not included in the merger. If there are n firms and two of them are merging, this is a $(n - 1)$ dimensional column vector. The system of first order condition solved by the surviving firms is:

$$\begin{bmatrix} 0 \\ 0 \end{bmatrix} = \begin{bmatrix} \mathbf{b} - \mathbf{c}_1 \\ \mathbf{b} - \mathbf{c}_2 \end{bmatrix} - \begin{bmatrix} (2 + \delta_1) \mathbf{q}_1 \\ (2\mathbf{I} + \mathbf{\Delta}_2) \mathbf{q}_2 \end{bmatrix} - \begin{bmatrix} 2 \cdot \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} \end{bmatrix} \begin{bmatrix} \mathbf{q}_1 \\ \bar{\mathbf{q}}_2 \end{bmatrix} \quad (\text{K.5})$$

□

Proof to Proposition 4. An atomistic firm z has the following cost function:

$$h(z) = \frac{q^2(z)}{2z} \quad (\text{K.6})$$

Let p_{n+1} be the price of the good sold by the atomistic firms (recall we assumed they produce a homogeneous product). Because firms behave competitively and price at marginal cost, it must be the case that:

$$q(z) = p_{n+1} \cdot z \quad (\text{K.7})$$

hence labor supply and profits are given by:

$$h(z) = \frac{p_{n+1}^2}{2} \cdot z; \quad \pi(z) = \frac{p_{n+1}^2}{2} \cdot z \quad (\text{K.8})$$

Because the atomistic firms pay an entry cost of one unit of labor, the productivity cutoff for entry, which we call z_{\min} , is given by:

$$1 = \frac{p_{n+1}^2}{2} \cdot z_{\min} \quad (\text{K.9})$$

Let us now compute aggregate output and aggregate labor

$$\begin{aligned} q_{n+1} &= \int_{z_{\min}}^{\infty} q(z) dF(z) = p_{n+1} \int_{z_{\min}}^{\infty} z \cdot f(z) dz = \sqrt{\frac{2}{z_{\min}}} \int_{z_{\min}}^{\infty} \frac{\beta-1}{z^{\beta}} dz = \sqrt{2} z_{\min}^{\frac{1}{2}-\beta} \\ h_{n+1} &= \int_{z_{\min}}^{\infty} h(z) dF(z) = \frac{p_{n+1}^2}{4} \int_{z_{\min}}^{\infty} z \cdot f(z) dz = \frac{1}{z_{\min}} \int_{z_{\min}}^{\infty} \frac{\beta-1}{z^{\beta}} dz = z_{\min}^{-\beta} \end{aligned} \quad (\text{K.10})$$

By writing the productivity cutoff z_{\min} in terms of aggregate output q_{n+1} and plugging it in the expression for aggregate cost h_{n+1} , we find the aggregate cost function:

$$h_{n+1} = \left[\left(\frac{q_{n+1}}{\sqrt{2}} \right)^{\frac{1}{\frac{1}{2}-\beta}} \right]^{(-\beta)} = \left(\frac{q_{n+1}}{\sqrt{2}} \right)^{\frac{2\beta}{2\beta-1}} \quad (\text{K.11})$$

by taking the limit $\beta \rightarrow 1^+$, we can see that the expression above converges to the quadratic form:

$$h_{n+1} = \frac{q_{n+1}^2}{2} \quad (\text{K.12})$$

□

Proof to Lemma 3. The markup is equal to:

$$\mu_i = \frac{p_i}{MC_i} = \frac{p_i}{c_i + \delta q_i} = \frac{p_i q_i}{c_i q_i + \delta q_i^2} = \frac{p_i q_i}{h_i + \frac{\delta}{2} q_i^2} \quad (\text{K.13})$$

we can then use the fact that $\pi_i = (1 + \delta_i/2) q_i^2 = (p_i q_i - h_i)$ to write:

$$\mu_i = \frac{p_i q_i}{h_i + \frac{\delta}{2(1+\delta_i/2)} \pi_i} = \frac{(2 + \delta_i) p_i q_i}{(2 + \delta_i) h_i + \delta_i \pi_i} = \frac{(2 + \delta_i) p_i q_i}{2 \cdot h_i + \delta_i p_i q_i} \quad (\text{K.14})$$

□