Pricing with Markups under Horizontal and Vertical Competition

Extended Abstract

José R. Correa
Dept. Ingeniería Industrial
Universidad de Chile
jcorrea@dii.uchile.cl

Roger Lederman
Graduate School of Business
Columbia University
rlederman13@gsb.columbia.edu

Nicolás E. Stier-Moses
Graduate School of Business
Columbia University
stier@gsb.columbia.edu

ABSTRACT

We model a market for a single product that may be composed of sub-products that face horizontal and vertical competition. Each firm, offering all or some portion of the product, adopts a price function proportional to its costs by deciding on the size of a markup. Customers then choose a set of providers that offers the lowest total cost. We characterize equilibria of the two-stage game and study the efficiency resulting from the competitive structure of the market.

Categories and Subject Descriptors

J.4 [**Social And Behavioral Sciences**]: Economics—*non-cooperative games, oligopoly, bundling, supply function equilibrium*

1. INTRODUCTION

Classical models of competition, through either prices or production quantities, have focused predominantly on markets of a single good. In this setting, producers of substitutes, either perfect or imperfect, compete *horizontally* for the same pool of customers. Recently, there has been increasing recognition that in some industries competition among customers has a significant combinatorial component, beyond the scope of traditional single-market models. In such industries, producers whose goods may be purchased in combination compete *vertically.* In this paper, we study a framework where each producer chooses a price schedule, having in mind both the actions of vertical and horizontal competitors. Each producer sets its pricing schedule according to a price function that is specified as a constant percentage markup over per-unit production costs. Producers set the pricing schedule for individual products and customers choose a bundle of products at minimum price. Customers are interested in bundles composed of products given by any path of a series-parallel network. Producers on parallel links compete horizontally, while connections in series represent vertical relationships. Our model allows for many infinitesimal customers, or a single, centralized buyer. The latter case applies to a monopsony where a single buyer purchases components from multiple producers. In this setting, price schedules can be cast as contracts that depend on quantities.

BQGT'10 May 14-16, 2010 Newport Beach, California USA
ACM 978-1-60558-919-0.

Industries with this structure include those where a physical or geographic network is explicitly present, as in the airline industry, as well as those where a network structure is defined implicitly by the available bundles. To offer an example of the latter, consider a computer consisting of a CPU, keyboard and monitor. Buyers either create their computer by selecting the parts separately from a set of providers in each category, or by selecting an integrated model.

There is a large literature concerning centralized pricing, and a more recent body of work on price competition by decentralized firms. See, e.g., [2, 8, 7, 9, 1] for the case of substitutes, or [3, 12] for more general market structures. The growing literature on competition in networks has been motivated largely by applications involving a physical network. The focus has been the role of prices in guiding users towards efficient paths through the network. A feature of physical networks is often that customers experience costs due to congestion, with the effects increasing in the number of customers sharing a path. Our model is more in line with traditional models of competition in that customers experience no costs outside of the price that is paid to the producer. However, producers themselves experience costs that are marginally increasing, and pass this cost structure on to the customers through their price schedules. In this way, demand is encouraged to spread across multiple paths as in a network with congestion. Marginally-increasing unit costs are a common assumption in industries where capacity is constrained or costly to increase in the short term.

Supply function equilibria, popularized by the work of [10], represent a generalization of Cournot and Bertrand models of competition. In each of these cases, producers commit to either prices or production quantities before observing their competitors' choices, leaving only one of these as a possible lever for responding to the market. In the case of supply function equilibria, producers choose a function relating the price to the quantity produced (i.e., the inverse of the price function). Then, after all such functions have been chosen, the firms can adjust to the market conditions by choosing a point along this supply function. In equilibrium, each producer sells according to a single quantity/price combination, selected from its supply function so that the market clears. In this way, supply function equilibria model the common scenario where both price and quantity are adjusted in response to the market conditions. The paper [10] shows that, in the case of a duopoly, competition in supply functions leads to equilibrium prices and quantities intermediate between those of Bertrand and Cournot competition. In related research, [5] analyzes mergers of firms with quadratic cost functions in a market of substitutes using the framework of [10]. For a market of substitutes with a deterministic demand and nonlinear cost functions, [7] identifies an equilibrium where price functions have the same structure as cost functions and study their properties, particularly the efficiency of the resulting market.

The model we study assumes that all producers face a quadratic cost function, so that marginal costs are linearly increasing. We allow for heterogeneity among producers by applying an efficiency parameter that scales each cost function according to the specifics of the firm's production technology. Producers are assumed to select a price function that maintains the same shape as their cost function so that the per-unit price charged by each producer is a linear function of the quantity produced. This simplifies each producer's decision to a single parameter, which we call its markup. We refer to supply function equilibria where each firm is restricted to playing a markup over its cost function as *markup equilibria*. To motivate the restriction to markup equilibria, we show that even when producers may choose any non-decreasing price function, any markup equilibrium remains an equilibrium in the unrestricted game. This generalizes a result of [7] for the case of substitutes. In addition, [10] considers the game with general supply functions in the symmetric duopoly case, and shows that of the many equilibria existing when demand is deterministic, uncertainty eliminates all but the unique markup equilibrium.

As indicated, our model applies to general series-parallel networks. The network is assumed to have a single source s and sink t, and the possible bundles are represented by all paths connecting s and t. Hence, these bundles may be supplied by various combinations of producers. All bundles are assumed to be perfect substitutes, so that customers choose only on price, selecting the cheapest path through the network. Demand is deterministic and initially considered to be fixed, although an extension to the case of an elastic linear demand function is discussed. Once the producers have selected their price functions, customers make buying decisions choosing the cheapest bundle. As price functions are increasing, the price of each bundle increases with the number of customers purchasing, and so each customer's preferred bundle is dependent on the consumption choices of all other customers. Thus, the allocation itself is modeled as a game played amongst the customers where, at equilibrium, all bundles will sell for the same price.

In this paper, we fully analyze the two-stage game in which producers select price functions in the first stage, followed by the allocation game in the second stage. To our knowledge, we are the first to address equilibria of supply function games in a market with both substitutes and complements. We show that equilibria can be computed, and present examples of markup equilibria, focusing in particular on instances where the effects of a change in the market structure differ qualitatively from what a localized, single-market model would predict. We present a necessary and sufficient condition for the existence of equilibria, and show that the equilibrium is unique when it exists. For a fixed, inelastic demand, an equilibrium exists only in networks that are 3-edge-connected (see e.g., [4] for background on graph connectivity). Notice that this condition depends entirely on the topology of the network, and is independent of the cost parameters. For a network of substitutes, this is equivalent to requiring at least 3 producers to compete in the market [7]. Surprisingly, this matches results of existence of equilibria in related models [11, 8]. For general series-parallel networks, this condition rules out the case in which two producers within a bundle act as "monopolies" in that no other firm can replace them in that bundle. A similar problem was discussed in the network competition model of [3], and both scenarios are reminiscent of double marginalization, which is widely recognized as a source of inefficiency. When demand is elastic, the outside option provides sufficient horizontal competition, but a weaker existence condition is still needed to address the potential for vertical instability.

The best-response functions of producers have a highly intuitive structure: the per-unit price equals the per-unit cost plus a markup whose functional expression depends only on the markups of everybody else. We present its closed-form description, highlighting the fact that a producer's markup decreases with the introduction of horizontal competitors and increases with the introduction of vertical competitors. Further study of the comparative statics of equilibria shows that the relationship of a producer's market share and profits to the efficiency of its vertical competitors is more complex. Inefficiencies in markets for complementary items allow a producer to extract higher markups while simultaneously reducing the producer's market share. The net effect would be difficult to predict without an equilibrium model that includes vertical competitors.

For games with inelastic demand, we study the efficiency-loss at equilibrium. To that effect, the production cost at equilibrium is compared to the cost when customers are allocated optimally with respect to real production costs. In the two-stage game, the loss of efficiency results from the fact that markups distort the cost structure, leading customers to purchase less from those producers with high markups than they would in an optimum allocation. We show that when there is a sufficient level of competition, markups are bounded, and so the resulting allocation is not too inefficient. In contrast to the single-market model, where the addition of competitors increases efficiency, the addition of vertical competitors increases markups and can lead to inefficiency.

2. MODEL AND GAME STRUCTURE

We model a market for complementary goods by considering demand for a single good that we will call a bundle. Customers face multiple options for purchasing a bundle, and while each is equivalent in the eyes of the customer, they may be the result of production from a number of separate producers, each selling some portion of the good. We recognize that there need not be any single definitive way to divvy up production of a bundle, and so our model is general enough to allow each purchase option to be subdivided among any number of producers. Furthermore, each subdivision defines a production niche that multiple producers may compete to fill. Finally, among the various options for filling any particular niche, we consider that each may be further subdivided in some fashion and split among more specialized producers.

In general, we look at markets that takes a series-parallel (SP) structure. The class of SP networks are exactly those that can be constructed recursively through link subdivisions in series and in parallel. That is, through repeated application of the operations $\mathbf{DivS}(\cdot)$ and $\mathbf{DivP}(\cdot)$, pictured

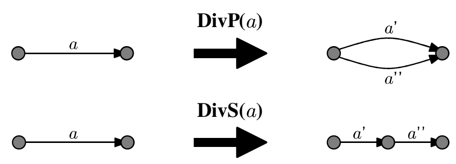

Figure 1: Operations defining an SP network.

in Figure 1. $\mathbf{DivS}(a)$ subdivides a link a into two links, connected in series, with a newly created node joining the head of a' and tail of a''. $\mathbf{DivP}(a)$ subdivides a link a into two links, a' and a'', each connecting the same two nodes as the original link a. Seeing as we do not limit the number or arrangement of subdivisions, SP networks provide a great deal of generality.

We model the set of available purchase combinations as paths from the source s to the sink t of an SP network, G, comprising a set of links $A_G = \{1, \ldots, n\}$. Each link $a \in A_G$ represents a producer, and each path through G a bundle that customers may purchase. Thus, denoting the set of available bundles by $\mathcal{B} := \{B_1 \ldots B_m\}$, we say for producer $a \in A_G$, that $a \in B_i$ if link a appears along path B_i in the network representation. In this way, the network defines a mapping of producers to purchase bundles. Each customer chooses a complete bundle, so that $\sum_{i=1}^{m} f_i = 1$ where f_i is the proportion of customers choosing bundle B_i. Then, the proportion of demand produced by producer a, is equal to $x_a = \sum_{B_i \ni a} f_i$. Because we interpret f_i and x_a as proportions, the total demand is normalized to one unit. We assume that individual customers are small so that demand is divisible among them, and each acts as a price taker. Although the discussion is for the case of inelastic demand, most results extend to the elastic case (see Section 3.1.1).

The per-unit production cost for each producer $a \in A_G$ is a function $u_a := \mathbb{R}_+ \to \mathbb{R}_+$ that depends on the production level x_a. We assume that all producers make use of similar 'technology' but some are more efficient than others. This is modeled by a cost function of the form $u_a(x_a) := c_a u(x_a)$ where the function $u(x_a)$ is an indication of the industry's unit cost for production level x_a, and the parameter c_a measures the efficiency of producer $a \in A_G$. This paper assumes that per-unit costs are linear, i.e., $u(x) = x$ (see Section 2.2.1). More generally, the model may assume that u is increasing, differentiable, and bijective (i.e., evaluates to zero at zero and grows to infinity). Furthermore, $xu(x)$ is convex; in other words, industries face increasing marginal production costs, which is the case, e.g., when labor or production capacity is scarce or when there is congestion. Putting all the elements together, the total cost to producer a of producing x_a units is $\kappa_a(x_a) := x_a u_a(x_a) = c_a x_a u(x_a)$, which is quadratic in this paper, and convex in general.

We consider a two stage game, where producers determine a pricing structure in the first stage, and customers choose a bundle of producers to purchase from in the second stage. In the first stage, producers commit to a price function $p_a(x_a)$ specifying the per-unit price to be charged at a specific level of production. Thus, both prices and production quantities are determined in the second stage, where the market clears. We assume that the price of a bundle is additive so that a customer purchasing bundle B_i pays a total of $\sum_{a \in B_i} p_a(x_a)$. Note that in the case of complementary items produced by the same producer, we would model purchase of both items by a single link. Thus, we are assuming additive pricing here only in the case of items purchased from competing producers. A critical feature of this structure is that the price a customer pays for a unit of production from producer a, depends on the total quantity that producer a produces, which itself is dependent on the consumption choices of all customers. This gives the second stage its interpretation as a game between customers.

We simplify the first-stage game by restricting the set of price functions a firm may choose. We consider only *markups*, in the sense that the producer a's price function $p_a(x_a) = \alpha_a u_a(x_a)$ for some positive factor $\alpha_a \geq 1$ that is chosen by the producer. We interpret α_a as a markup, due to the fact that α_a represents the ratio of price to production cost for producer a. This model is consistent with cost-plus pricing policies that are often employed in practice. Within this framework, the shape of all price functions is determined exogenously through the cost structure, and producers compete by selecting a single a parameter. This is not as restrictive an assumption as it may seem. Even in the setting where producers may choose any non-decreasing price-function, it can be shown that while there are in general many equilibria for the game, at any equilibrium in price-functions it is a best response to play a price function that is a markup of the producer's cost function. This robustness result was shown for a network of parallel links in [7]. The full paper presents an extension to the setting of SP networks. By exogenously setting the shape of all price functions as we have, we allow each to be described completely by a single *price multiplier* $w_a := c_a \alpha_a$. The actual unit price for product a is then given by $w_a u(x_a)$.

We seek to analyze the assignment of demand to specific producers. An assignment is described through either consumption decisions, using the vector $\vec{f} \in \mathbb{R}^m$, or through production quantities, as represented by the vector $\vec{x} \in \mathbb{R}^n$. The heterogeneity in our problem is across producers only, and so we will be primarily concerned with the production assignment $\vec{x}$. Note that for a given assignment $\vec{x}$, there may be multiple consumption allocations that give rise to $\vec{x}$. In particular, when we discuss uniqueness of an optimal or equilibrium production assignment, this need not imply uniqueness of the consumption assignment. We denote the set of possible production-consumption pairs by

$$\mathcal{F} := \Big\{ (\vec{x}, \vec{f}) \in \mathbb{R}_+^{(n+m)} : \sum_{i=1}^{m} f_i = 1, x_a = \sum_{B_i \ni a} f_i \;\forall a \in A_G \Big\}.$$

We say that a production allocation $\vec{x}$ is feasible if there exists a consumption assignment $\vec{f} \in \mathbb{R}_+^m$ such that $(\vec{x}, \vec{f}) \in \mathcal{F}$.

2.1 Submarket Structure

We now define the concept of a submarket. It will be helpful to introduce the composition operations $\mathbf{S}(\cdot)$ and $\mathbf{P}(\cdot)$, each of which takes as input a set $\mathcal{G}$ of SP networks, and returns a single SP network. In the case of $\mathbf{S}(\mathcal{G})$, the input networks are composed in series with the sink of one network doubling as the source node of the next. In the case of $\mathbf{P}(\mathcal{G})$, the input networks are composed in parallel so that all share a common source and sink.

As our market connects the source and sink nodes of an SP network, so a *submarket* is defined by a subnetwork connecting two nodes of G. Formally, a submarket g is a connected subnetwork of G, with two terminal nodes, a source

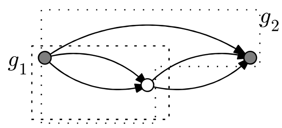

Figure 2: Example of a submarket (g_1), and a network that is not a submarket (g_2).

and sink, chosen from among the nodes in G, and the property that for any non-terminal node in g, all incident links are included in g as well. See Figure 2 for an example. The submarket g is self-contained in that it defines a product offering such that the output of any producer within g can be purchased only as part of that larger offering. The full market G is a submarket, as is any individual producer a. The flow into the source node of g represents the demand for the product this submarket produces. Were this quantity fixed, then competition on g would fit the form of our general model. As it is, a submarket strictly smaller than G faces an elastic demand, decreasing in the price of its offering. The price-sensitivity of this demand is determined by the price functions chosen by producers in $A_G \setminus A_g$.

The SP structure of G dictates that submarkets are arranged in a nested fashion. Each submarket g can be characterized as either a *series submarket*, indicating that $g = \mathbf{S}(\mathcal{G})$ for some set $\mathcal{G}$ of submarkets, or a *parallel submarket*, composed as $g = \mathbf{P}(\mathcal{G})$. We can give the set $\mathcal{G}$ of component markets comprising g an explicit name, denoting this set by $\psi(g)$. To avoid ambiguity, we require when g is a series submarket that all elements of $\psi(g)$ be parallel submarkets, and vice versa, so that $\psi(g)$ represents the largest (by cardinality) set of submarkets from which g can be formed in a single composition. In defining $\psi(\cdot)$, we have implicitly defined a tree structure that captures all submarkets; see Figure 3 for an example. Beginning with G as the root, $\psi(\cdot)$ determines a set of successors for each node. Every submarket appears as a node in this tree representation, with each producer $a \in A_G$ appearing as a leaf node. By convention, we will think of individual producers as series (parallel) submarkets, when their predecessor is parallel (series).

For an arbitrary vector $\vec{v} \in \mathbb{R}^n$ defined on the full set of producers, we use the notation $\vec{v}_A$ for the vector restricted to some set $A \subseteq A_G$. When g is an SP network representing some market, we abuse notation by referring directly to $\vec{v}_g$, with the understanding that this vector contains values for producers in A_g. In this respect, for two markets g and g', we have $g' \subseteq g$ if $A_{g'} \subseteq A_g$, and $g \setminus g'$ denotes the set of producers contained in A_g, but not in $A_{g'}$.

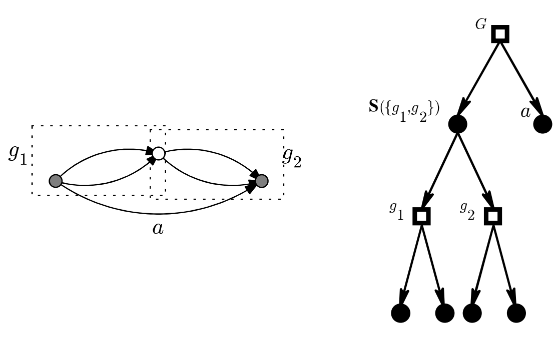

Figure 3: The network and submarket representations of a market G.

2.2 Optimal and Equilibrium Assignments

To quantify the quality of an assignment, we consider the total production cost $C(\vec{x}) := \sum_{a \in A} c_a x_a u(x_a)$ as a social cost function. This function captures whether customers are matched to the producers that are most efficient. Notice that payments are not considered in this function because they are internal transfers. The socially optimal assignment x^{OPT} is the unique production assignment minimizing $C(\vec{x})$. In other words,

$$(x^{\mathrm{OPT}}, f^{\mathrm{OPT}}) := \arg\min_{(\vec{x},\vec{f})} \left\{ C(\vec{x}) : (\vec{x}, \vec{f}) \in \mathcal{F} \right\}. \quad (1)$$

The production assignment is unique because $u(\cdot)$ is such that $x_a u(x_a)$ is convex for all $a \in A_G$.

An equilibrium for producers is a vector of markups $\vec{\alpha}$ that maximize the profits of all producers simultaneously, and an equilibrium for customers is an assignment f^{NE} such that all customers are buying at minimal price. These two games are played sequentially, making it a Stackelberg game. It will be convenient to think of the producers as setting price multipliers, leaving markups defined implicitly. So, in the markup game, producers first choose $\vec{w}$, followed by a second stage in which customers determine $f^{\mathrm{NE}}(\vec{w})$, and consequently determine a production assignment $x^{\mathrm{NE}}(\vec{w})$ as well. For a fixed $\vec{w}$, producer a realizes profits

$$\pi_a(\vec{w}) := (w_a - c_a)(x_a^{\mathrm{NE}}(\vec{w}))u(x_a^{\mathrm{NE}}(\vec{w})). \quad (2)$$

The equilibrium conditions imply that a tuple $(\vec{w}, x_a^{\mathrm{NE}}(\vec{w}))$ representing the two stages is at equilibrium if and only if

$$w_a \in \Phi_a(\vec{w}_{-a}) := \arg\max_{w \geq 0} \{\pi_a(w, \vec{w}_{-a})\} \text{ for all } a \in A_G \quad (3)$$

where $\Phi_a(\cdot)$ is the best response function of producer a to the price multipliers of all other producers, denoted by $\vec{w}_{-a}$. Here, the second stage assignment $(x^{\mathrm{NE}}(\vec{w}), f^{\mathrm{NE}}(\vec{w})) \in \mathcal{F}$ is defined for an arbitrary vector $\vec{w}$ and satisfies the condition

$$\textstyle\sum_{a \in B_i} w_a u(x_a^{\mathrm{NE}}(\vec{w})) \leq \sum_{b \in B_j} w_b u(x_b^{\mathrm{NE}}(\vec{w})) \quad (4)$$

for all $B_i, B_j \in \mathcal{B}$ such that $f_i^{\mathrm{NE}}(\vec{w}) > 0$. The above inequality says that in any equilibrium of the second-stage game, all bundles sell, if at all, at a single minimal price.

The uniqueness of $\vec{w}$ is established later, but at this point it is clear that $x_a^{\mathrm{NE}}(\vec{w})$ is unique for any $\vec{w}$ because the function $u(\cdot)$ is strictly increasing [6]. Notice that price distortions driven by producers with market power, as well as potential negative externalities in the second stage, make it such that the markups $\vec{w}$ may not give rise to the most efficient equilibrium assignment. Rather, it is likely that $C(x^{\mathrm{NE}}(\vec{w})) > C(x^{\mathrm{OPT}})$.

2.2.1 Linear Unit Cost Functions

We now specialize to the case of quadratic cost functions, which we generate by considering $u(x) = x$. Hence the total cost for producer a has the form $\kappa_a(x_a) = c_a x_a^2$. This assumption will remain in force throughout. From a technical point of view, this assumption allows us to explicitly characterize the optimal assignment and the unique assignment corresponding to a given vector of markups. Indeed, in this situation $C(\vec{x})$ is a convex function and the absence of any fixed costs ensures that all producers are active under both assignments. Furthermore, the restriction to linear unit costs is sufficient to ensure that customers are efficient in the sense that for fixed $\vec{w}$, $x^{\mathrm{NE}}(\vec{w})$ minimizes $\sum_{a \in A} x_a w_a u(x_a)$. Thus, their behavior in the game is consistent with that of a centralized buyer with the ability to split consumption optimally among the bundles. In this case, our model can be viewed as a demand-dependent contract between producers and a monopsonistic buyer.

Seeing as customers behave efficiently, any inefficiency in the assignment is the result of distortion of the true cost functions due to producer markups. It follows that when markups are not distortionary, the equilibrium assignment will match the socially optimal assignment.

THEOREM 1. *The unique socially optimal assignment x^{OPT} is equal to $x^{NE}(\vec{c})$.*

Thus, the optimal assignment matches the second stage equilibrium that results when producers charge their actual costs without markup. In the next section, we characterize the second-stage equilibrium assignment for any fixed set of price functions, which will include the optimal assignment as a special case.

2.3 Assignment Game

In this section we present a precise functional form of the second-stage assignment $x^{\mathrm{NE}}(\vec{w})$, where $\vec{w}$ is a fixed vector of price multipliers. To start, we introduce the *network price multiplier* $R_g(\vec{w}_g)$, which generalizes w_a to a submarket g. When a demand of x_g is assigned to g according to (4), the market price for a bundle is $R_g(\vec{w}_g)x_g$. As $\vec{w}$ is fixed prior to the assignment game, we use the notation R_g with the understanding that the multiplier reflects the combined effects of a set of price functions selected by individual producers in the first-stage game. Since demand for market G is normalized to one unit, R_G is also the equilibrium price of a bundle under $\vec{w}$.

For an individual producer, $R_a = w_a$. For a larger submarket g, R_g depends ultimately on the proportions in which customers choose from among the bundles in g. A full characterization is obtained inductively by:[1]

$$R_{S(\mathcal{G})} = \textstyle\sum_{g \in \mathcal{G}} R_g, \quad \text{and} \quad R_{P(\mathcal{G})} = (\sum_{g \in \mathcal{G}} 1/R_g)^{-1}. \quad (5)$$

which follow since customers allocate to parallel submarkets in inverse proportion to their price multipliers.

We say that a producer a *spans* the market if link a connects s and t directly. In this case, all other bundles are substitutes for a, and x_a is increasing in the multipliers of all competitors. If, on the other hand, producer a faces vertical competition, the residual demand for product a is shifted downwards as the markups on complementary items increase.

We now express x_a in terms of aggregate measures of the horizontal and vertical competition faced by a. Our approach is to redefine the market by pivoting G so that the nodes incident to a become the source and sink. In this reformulation, denoted $G \odot a$, a spans the market and all competition with a is horizontal. To interpret, the market spanned by a is one in which all customers come to market in possession of a bundle that is perfectly complementary to

[1]Equation (5) matches that used for electrical circuits to compute the equivalent resistance when placing resistors in series and parallel. Ohm's law, *Voltage = Current · Resistance*, is analogous to the price function $p_a = x_a R_a$.

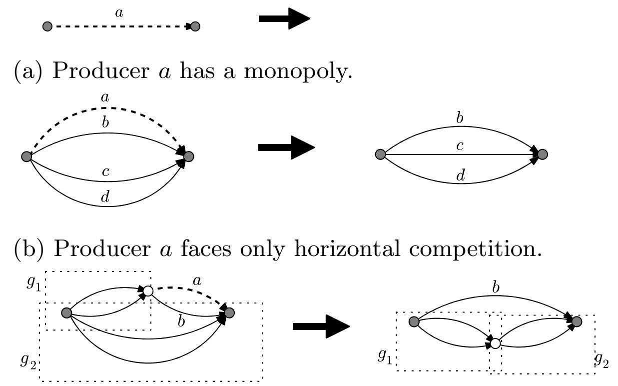

(a) Producer a has a monopoly.

(b) Producer a faces only horizontal competition.

(c) Producer a compites horizontally and vertically. Links in g_1 are reversed to form $G \ominus a$.

Figure 4: The producer's substitute network $G \ominus a$.

a. To complete their bundle, customers may purchase from a or one of its direct competitors. In addition, there is the option to 'sell back' the complementary items and purchase a new bundle. Any such action can be represented as a path through $G \odot a$. In the course of pivoting G, any complementary links to a; i.e., those on a path from s to a or a to t, are reversed in direction to reflect that these products are sold back to producers at the prevailing market price. Any combination of sales/purchases that forms a path through the pivoted network will leave the customer with a complete bundle, and is in effect a perfect substitute to a. Accordingly, we call the network created by removing a from $G \odot a$, the *substitute network* for producer a. The substitute network is denoted by $G \ominus a$, and its construction is demonstrated in Figure 4. The example in (c) contains vertical competition and so requires pivoting. The total demand in the reformulated game is subject to an adjustment factor for the amount of vertical, rather than horizontal, competition that is faced by producer a. The scaling factor $\mu_a(\vec{w}_{-a})$ is the demand for producer a when w_a is equal to zero. When a faces vertical competition, this will be strictly less than one, and decreasing in the markups demanded for complements of a. In general, the factor $\mu_a(\vec{w}_{-a})$ may be increasing, decreasing, or unaffected by w_b, depending on whether b is largely a substitute or a complement of a. Since $\vec{w}$ is fixed we suppress the dependency and will use the notation μ_a.

The uniqueness of the niche that producer a fills will determine the multitude of paths in $G \ominus a$, and play a key role in determining market power. A measure of this is $R_{\ominus a}$ defined for $G \ominus a$ according to (5). If producer a is a monopolist, then $G \ominus a$ is empty, indicating that customers have no choice but to purchase from a. Since there is no price at which a substitute can be purchased, we say $R_{\ominus a} = \infty$ in this case. In general, $R_{\ominus a}$ measures the market power of producer a in equilibrium, with a higher multiplier indicating a relative absence of attractive alternatives to producer a. In Proposition 1 we show that $R_{\ominus a}$ in fact determines the slope of producer a's residual demand with respect to its markup, α_a.

PROPOSITION 1. *For any SP network with price functions fixed according to $\vec{w}$, and for any link a, the equilibrium* *assignment* $x^{NE}(\vec{w})$ *takes the form*

$$x_a^{NE}(\vec{w}) = \mu_a\left(\frac{R_{\ominus a}}{R_{\ominus a} + w_a}\right). \quad (6)$$

By describing the assignment as in (6), we can in turn describe the total production costs as a function of the markups. According to Proposition 1

$$C(x^{\mathrm{NE}}(\vec{w})) = \sum_{a \in A_G} c_a \left[\mu_a\left(\frac{R_{\ominus a}}{R_{\ominus a} + w_a}\right)\right]^2. \quad (7)$$

From Theorem 1, we know that x^{OPT} is precisely $x^{\mathrm{NE}}(\vec{c})$. Recalling that $R(\vec{w})$ is the price of any bundle in equilibrium, the optimal production cost is

$$C(x^{\mathrm{OPT}}) = \textstyle\sum_{i:B_i \in \mathcal{B}} f_i(R_G|_{\vec{w}=\vec{c}}) = R_G|_{\vec{w}=\vec{c}}\,. \quad (8)$$

For an arbitrary equilibrium assignment, the total payment by customers is $R_G = \sum_{a \in A_G} w_a \left[x_a^{\mathrm{NE}}(\vec{w})\right]^2$. In general, it is not the case that $C(x^{\mathrm{NE}}(\vec{w})) = R_G$, since some portion of these payments is kept by the producers as profit. To study efficiency of an assignment $x^{\mathrm{NE}}(\vec{w})$, we compare the cost $C(x^{\mathrm{NE}}(\vec{w}))$ to $C(x^{\mathrm{OPT}})$. Ultimately, the degree of inefficiency will depend on $\vec{\alpha}$, which is the outcome of the strategic choices taken by producers in the first-stage game.

The size of producer a's markup in the first stage will depend on its market power, as made explicit here.

PROPOSITION 2. *In the first-stage game, the best response function for any producer* a *satisfies*

$$\Phi_a(\vec{w}_{-a}) = 2c_a + R_{\ominus a}\,. \quad (9)$$

In terms of $R_{\ominus a}$, the per-unit price that producer a will charge in equilibrium is $p_a(x_a) = w_a x_a = 2c_a x_a + R_{\ominus a} x_a$. Producer a's costs are given by $\kappa(x_a) = c_a x_a^2$, yielding a marginal cost of $\partial\kappa(x_a)/\partial x_a = 2c_a x_a$. Thus, equilibrium prices can be interpreted intuitively to consist of marginal costs of production, plus a markup of $R_{\ominus a} x_a$. Furthermore, using $R_{\ominus a}$ as a measure of market power, the markup that can be extracted is directly related to the level of competition faced. As the competition faced by producer a increases, $R_{\ominus a}$ will decrease. In the extreme case, where $R_{\ominus a}$ grows small, the price will approach the marginal cost, in accordance with the interpretation as a competitive market. If $R_{\ominus a}$ grows small for all producers, then the markup vector approaches $2\vec{c}$, so that the equilibrium assignment approaches x^{OPT}. That is, perfectly competitive markets are efficient.

2.3.1 Submarket Price Functions

The result of Proposition 1 will extend to any submarket g, with R_g in place of w_a. The substitute network $G \ominus g$ for any submarket g is defined analogously to $G \ominus a$. Furthermore, we can consider alternatives to g within a larger submarket g', by constructing the graph $g' \ominus g$, and assigning the price multiplier $R_{g' \ominus g}$ to this market. For convenience, when the outer submarket in this construction is the full market G, we use the shorthand notation $R_{\ominus g}$, leaving the specification of G as an implicit assumption.

Once price multipliers have been fixed, a submarket g is equivalent to a single producer with a price function of $p_g(x_g) = R_g x_g$. We can also define the aggregate choice of R_g by a *submarket response function* $\phi_g(R_{\ominus g})$, reflecting the interaction of producers in g, although this need not take the form of (9). We show that this is indeed a single-valued function. Existence of the underlying equilibrium is explored in the next section.

For the case of a single producer a, $\Phi_a(\vec{w}_{-a}) = \phi_a(R_{\ominus a}) = 2c_a + R_{\ominus a}$, so that $\phi_g(\cdot)$ generalizes the best response function $\Phi_a(\cdot)$ while making the dependence on the substitute network explicit in the definition. Accordingly, a markup equilibrium $\vec{w}$ satisfies $w_a = \phi_a(R_{\ominus a})$ for all producers a.

3. EQUILIBRIUM OF MARKUP GAME

Each producer selects w_a to satisfy the best-response map $\Phi_a(\vec{w}_{-a}) = 2c_a + R_{\ominus a}$. A Nash equilibrium of the markup game is a vector $\vec{w}$ satisfying $w_a = \Phi_a(\vec{w}_{-a})$ for all $a \in A_G$. It is clear from (5) that $R_{\ominus a}(\cdot)$ is a continuous function, and so $\Phi_a(\cdot)$ is continuous and single-valued. Combining the producers' individual best response functions yields a continuous vector-valued function $\Phi(\vec{w})$ whose fixed points, if any exist, correspond to equilibrium markups. If the image of $\Phi(\cdot)$ over the domain $\vec{w} \in \mathbb{R}_+^n$ is contained within $X = \prod_{a \in A_G} X_a$ with $X_a \subset [2c_a, \infty)$, we can, without loss of generality, define markup equilibria as fixed points of the function $\tilde{\Phi} : X \to X$ where $\tilde{\Phi}_a(\vec{w}) := \Phi_a(\vec{w}_{-a})$. Making use of Brouwer's fixed point theorem, a sufficient condition for existence of a fixed point of $\tilde{\Phi}(\cdot)$ is compactness of X. If producer markups are bounded so that $w_a \le \bar{w} < \infty$ for all producers, then we define X_a by the compact set $[2c_a, \bar{w}]$, and apply the fixed point theorem. We proceed in this section by deriving conditions which guarantee the existence of $\bar{w}$. Essentially, an equilibrium requires sufficient competitive pressure to prevent any producer from continually increasing the size of their markup.

3.1 Competition and Graph Connectedness

In this section, we explore the existence of an upper bound on $\vec{w}$ in a market with inelastic demand. We will see that the critical property in establishing a bound is the degree of connectedness of the network structure. A set of links whose removal disconnects the graph is a *cut*. A graph is k*-edge-connected* if there are no cuts containing less than k links [4]. For example, Figure 5 shows a 2-connected network.

The intuition for studying connectedness was provided in Section 2.3, where we noted that a producer's market power is directly related to the uniqueness of the producer's niche, which is reflected in the number (and ultimately, price) of alternative paths available for joining the nodes that the producer connects in G. The connectedness of the graph indicates the smallest set of producers such that one must be used in any path connecting some pair of nodes in the network. A high degree of connectedness should translate to some bound on the market power of any individual producer. In this section we formalize this idea.

For a submarket g, the connectedness $Q(g)$ is the largest k for which g is k-edge-connected. A directed cut is one that

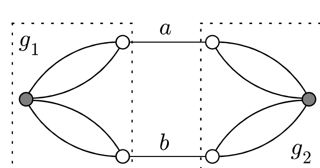

Figure 5: A 2-connected network. Producers a and b make up a cut.

divides the graph so that the source and sink are disconnected. A cut that does not separate the source and sink is a *vertical cut* in that the producers in the cut belong to some common bundle and compete vertically. If $Q(g) = k$, then there can be neither any directed cuts, nor any vertical cuts, that contain less than k links. As such, link directions play no role in determining $Q(g)$. Redefining connectedness in terms of vertical cuts alone gives the vertical connectedness $V(g)$. In general, composing g with producers in parallel can increase $Q(g)$, but $V(g)$ provides an upper bound on the connectedness of any market within which g is nested.

It is clear from the response functions that in the case of a duopoly, the combined sensitivity of the producers leads to an infinitely increasing sequence of markups. This applies as well to any network with $Q(G) < 3$. Although G is directed, instability can result from both directed and vertical cuts. Essentially, stability requires that the substitute network for any producer is 2-connected in its directed form. The cause of instability in any case where $G \ominus a$ is not 2-connected is similar to that of the duopoly case, where $G \ominus a$ consists of a single link for either producer. In the general case, it is producer a and the producer that disconnects $G \ominus a$ that combine to drive the instability. We will show that when the graph is 3-connected, there is enough competition to ensure that markups are bounded.

THEOREM 2. *A markup equilibrium exists iff the network is 3-edge-connected.*

By bounding $\vec{w}$, we restrict the image of $\tilde{\Phi}(\vec{w})$ to a compact set, assuring the existence of a markup equilibrium. We observe further that $\Phi_a(\vec{w}_{-a})$ is increasing in w_b for all $b \neq a$. As a result, any sequence $\{\vec{w}^\tau\}$ with $\vec{w}^\tau = \tilde{\Phi}(\vec{w}^{\tau-1})$ will be increasing element-wise. Starting at $\vec{w}^0$ with $w_a^0 = 2c_a$ for all $a \in A_G$, we generate a sequence of markups that must converge to a markup equilibrium. Applying iterated best responses, we are able to compute the unique markup equilibrium in this way for any game that satisfies the 3-connectedness condition.

COROLLARY 1. *If a markup equilibrium exists, it can be approximated by iterating best responses.*

At an equilibrium, $\vec{w}$, we have $\phi_a(R_{\ominus a})/R_{\ominus a} = w_a/R_{\ominus a} = w_a/\phi_{G\ominus a}(w_a)$. In Theorem 3, we show that the first and last ratios in the equality are monotonically decreasing and increasing in $R_{\ominus a}$, respectively. This ensures that an equilibrium can exist for at most one value of $R_{\ominus a}$, and so for at most one value of w_a.

THEOREM 3. *If a Nash equilibrium, $\vec{w}^{NE}$, exists in the markup game, then it is a unique equilibrium.*

3.1.1 Extension to Elastic Demand

The network multiplier R_G is itself the result of a response function $\phi_G(\cdot)$ whose argument is set exogenously, and reflects the price multiplier of an option outside of the market. To this point, by assuming that no alternatives to G exist, we have used implicitly that $R_{\ominus G} = \infty$. By choosing $R_{\ominus G} < \infty$ we allow for an elastic demand. The price of an outside alternative, g_0, and by extension the willingness to pay for market G, can be defined by a fixed multiplier $R_{\ominus G}$ applied to the function $u(1 - x_G)$ where x_G is the demand assigned to market G. With linear unit costs, the demand takes the form $p_G = R_{\ominus G}(1 - x_G)$, or $x_G = 1 - p_G/R_{\ominus G}$, yielding a model of linear demand and quadratic total costs.

Here we extend Theorem 2 to a market with elastic demand. In network form, the elastic demand model includes an outside option, spanning the entire market, with a fixed markup. Adding an outside option does not introduce any instability in markups, but may provide enough competition that markups are bounded for a market structure that does not yield an equilibrium with inelastic demand. The conditions for existence are thus weaker than in the inelastic case.

COROLLARY 2. *For a market G with elastic demand, define G^+ as $\mathbf{P}(G, g_0)$, where the outside option g_0 consists of two parallel links. A markup equilibrium exists iff G^+ is 3-edge connected.*

The corollary includes the possibility that an equilibrium does not exist, even in an elastic demand market. The outside option provides stability when there is a shortage of purchase options, as with a monopoly or duopoly, by assigning some value to not purchasing. If there is a lack of competition vertically within some bundle, this will persist in the elastic case. Thus, an equilibrium exists in an elastic demand market if and only if the vertical connectedness is at least 3.

If there is some vertical instability, but the directed network remains 3-connected, there may still be an equilibrium on some subnetwork spanning the market (this holds trivially for an elastic demand market, because the price never exceeds R_{g_0}). In the next section, we see that in such cases we can simply ignore the paths with vertical instability.

3.1.2 Irrelevance of Inefficient Submarkets

In both inelastic and elastic demand markets, when the competition in a market is insufficient, then producers will continually have an incentive to increase markups, so that no equilibrium exists. Yet, because we allow for asymmetric market structures, it may happen that some producers face sufficient competition while others do not. Intuitively, if some subnetwork of producers, spanning the market vertically, supports an equilibrium, while other producers raise their markups infinitely, we expect that eventually all customers will abandon the unstable producers and adopt the equilibrium assignment consistent with the stable set of production bundles.

This observation allows us to study some markets that are not 3-connected, but do have a 3-connected substructure embedded within. That is, we consider a market G^+ that is an *extension* of a 3-connected market G. To ensure that G spans the market vertically, we define an extension as the addition of competition in parallel to a submarket of G. Formally, we form G^+ by replacing some submarket g' of G with $\mathbf{P}(g', g^+)$, where g^+ is an SP network.

First, we consider the case where G^+ remains 3-connected, but costs for producers in g^+ are prohibitively large. We show that g^+ can be ignored.

THEOREM 4. *Let $\vec{w}^*$ be an equilibrium of the markup game on G, and G^+ be a 3-connected extension of G, with $c_b^l \to \infty$ simultaneously for all $b \in G^+ \setminus G$, when $l \to \infty$. Then, the sequence $(\vec{w}_G^l, \vec{w}_{G^+\setminus G}^l)$ of extended equilibria converges to $(\vec{w}, \infty)$ as $l \to \infty$.*

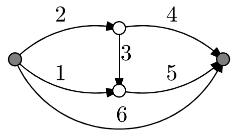

Figure 6: A non series-parallel 3-connected network.

In Theorem 4, G^+ is 3-connected, so each instance of the extended network is one that we can analyze on its own. In contrast, we look next at the case where g^+ introduces some instability into the market through its competitive structure. We note that g^+ needs not be 3-connected for G^+ to be so. For instance, adding a single link to G cannot reduce the connectivity. However, it may be the case that g^+ is vertically unstable, in that removing two producers from g^+ disconnects g^+ into three disconnected markets. In this case, no extension formed from g^+ will produce an equilibrium.

The extended network is no longer 3-connected, and when best-responses are iterated, it must be that $\vec{w}^{\tau}_{G^+\setminus G} \to \infty$, where τ is the number of iterations applied. (Unlike in Theorem 4, costs in the extended network remain constant throughout.) We will show that while the sequence $\{\vec{w}^{\tau}_{G^+\setminus G}\}$ is unbounded, $\vec{w}^{\tau}_G \to \vec{w}^*$, with the limit an equilibrium in G. In this way, we expand the set of networks we may analyze to include all extensions of 3-connected networks.

THEOREM 5. *Let $\vec{w}^*$ be equilibrium markups for G. If $G^+ \ominus G$ is vertically unstable, then multipliers on G converge to $\vec{w}^*$ when best responses are iterated on G^+.*

The result supports the use of iterated best responses to analyze any market in which there is an embedded 3-connected network spanning from source to sink. We have defined an extension as the addition of a single submarket, but repeated application of Theorem 5 allows for more general structures.

3.2 Other Network Structures

As the following example will demonstrate, Theorem 2 does not immediately generalize to networks that are not Series-Parallel. Figure 6 presents a very simple network structure that is 3-edge-connected, but violates the restriction to SP structure. No markup equilibrium exists for this network.

Critically, when the network is not SP, we cannot guarantee that all producers are active in equilibrium. In Figure 6, producer 3 is offering a contribution to the bundle that is evidently being offered by producers 1 and 4 as well. Here producer 1 is offering the equivalent of products 2 and 3 in combination. Similarly, producer 4 is offering the equivalent of products 3 and 5 in combination. If the markups and demand allocation are such that the prices for products 1 and 4 are less than the prices of products 2 and 5, respectively, then producer 3 is in effect excluded from the market. There is no markup that producer 3 can choose for which customers will purchase product 3.

When this the case, the price function for product 3 does not influence the second stage results, and as such does not factor into the profits of other producers. Consequently, when producer 3 is not active, we can eliminate them from the analysis entirely, with no affect on the equilibrium. The remaining producers then comprise a series-parallel network that is not 3-connected. There is no equilibrium in such a network, so producer 3 must be active in any equilibrium.

We show that, regardless of the cost structure, no equilibrium can exist in which producer 3 is active, and thus no markup equilibria exist for this 3-connected network.

4. SENSITIVITY ANALYSIS

In this section we study the effects of changes to market parameters and structure. We pay particular attention to insights that can only be obtained through a model that includes vertical competition. We begin by looking at the market price for a bundle, R_G, which is closely related to the local competitiveness of submarkets. This relationship is not as clear when considering the social cost criterion, and we look at the implications for understanding efficiency in the market. We show how a model of vertical competition can benefit a regulator who is analyzing a particular submarket. Finally, we consider the role of vertical competition in a predictive model used by a producer. We illustrate the effect this can have on a producer's decision-making.

Recall that R_G is the market price in equilibrium for a given network, and $C(x^{\text{OPT}}) = R_G|_{\vec{w}=\vec{c}}$ is the cost of satisfying demand in a socially optimal manner. A comparison of these terms gives a measure of the extent to which bundles have been marked up. In particular, $R_G/(R_G|_{\vec{w}=\vec{c}})$ measures the 'average' markup, and $R_G - R_G|_{\vec{w}=\vec{c}}$ is equal to the total producer profit. In terms of social cost, $C(x^{\text{NE}}(\vec{w}))$ evaluates a markup vector $\vec{w}$, and the ratio $C(x^{\text{NE}}(\vec{w}))/(R_G|_{\vec{w}=\vec{c}})$ determines the inefficiency of that vector. As such, we are particularly interested in changes to the market structure for which $R_G|_{\vec{w}=\vec{c}}$ remains constant, so the effects on profits and efficiency can be observed through R_G and $C(x^{\text{NE}}(\vec{w}))$ alone. One interpretation of such a change, focusing only on network structure, is that of a merger. Here, the underlying cost structure of the market remains unchanged, but we consider changes in the ownership of production capacity. A formal definition is given below.

4.1 Price of a Bundle and Industry Profits

We will we look at the ramifications of a shift in the response function for a single producer or some subnetwork of producers. By a shift, we mean that $\phi_g(R_{\ominus g})$ is replaced by a function $\hat{\phi}_g(R_{\ominus g})$ such that (for an upwards shift) $\hat{\phi}_g(R) \geq \phi_g(R)$ for all R in the domain. For a single producer, a shift in $\phi_a(\cdot)$ can result only from a change in c_a, but for a submarket g it can be the result of any number of structural or parametric changes within g. We show that an upwards (downwards) shift will always lead to an increase (decrease) in the equilibrium price, R_G.

LEMMA 1. *An upwards shift in $\phi_g(R_{\ominus g})$, for any submarket g of G leads to an increase in the equilibrium price, R_G.*

COROLLARY 3. *An increase in c_a for any producer a leads to an increase in equilibrium price for G. When a new link is added in parallel to any submarket, the equilibrium price for G decreases.*

The corollaries follow immediately from Lemma 1. They illustrate a consistency between the local competitiveness of sub-products and the competitiveness of the market as a whole, with respect to fundamental changes in the production capabilities. Locally, it is clear that we are decreasing efficiency in the first part of Corollary 3, and adding competition in the second. These effects extend directionally to the entire network. We next consider mergers, where the production capacity is preserved in a certain sense. In this setting, any changes in the bundle price will result entirely from changes in the way the producers interact.

We define a *merger* as a change in the network structure where multiple links are combined in a way that preserves the aggregate cost structure. The cost of the new link should match the cost of using the subnetwork that it replaces, assuming that flow is allocated optimally within the original subnetwork. We denote the optimally aggregated cost of a submarket g by c_g. The procedures for aggregating costs optimally are identical to those for aggregating price multipliers:

$$c_{S(\mathcal{G})} = \sum_{g\in\mathcal{G}} c_g, \quad \text{and} \quad c_{P(\mathcal{G})} = (\sum_{g\in\mathcal{G}} 1/c_g)^{-1}. \qquad (10)$$

We first look at horizontal mergers, where two parallel links, a_1 and a_2 are combined to form a_P. We denote the cost of the merged producer by c. Letting $p := c_{a_2}/(c_{a_1} + c_{a_2})$, we have $c_{a_1} = c/p$ and $c_{a_2} = c/(1-p)$. Any horizontal merger can be described in this way for some $p \in (0, 1)$. We show that any such merger results in an upward shifted response function, $\phi_{a_P}(\cdot)$ relative to the aggregated response function $\phi_{g_P}(\cdot)$, where $g_P = \mathbf{P}(\{a_1, a_2\})$. In fact, for any such parallel pair, prices increase with p, yielding the lowest prices in the symmetric case.

THEOREM 6. *Horizontal mergers increase the equilibrium price of a bundle.*

We complete our discussion of mergers with the case of mergers involving vertical competitors. The ability to analyze vertical mergers is a distinct benefit of our SP market definition. Two producers in series is an unstable configuration, so we will not analyze mergers originating from this structure. Rather, we look at the case of a single producer in series with a set of producers who compete with each other horizontally. We consider the effect of consolidating all of these producers to a single link. We interpret this as a scenario where the production being offered by the parallel competitors is carried out in-house by the producer occupying a single link. In this way we study the effect of *vertical integration*.

Let c be the cost of the integrated producer a_S. We consider parallel producers a_1 and a_2, comprising a submarket g_P. The submarket g_P is connected in series with a third producer a_V to form $g_S = \mathbf{S}(\{a_V, g_P\})$. We require $c_{g_S} = c$, and in particular, for $p, q \in (0, 1)$, let $c_{a_1} = \frac{cq}{p}$, $c_{a_2} = \frac{cq}{1-p}$, and $c_{a_V} = 1 - q$. For any choice of p and q, $\phi_{a_S}(\cdot)$ is a downward shift of $\phi_{g_S}(\cdot)$, so that vertical integration results in a lower price than the subcontracting setup. The size of the effect is increasing in q, and thus greatest when the local monopolist, a_V, controls a small portion of the market g_S.

THEOREM 7. *Vertical integration decreases the equilibrium price of a bundle.*

4.2 Social Cost

Lemma 1 provides a formal connection between the local competitiveness of a submarket and the overall size of the markup on a bundle. The equivalent connection need not exist for costs, providing a clear motivation for consideration of vertical competition in a regulatory context.

For a producer a competing in a parallel submarket g of G, that is $g = \mathbf{P}(\mathcal{G})$ with $a \in \mathcal{G}$, an alternative to a full model of G would be to estimate the parameters μ_g and $R_{\ominus g}$, or equivalently to estimate the demand function for the submarket g, and treat g as a market of exclusively horizontal competition, subject to an elastic demand. This may be a reasonable approach, and is in line with the way in which producers choose their markups in our model (actually we assume an even more restricted viewpoint in which $R_{\ominus a}$ is held fixed to generate a residual demand). However, a localized approach of this type can be misleading with regard to social cost.

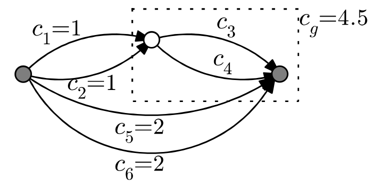

Figure 7: Total production cost in this market is smaller when c_3 and c_4 are unequal.

Relative to total profit, social cost depends on the symmetry, rather than the size of markups. So, when a bundle consisting of product g is inherently more expensive to produce than substitute bundles, likely leading to high markups on those substitutes, what is perceived locally as an inefficiency in the market for g may be a force that drives markups on g closer to those on substitutes, in effect reducing the degree of distortion in the overall market. We proceed with an example to demonstrate this possible effect.

Consider the market G in Figure 7, with submarket g. The market for product g is a duopoly, and for a fixed multiplier R_{-g}, the two producers face an elastic combined demand. For a regulator considering the market for g as such, the most efficient configuration of the market, would appear to be the symmetric one. For comparison, we consider the efficiency that results in G when the symmetry of the submarket g is adjusted, but the aggregate cost structure, as measured by $(\frac{1}{c_3} + \frac{1}{c_4})^{-1}$, is held constant.

When costs in g are symmetric, $c_3 = c_4 = 9$. In this case, the optimal allocation has 83.3% of customers purchasing from producers 5 and 6. The average cost of a bundle under the optimum is 0.83. The equilibrium allocation is presented in Table 1. Producers 5 and 6, each being more efficient than the other purchase combinations, apply relatively large markups of $\alpha_5 = \alpha_6 = 5.1$ to their products. In comparison, the price of a combination purchase from the other producers is only 3.1 times larger than the cost of 10. This distortion encourages a larger proportion of costly combination purchases, and the average cost of a bundle in equilibrium is 0.87.

In another scenario, producer 3 is considerably more efficient than producer 4. If $c_3 = 5$ and $c_4 = 45$, the aggregate cost structure is unchanged; the optimally aggregated cost is 4.5 as in the symmetric case, while the relative efficiency becomes unequal. Studying g in isolation would suggest that this arrangement is inefficient. Yet, the higher markups applied by producer 3 raises the price of a combination purchase to 3.9 times the cost. This shifts some demand back to producers 5 and 6, so that the average cost of a bundle falls to 0.86, despite the asymmetry. Although the difference in social cost between these two scenarios is rather small, the direction of change is surprising as it goes contrary to what a local model of market g suggests.

Social Optimum	Pr. 1	Pr. 2	g		Pr. 5	Pr. 6	Market (G)
Efficiency (c)	1	1	4.5		2	2	.833
Market Share (x)	.083	.083	.167		.417	.417	1
Cost (cx^2)	.007	.007	.126		.347	.347	.833
Symmetric Costs	Pr. 1	Pr. 2	Pr. 3	Pr. 4	Pr. 5	Pr. 6	Market (G)
Efficiency (c)	1	1	9	9	2	2	.833
Markup (α)	6.96	6.96	2.70	2.70	5.08	5.08	4.60
Market Share (x)	.123	.123	.123	.123	.377	.377	1
Cost (cx^2)	.015	.015	.135	.135	.285	.285	.870
Asymmetric Costs	Pr. 1	Pr. 2	Pr. 3	Pr. 4	Pr. 5	Pr. 6	Market (G)
Efficiency (c)	1	1	5	45	2	2	.833
Markup (α)	7.56	7.56	3.69	2.14	5.50	5.50	5.14
Market Share (x)	.111	.111	.186	.036	.389	.389	1
Cost (cx^2)	.012	.012	.174	.058	.303	.303	.861

Table 1: Social cost comparison for symmetric and asymmetric costs in g.

4.2.1 *Inefficiency Bound*

Let $\vec{\alpha}$ be the unique markup equilibrium for an arbitrary 3-edge-connected market structure. Let the scalar $\bar{\alpha}$ be the upper bound on markups for that market, which is guaranteed to exist when $\bar{w}$ exists. Then:

$$\frac{C(x^{\mathrm{NE}})}{C(x^{\mathrm{OPT}})} \leq \frac{\bar{\alpha}}{2}\,. \tag{11}$$

Here we take $\bar{\alpha}$, the upper bound on producer markups, to represent $\max_a\{\alpha_a\}$, making the inequality as tight as possible. For this bound to be meaningful, we would like to express, or at least bound, $\bar{\alpha}$ as a function of the model primitives. To do this, we introduce the term σ_a, which is an indicator of producer a's market power. We define $\sigma_a := (R_{\ominus a}|_{\vec{w}=\vec{c}})/c_a$. We then define $\sigma := \max_a\{\sigma_a\}$, as a measure of asymmetries in the network as a whole.

From (9), $\alpha_a = 2 + R_{\ominus a}/c_a \leq 2 + \sigma\bar{\alpha}$, This implies that $\bar{\alpha} \leq 2 + \sigma\bar{\alpha}$, from which we deduce that $\bar{\alpha} \leq 2/(1-\sigma)$ for $\sigma < 1$. Plugging into (11) establishes:

$$\frac{C(x^{\mathrm{NE}})}{C(x^{\mathrm{OPT}})} \leq \frac{1}{1-\sigma}\,. \tag{12}$$

We cannot use this bound for $\sigma \geq 1$. Clearly, it is not tight for σ close to 1 as well, as the right hand side blows up approaching 1 from below. If the market is very competitive, then σ will be close to zero, at which point the bound approaches 1. When general series-parallel markets are considered, σ can be typically made large by introducing additional vertical competition though link subdivisions. Some structural restrictions are thus necessary to guarantee any level of efficiency.

We note that slackness is introduced in our bound by the maximum in the definition of σ. This is necessary due to the possible asymmetry of our market structure. When producers compete horizontally in a single market, this analysis can produce a much tighter bound, due to the fact that the degree of competition faced by each of the producers is closely linked in that type of setting (see [7]). Other structural symmetries can be imposed to produce a similar effect.

4.3 Producer Market Position

Finally, we look at the market from the perspective of an individual producer. The producer is concerned with its own profits, and needs to predict the effect that changes in the market structure (initiated on its own or by competitors) will have on this quantity. Here again the story is complicated by the potential asymmetry of the network. This opens the possibility that network alterations will increase the competition for some, while decreasing it for others. We have shown that all markups increase when a subnetwork becomes less efficient. To contrast, profits may decrease for some producers when their competitors become less efficient. The direction of the change in these terms will depend on the relative position of the producers in question to the altered market. At a high level, the nature of the changes will depend on whether the two products are competing in a fashion that is 'more horizontal' or 'more vertical.'

Besides allowing a producer to understand the changes to its competitive environment that result from potential changes to vertical competitors, variance in the responses of both horizontal and vertical competitors combine to produce a large range of potential outcomes in response to any change the producer may make in its production process (we model such changes as a change in the efficiency parameter c_a). A producer considering some form of investment may get a very different projection of the returns when using a localized model of competition. From a producer's perspective, this can be a strong motivation to model vertical competition.

Acknowledgments

This research was partially funded by CIBER at Columbia University and by Conicyt Chile grant FONDECYT 1090050.

5. REFERENCES

[1] D. Acemoglu, K. Bimpikis, and A. Ozdaglar. Price and capacity competition. *Games and Economic Behavior*, 66(1):1–26, 2009.

[2] D. Acemoglu and A. Ozdaglar. Competition and efficiency in congested markets. *Mathematics of Operations Research*, 32(1):1–31, 2007.

[3] D. Acemoglu and A. Ozdaglar. Competition in parallel-serial networks. *IEEE Journal on Selected Areas in Communications*, 25(6):1180–1192, 2007.

[4] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin. *Network Flows: Theory, Algorithms, and Applications*. Prentice-Hall, Englewood Cliffs, NJ, 1993.

[5] U. Akgün. Mergers with supply function. *Journal of Industrial Economics*, 52(4):535–545, 2004.

[6] M. J. Beckmann, C. B. McGuire, and C. B. Winsten. *Studies in the Economics of Transportation*. Yale University Press, New Haven, CT, 1956.

[7] J. R. Correa, N. Figueroa, and N. E. Stier-Moses. Pricing with markups in industries with increasing marginal costs. Columbia Working Paper, 2008.

[8] R. Johari and J. N. Tsitsiklis. Parameterized supply function bidding: equilibrium and welfare. Working paper, 2008.

[9] R. Johari, G. Y. Weintraub, and B. Van Roy. Investment and market structure in industries with congestion. *Operations Research*, 2010. Forthcoming.

[10] P. D. Klemperer and M. Meyer. Supply function equilibria in oligopoly under uncertainty. *Econometrica*, 57(6):1243–1277, 1989.

[11] A. S. Kyle. Informed speculation with imperfect competition. *Review of Economic Studies*, 56(3):317–355, 1989.

[12] C. Papadimitriou and G. Valiant. A new look at selfish routing. *Innovations in Computer Science (ICS)*, 2010.