Math. Program., Ser. A (2014) 146:143–184
DOI 10.1007/s10107-013-0682-8

FULL LENGTH PAPER

Pricing with markups in industries with increasing marginal costs

José R. Correa · Nicolás Figueroa · Roger Lederman · Nicolás E. Stier-Moses

Received: 20 January 2011 / Accepted: 25 April 2013 / Published online: 9 May 2013

Abstract We study a game that models a market in which heterogeneous producers of perfect substitutes make pricing decisions in a first stage, followed by consumers that select a producer that sells at lowest price. As opposed to Cournot or Bertrand competition, producers select prices using a *supply function* that maps prices to production levels. Solutions of this type of models are normally referred to as supply function equilibria. We consider a market where producers' convex costs functions are proportional to each other, depending on the efficiency of each particular producer. We provide necessary and sufficient conditions for the existence of an equilibrium that uses simple supply functions that replicate the cost structure. We then specialize the model to monomial cost functions with exponent $q > 0$, which allows us to reinterpret the simple supply functions as a markup applied to the production cost. We prove that an equilibrium for the markups exists if and only if the number of producers in the market is strictly larger than $1 + q$, and if an equilibrium exists, it is unique. The main

J. R. Correa
Departamento de Ingeniería Industrial, Universidad de Chile, Santiago, Chile
e-mail: jcorrea@dii.uchile.cl

N. Figueroa
Instituto de Economía, Pontificia Universidad Católica de Chile, Santiago, Chile
e-mail: nicolasf@uc.cl

R. Lederman · N. E. Stier-Moses (✉)
Graduate School of Business, Columbia University, New York City, NY, USA
e-mail: stier@gsb.columbia.edu

R. Lederman
e-mail: rlederman13@gsb.columbia.edu

N. E. Stier-Moses
School of Business, Universidad Torcuato Di Tella and CONICET, Buenos Aires, Argentina

result for monomials is that the equilibrium nearly minimizes the total production cost when the market is competitive. The result holds because when there is enough competition, markups are bounded, thus preventing prices to be significantly distorted from costs. Focusing on the case of linear unit-cost functions on the production quantities, we characterize the equilibrium accurately and refine the previous result to establish an almost tight bound on the worst-case inefficiency of equilibria. Finally, we derive explicitly the producers' best response for series-parallel networks with linear unit-cost functions, extending our previous result to more general topologies. We prove that a unique equilibrium exists if and only if the network that captures the market structure is 3-edge-connected. For non-series-parallel markets, we provide an example that does not admit an equilibrium on markups.

Keywords Imperfect competition · Supply function equilibrium · Pricing · Game theory · Allocation efficiency

Mathematics Subject Classification 90B10 · 90B06 · 91B24 · 91B26

1 Introduction

Even though in most markets one can only observe the quantities produced and the prices chosen by firms, as represented by the most traditional competitive models of perfect substitutes like Bertrand's and Cournot's [30], the underlying strategic decisions that firms face need not be this simple. Because market conditions affect prices, firms can be thought as choosing a function that maps how much to charge for different demands. In some cases, these functions take a specific role in the mechanism that clears the market but in other cases, firms make this consideration internally and an outside observer only sees the realized quantities and prices after the market clears. This seminal model where firms consider supply functions (or equivalently price functions) was popularized by the work of Klemperer and Meyer [26] and its outcomes have been referred to as *supply function equilibria* (SFE). Because SFE accommodates both Cournot and Bertrand competition as special cases, SFE allows for price and quantity competition to appear endogenously, as the result of strategic decisions taken by firms.

While Klemperer and Meyer [26] emphasized the importance of these strategies in environments with uncertainty, supply function equilibria are relevant to various production and service industries even when uncertainty is not modeled explicitly. One obvious example is the case of centralized markets, where firms have to submit an actual supply function to a coordinating agency. The agency dynamically adjusts prices to the market-clearing ones, which in turn fixes the quantities transacted by all participants. The most prominent example is given by an electricity market in which generators quote prices contingent on the amount of electricity it will produce. However, many decentralized markets also fit this framework although the interactions among agents may not be defined precisely. For example, in freight transportation, firms can strategically choose a supply function that relates quantity demanded to the price charged for the services. A last example, suggested by Klemperer and Meyer [26] and further discussed by Vives [39] is the consulting industry, where the firm has some flexibility when deciding whether some tasks are part of the project or not. In all these situations, it is particularly important to understand the effect of the strategic behavior of firms.

In this paper we consider an industry with an arbitrary number of asymmetric and strategic firms that produce perfect substitutes of an homogeneous good. Many models of market competition assume constant marginal costs, which does not capture scarce capacity. Instead, we assume that producers have decreasing returns to scale (or increasing marginal costs) and use a similar "technology" although some may be more efficient than others. A typical example of this is electricity generation, where firms use their efficient generators first, and turn on their less efficient ones only when the demand is high enough to deplete the capacity of the more efficient generators. Similarly, freight transportation companies mainly differ in the size of their fleet and in their ability to operate efficiently. Marginal costs are increasing because new ships are expensive, and to go over capacity they need to lease additional space incurring in larger marginal costs. To elaborate further, even if demand is deterministic, a mechanism based on supply functions has advantages over classic mechanisms such as Bertrand or Cournot competition. In those classic mechanisms, each firm has to submit a scalar which is optimal only if competitors also behave exactly as specified by the equilibrium of the game. Instead, a supply function allows the firm to exploit others' deviations from the equilibrium prediction to get to a better outcome for them, and a more robust solution for the mechanism.

In our model, firms make pricing decisions forecasting the demand they will face under each combination of supply functions offered by the different producers. In a second phase, an assignment of producers to consumers is chosen by a central planner, as it is the case in electricity markets. Alternatively, in a decentralized market, consumers learn the price functions chosen during the first phase and converge to an equilibrium in which they select producers selling at lowest price (we assume that consumers are small enough so they act as price takers). As it will become clear, both situations are equivalent under our modeling assumptions.

The demand is deterministic, inelastic and publicly known. Because of the assumption that the heterogenous firms face cost functions with a similar structure, we consider per-unit cost functions $u_a(x_a) := c_a u(x_a)$, parameterized with a single number c_a. Here, x_a is the production quantity of the firm and u is a function that describes the cost structure of the industry.

We start by analyzing the existence of equilibria in a game where producers are constrained to choose a supply function from a family parameterized by one parameter. The supply functions in this family, which we refer to as "simple", replicate the firms' production costs. Indeed, firms choose a parameter β_a and bid the supply function $S_a(p) := \beta_a u^{-1}(p)$. We prove that an equilibrium has to exist if enough producers participate in the market, where the threshold depends only on the function u. This equilibrium is monotonic in the sense that more efficient firms bid higher production quantities (for each given price) and capture a bigger market share. When the number of firms is too low, an equilibrium fails to exist because the amount of competition is not enough to curb sale prices and prevent firms from overcharging. In this situation, a best response to the prices of other firms is to charge more, thus preventing the existence of a fixed point. In the case of linear marginal costs and symmetric firms, Baldick and Hogan [9] prove that the only stable equilibria are achieved with simple supply functions. We provide further support for simple supply functions by showing that, under our assumptions, this equilibrium is immune to arbitrary deviations: firms cannot increase their profits by deviating and choosing an arbitrary increasing supply function (not necessarily simple). Therefore, an equilibrium supported by simple supply functions is also one for the larger strategy-space of increasing functions. From a practical point of view, in electricity markets, the regulator has a good estimate of the cost structure of generators and therefore can impose that the supply functions generators submit mimic the cost pattern. This can be used, for instance, to disallow *hockey-stick* supply functions that may lead to inefficient outcomes [37].[1] On the other hand, in decentralized markets, an equilibrium where producers adopt price functions that imitate the shape of their production costs is justified by the widely used practice of setting prices by applying a fixed markup to the cost (normally referred to as 'cost-plus pricing').

It is worth mentioning that with deterministic demand there is a very large number of equilibria. Technically speaking, any set of supply functions for the producers with the *correct* value and derivative at the market clearing price will constitute an equilibrium. This contrasts with the result of Klemperer and Meyer [26] who proved that uncertain demand dramatically reduces multiplicity of equilibria.

After proving existence, we offer an explicit characterization of the equilibrium, which allows us to compute it and to study its properties. For this, we focus on the case of monomial cost functions whereby $u(x) = x^q$ for a fixed $q > 0$. In this case, a simple supply function can be reinterpreted as a price function that includes a markup applied to the production cost, and the outcome of the game can be seen as a *markup equilibrium*. Multi-stage games like the one we analyze in this paper frequently become intractable when general cost functions are used (see, e.g., [2,17,25,41,43]). Monomials are tractable and general enough so one can still make a good first-order approximation by fitting the value and derivative of any function close to the equilibrium situation. Within this class, linearly-increasing marginal costs are particularly interesting because they are relevant to practice. For instance, Baldick e al. [8] provide a detailed explanation about why linear cost and price functions in electricity markets provide accurate results. We prove that a markup equilibrium exists for general $q > 0$ if and only if the number of competitors is strictly larger than $1 + q$. Moreover, whenever an equilibrium exists, it is unique. Note that if the decreasing returns to scale are steeper, more firms are needed for an equilibrium to exist. As marginal costs rise faster, firms can be less aggressive in trying to obtain large market shares, so more of them are needed to ensure that there is a best response with bounded prices.

We then study the welfare implications of imperfectly-competitive markets satisfying our assumptions. An equilibrium is not necessarily efficient, meaning that it need not minimize the total production cost because of the presence of negative external-

[1] This is the case, among others, of the Chilean system, which operates with "audited costs" (a particular case of what is known in the literature as cost-based bids). In this system, the central dispatcher may audit firms that submit non-credible supply functions. Basically, the regulator knows the shape of the cost function (because he knows the technology used) but not the exact values. This could be the case, for example, because the regulator does not know the private contracts between the firm and its suppliers (coal, fuel, etc.).

ities. It is natural to study the extent of the inefficiency and how it depends on the market power of producers. At equilibrium, supply functions could be overstated and distortion in prices could lead to oversupply or undersupply to particular producers with respect to the optimal assignment, thus inflating the total production cost for the economy. To quantify the inefficiency, we consider the worst possible ratio between the total production cost at equilibrium and that of an efficient allocation, taken over all possible instances to the problem. This ratio has been referred to as the *Price of Anarchy* [27], and has been studied in several games relevant to operations research, operations management, computer science and economics. When the mentioned ratio is small, a planner can be sure that, independently of the details of the market structure, there is no big loss in welfare due to market power.

Our main conclusion is that the distortion in prices created by firms acting strategically create inefficiencies, but in the context of inelastic supply and an arbitrary number of firms with monomial cost functions, this impact is limited. Actually, we compute the inefficiency of equilibria parameterized by the *competitiveness* of the market, measured by the norm of a vector whose components are $\min_i c_i/c_a$, i.e., costs normalized by the most efficient firm. This norm satisfies the properties expected of a measure of competitiveness: it increases if a new firm enters the market, it increases more if the new firm has lower costs, it increases if any incumbent which is not a leader reduces its costs and this effect is bigger if the reduction happens to a more competitive firm. Finally, it decreases if the leader decreases its marginal cost, since in that case the gap between the leader and the follower decreases. Moreover, in the mergers' literature (see, e.g., [5,31,34]), marginal costs are assumed to be $c_a = 1/K_a$ for a capital stock of firm a's equal to K_a. Then, a merger between two firms, which just combines their capital stock, always decreases competitiveness. Our measure corresponds to the aggregation of the capital of all firms, representing the market as a whole.

We provide an upper bound to the efficiency-loss for any competition level higher than $(1+q)^q$, which is a sufficient condition for the existence of an equilibrium. Although with very low competition equilibria can be arbitrarily inefficient, the price of anarchy is bounded by a small constant if competition is competition is relatively high. Furthermore, an equilibrium assignment is nearly efficient as the competitiveness of the market tends to infinity. Note that, in practice, it is likely that an industry is competitive whenever entry costs are small, since a non-competitive industry with high profits will induce entry. A basic idea behind these results is to show that although the most efficient producers are more profitable than less efficient ones (since the market structure supports larger markups for them), when there is enough competition, markups are bounded and cannot be infinitely large. For the case of linearly increasing marginal costs, we establish a bound on the price of anarchy with an error of at most 0.3 %. Evaluating the bound numerically, we get that the production cost at equilibrium is at most 50 % worse than the optimal one for reasonable values of competitiveness (i.e., the price of anarchy is 3/2). The worst-case gap between the two assignments decreases rapidly as competition increases. For instance, the inefficiency is already below 6.2 % when the competition level equals 3. On the other hand, using this bound, we construct (asymptotically) worst-case instances. These results come as a subproduct of a procedure that we design to find an equilibrium in the linear case. One of the main ideas behind it is to observe that we can normalize any instance so the equilibrium equations become significantly simpler. Although this normalization does not lead to a closed-form solution for the equilibrium, it does provide an efficient procedure to compute one. With this simplification we can write the price of anarchy of all instances with linear costs explicitly as a nonconvex program. This can be reduced further to a nonconvex integer programming problem that just has seven variables and a very small integrality gap.

It is important to note that the ratio between profits that firms experience at equilibrium and those that would be achieved if producers were *non-strategic* can be much larger than the ratio between the corresponding social costs. In that case, there is surplus extraction by the producers from the consumers, but since demand is inelastic, there is little efficiency loss. Indeed, for the linear case, we prove that when the competitiveness of an instance tends to 2, the ratio of the social cost at equilibrium to that of the social optimum remains bounded by 3/2 although markups and profits may grow to infinity. This is of course is caused by our assumption of inelastic demand.

One of the limitations of our model, particularly relevant in the transportation industry, is that so far we only considered producers that provide substitute goods. To address this we extend our model to allow competition among producers that supply complements, in addition to substitutes, leading to a network that captures the market structure (applications with this structure are discussed further in Sect. 6). In this extension, producers facing linear marginal costs compete to provide all or some portion of the product to customers, who choose a set of producers offering the lowest combined price (path in the network). For arbitrary series-parallel networks (in graph theoretic terms), we derive explicitly the producers' best response, using a network transformation. This allows us to fully characterize the markups chosen by producers at equilibrium. We prove that a unique equilibrium exists if and only if the network that captures the market structure is 3-edge-connected, extending our earlier result for substitute goods. For non-series-parallel markets, we provide an example that does not admit an equilibrium on markups. We acknowledge these results only apply to the case of linear marginal cost. The main difficulty in extending this to more general cost functions relies in solving the second stage game in closed form (which is not hard in the substitutes case).

The rest of the paper is structured as follows. We now discuss the relations to the literature. Section 2 introduces the supply function equilibrium model. In Sect. 3 we look at the case of general cost functions, while in Sect. 4, we concentrate on monomial cost functions. The analysis of the efficiency-loss incurred by solutions at equilibrium is done in Sect. 5. We generalize our model so it can handle complements besides substitutes in Sect. 6. Finally, Sect. 7 concludes by presenting directions for future research.

1.1 Related literature

The literature of supply function equilibria goes back to the conjectural variations model of Bowley [11], re-emerging with papers by Grossman [20], Robson [35], Wilson [42], and Turnbull [38]. This early literature on supply function equilibria focus mainly on divisible goods auctions and assumes that demand is deterministic. Robson [35] and Turnbull [38], and the influential paper by Klemperer and Meyer [26] incorporate uncertain demand. Although demand uncertainty is an important motivation of supply function equilibrium models, this paper—like those mentioned earlier, and like others such as Yang and Hajek [44] and Johari and Tsitsiklis [24]—assumes that demand is deterministic.

Klemperer and Meyer [26] consider a supply function equilibrium model with uncertain demand and prove that the infinitely-many equilibria that exist in the space of supply functions with deterministic demand collapse into a single one inside the support of the uncertain demand. Furthermore, they prove that the equilibrium has the same structure as costs functions.

Rudkevich et al. [36] and Anderson and Philpott [7] study a supply function equilibrium model with uncertain but inelastic demand and find the explicit form of an equilibrium for symmetric firms. The former considers piecewise linear convex costs functions, while the latter generalizes the assumptions to arbitrary convex ones. Baldick et al. [8] and Anderson and Hu [6] consider the case of asymmetric firms, and study procedures to find supply function equilibria. The asymmetry and computational focus are features in common with our work, although they also consider stochastic demand and capacity constraints.

Related to our definition of competitiveness, Akgün [5] models a merger as the appearance of a new firm with a reduced cost function, since the new firm can avoid more easily the decreasing returns to scale by allocating production efficiently among different plants. His model specification considers linear unit cost functions and elastic demand, and in such a context, he finds that equilibria always exist and that mergers decrease total welfare but increase the profits of merging firms. Hendricks and McAfee [21] consider a more general framework where producers *and* consumers have market power and submit supply/demand functions. They develop a new measure of concentration, which can be related to equilibrium markups, profits and market shares. Acemoglu et al. [1], Acemoglu and Ozdaglar [2], and Johari et al. [25] also make the pricing decisions endogenous in the game and consider other elements like supply quantities, investment, or entry. Although they do not consider supply functions, they study a question similar to ours in spirit and provide bounds for the efficiency loss in their games.

Closest to our work is the paper by Johari and Tsitsiklis [24], who also consider supply function equilibria from the perspective of studying the worst-case inefficiency at equilibrium. They consider a model in which producers, also facing a unit inelastic demand, are restricted by a mechanism designer to choose a parameter w and submit a supply function of the form $S(p) = 1 - w/p$. By restricting the strategy space, they are able to prove existence of equilibrium for any market with more than three firms and very general cost functions, and to prove that both welfare and profits are close to those of the efficient assignment. In contrast, we do not restrict a firm's strategy space, but impose more restrictions on the cost functions. Allowing firms to use general supply functions is in line with our assumption of a decentralized market, since in that case a predefined form for the supply functions is hard to enforce. For the unrestricted space of supply functions, it is harder to prove the existence of equilibrium since there are more possible deviations. Nevertheless, we prove the existence of an equilibrium where each firm submits a supply function that mimics its cost structure. Another significant difference is that if the strategy space is not restricted a priori, firms' profits can be arbitrarily larger than in the socially efficient allocation. Allowing for arbitrary supply functions effectively allows more market power, but the welfare loss associated to this (which comes from the distortion in market shares, since the demand is inelastic) remains bounded. In a regulated market, where a mechanism designer can choose the "rules of the game," imposing a parameterized family of supply functions as a strategy space is reasonable, and the work of Johari and Tsitsiklis [24] shows that this can curb the firm's market power. In decentralized markets, such a restriction may be unfeasible, and our work shows that anyhow the welfare loss associated to firms acting strategically is bounded.

Finally, our work with series-parallel network structures relates to prior literature involving complementarities. A classic result of Cournot [15] states that price-setting monopolists split profits equally among them. Some recent work achieves more general outcomes through the use of network structures. For example, the model of Casadesus-Masanell et al. [12] expands on Cournot's original model by allowing competition and vertical differentiation in one of the markets. Closer to our setup is the literature on decentralized assembly systems (see, e.g., [19,40], and the references therein), in which an assembler purchases a set of components from multiple strategic suppliers. In particular, Jiang and Wang [22] avoid profit symmetry by allowing competition within individual component markets. In their model, competition is Bertrand, as suppliers compete by fixing their wholesale price. Constant marginal costs ensure that a single firm produces each component. Lastly, there is a growing literature on competition in networks that is focused on the role of prices in guiding users towards efficient selection of paths through the network. The literature on centralized pricing strategies has recently expanded to include work on price competition by decentralized firms. See, e.g., Acemoglu Ozdaglar [3], Chawla and Roughgarden [14], and Papadimitriou and Valiant [33] for market structures involving complementarities.

2 The model

We consider a market in which producers in $A = \{1, \dots, n\}$ sell identical goods. The per-unit production cost for each producer $a \in A$ is an increasing and differentiable function $u_a : \mathbb{R}_+ \to \mathbb{R}_+$ that depends on the production quantity $x_a \geqslant 0$. We assume that all producers make use of similar 'technology' but some are more efficient than others. This is modeled by letting the cost function be equal to $u_a(x_a) := c_a u(x_a)$ 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 $a \in A$. Without loss of generality, we order producers such that $c_1 \leqslant \dots \leqslant c_n$.

Since we consider the case of industries with increasing marginal production costs—which is the case, e.g., when labor or production capacity is scarce or when there is congestion—we also assume that $xu(x)$ is convex. Furthermore, we consider a market coverage condition to simplify the characterization of equilibria in the second stage of the game, which prompts us to assume that u is bijective (i.e., evaluates to zero at zero and grows to infinity). The restriction that the cost at zero is zero is technical but not very restrictive because companies face positive demands at equilibrium; it amounts to considering only big-enough firms (modeling entry and endogenous participation would add another layer of complexity that would easily make the model intractable). Note that as u is bijective, its inverse u^{-1} is well defined. To guarantee that producers' optimal welfare is achieved when the derivative of their utility vanishes, we make the additional technical assumption that $pu^{-1}(p)$ is convex. This is equivalent to the following two conditions: $u(x)u''(x) < 2(u'(x))^2$ and u^2/u' is increasing. Some examples of unit cost functions that satisfy all assumptions are polynomials with nonnegative coefficients, the exponential function $e^x - 1$, and the logarithmic function $\ln(1 + x)$. Putting all the elements together, the total cost of producing x_a units of the good is $\kappa_a(x_a) := x_a u_a(x_a) = c_a x_a u(x_a)$, which is convex.

The game we consider consists of two stages. The first stage is a pricing game among the producers and the second stage is an assignment game among the consumers, who decide from whom to buy. In the first stage of the game, producers select a *supply function* $S_a(p)$ which maps the quantity they are willing to produce to the corresponding unit price and inform consumers of their supply function. Equivalently, producers could consider a *price function* $p_a(x) = S_a^{-1}(x)$ because, being the inverse of the supply function, it provides the same information. If producer a receives a total order of x_a units of the good from the consumers, each unit will be sold at price $S_a^{-1}(x_a)$. Our only assumption is that producers are limited to choose supply functions that are increasing and concave. The supply function is chosen to balance the tradeoff between high per-unit revenue and low demand, or vice-versa. The goal of the producer is to maximize its profit, which equals $x_a(S_a^{-1}(x_a) - c_a u(x_a)) \geqslant 0$.

In the second stage, consumers select their suppliers. We assume that there are infinitely many consumers that require an aggregated demand of one unit. The assumption of a unit demand is just for simplicity; the structure of cost functions makes the choice of total demand irrelevant. Furthermore, we assume that each consumer is small compared to the market—implying that all of them act as price takers—and that the demand is inelastic, although both assumptions can be relaxed. Consumers satisfy their demands with producers that sell at minimal price, taking the supply functions of each firm as given. Throughout the paper we represent the aggregate consumption decisions by the vector $x \in \mathbb{R}^A$, which can be viewed as the market shares of the producers.

2.1 Nash equilibria

A *supply function equilibrium* of the producers' game is a vector of supply functions $(S_a)_{a \in A}$ that satisfies the Nash equilibrium condition: no producer can increase the profit by switching to another supply function when the rest of the supply functions are fixed. An equilibrium in the consumers' game is an assignment x^{NE} such that all consumers are buying at minimal price. These two games are played sequentially, making it a Stackelberg game in which producers are the leaders and consumers are the followers.

The second stage is simply a market-clearing game, in which the quantity x_a that producer a sells equals $S_a(p^*)$, where p^* is the market-clearing price. Since the total demand equals one, the market-clearing price p^* is the unique solution to the equation

$$\sum_{a \in A} S_a(p^*) = 1.$$

Anticipating the market-clearing process, producers choose supply functions to maximize their profits, which can be written as a function of p^*:

$$x_a^{\mathrm{NE}}(S_a^{-1}(x_a^{\mathrm{NE}}) - c_a u(x_a^{\mathrm{NE}})) = S_a(p^*)(p^* - c_a u(S_a(p^*))).$$

Interestingly, the previous equation implies that an equilibrium is completely determined by the choice of supply functions. Indeed, the vector of functions determines a unique market-clearing price p^*, which in turn determines unique market shares $x_a^{\mathrm{NE}} = S_a(p^*)$ for the producers. We thus have the following definition.

Definition 1 A vector of supply functions $(\bar{S}_a)_{a \in A}$ is a Nash equilibrium for the producers' game if and only if

$$\bar{S}_a(\bar{p})(\bar{p} - c_a u(\bar{S}_a(\bar{p}))) \geqslant S_a(p)(p - c_a u(S_a(p))), \tag{2.1}$$

for all $a \in A$ and for all increasing and concave supply functions $S_a(\cdot)$. Here, the market-clearing price at equilibrium $\bar{p}$ satisfies $\sum_{a \in A} \bar{S}_a(\bar{p}) = 1$, while the price p under the alternative strategy $S_a(\cdot)$ is the unique solution to $S_a(p) + \sum_{i \neq a} \bar{S}_i(p) = 1$.

2.2 Optimal assignment

To quantify the quality of an assignment, we let the total production cost $C(x) := \sum_{a \in A} \kappa_a(x_a) = \sum_{a \in A} c_a x_a u(x_a)$ be our social cost function. This function captures whether consumers are matched to the most efficient producers. Payments are not considered in this function because they are internal transfers that do not affect the welfare of the system. The socially-optimal assignment x^{OPT} is the unique minimizer of $C(x)$ given by

$$x^{\mathrm{OPT}} := \arg\min \left\{ C(x) : \sum_{a \in A} x_a = 1,\ x_a \geqslant 0 \right\}.$$

It is worth observing that an optimal assignment is achieved if producers charge their marginal cost. Indeed, producer a charges its marginal cost when its supply function is the inverse of $\kappa_a'(x_a)$. Indeed, $S_a(p) = (\kappa_a')^{-1}(p)$ leads to a market-clearing price p^* satisfying $\sum_{a \in A}(\kappa_a')^{-1}(p^*) = 1$. A simple calculation shows that the optimal assignment $x_a^{\mathrm{OPT}} = (\kappa_a')^{-1}(p^*)$. However, the distortion of costs introduced by supply functions at equilibrium can lead to an assignment that does not necessarily minimize $C(x)$ because of the existence of negative externalities. One of our goals is to find conditions under which the distortions of costs and the increase of the total cost at equilibrium are not too large.

3 Supply function equilibria when marginal costs are increasing

We now turn into characterizing the equilibria of the game played among producers. In this section we consider general cost functions that satisfy the assumptions described in the previous section. We first focus on simple supply functions that replicate the cost structure of producers. Restricting the search to this type of functions allows producers to greatly simplify the problem of finding an optimal supply function since they have to consider a single degree of freedom, as opposed to potentially searching in a space of infinite dimensions.[2] We prove that an equilibrium exists as long as there are enough producers in the market and an equilibrium does not exist if there are too few producers in the market. Next, we justify this choice of supply functions by proving that the equilibria we find are still at equilibrium when producers can pick their functions from the set of increasing functions. The main implication of this result is that these equilibria, supported by simple supply functions, are robust.

3.1 Equilibria with simple supply functions

In this section, we assume that producers restrict their consideration to simple supply functions of the form $S_a(p) = \beta_a u^{-1}(p)$ for a parameter $\beta_a \geqslant 0$ chosen by them. Notice that these functions charge prices that replicate the cost structure of the industry. Indeed, the corresponding price function (i.e., the inverse of the supply function) is $p_a(x_a) = u(x_a/\beta_a)$ so producers select an amplification factor for the demand, and charge the industry's cost evaluated at this amplified demand. We characterize best responses on the space of β_a, and establish sufficient conditions for an equilibrium to exist and for it to not exist.

Under the assumption that all producers bid simple supply functions, let us consider arbitrary but fixed parameters $\beta_i > 0$ for producers $i \neq a$. First, we compute the best response β_a of producer a as a function on the values of β_i for the other producers $i \neq a$.

The market-clearing price p satisfies demand, that is $\sum_{i\in A} \beta_i u^{-1}(p) = 1$, from where $p = u(1/\sum_{i\in A} \beta_i)$. Replacing this price into the profit function shown in (2.1), the profit of producer a equals

$$P_a(\beta_a) := \frac{\beta_a}{\sum_{i\in A} \beta_i}\left(u\left(\frac{1}{\sum_{i\in A} \beta_i}\right) - c_a u\left(\frac{\beta_a}{\sum_{i\in A} \beta_i}\right)\right). \tag{3.1}$$

The next proposition optimizes the profit function over β_a, and establishes that although it is possible that $P_a(\beta_a)$ is not concave, there is a unique solution to the maximization problem.

Proposition 3.1 *There is a unique solution to the problem* $\max\{P_a(\beta_a) : \beta_a \geqslant 0\}$, *which is achieved where* P_a' *vanishes.*

[2] In fact, Klemperer and Meyer [26] already noted that in the case of deterministic demand an equilibrium is supported by supply functions that have the right value and slope at the right price.

Proof Note that P_a is continuous and differentiable, and satisfies that $P_a(0) = 0$ and $P_a(\beta_a) < 0$ for $\beta_a \to \infty$. Therefore, $P_a(\beta_a)$ is maximized in $(0, \infty)$ at a point where the derivative vanishes. To simplify notation, we make $B_{-a} := \sum_{i \neq a} \beta_i$. Using the change of variable $p := u(1/(\beta_a + B_{-a}))$, the producer a can be viewed as choosing the market-clearing price that corresponds to a value of β_a. Rewriting the profit of producer a as a function of this price p, we get that

$$P_a(p) := (1 - B_{-a}u^{-1}(p))(p - c_a u(1 - B_{-a}u^{-1}(p))) = \omega p - c_a g(\omega), \tag{3.2}$$

where we used $g(x) := xu(x)$ and $\omega = 1 - B_{-a}u^{-1}(p)$ to simplify notation. The derivative of the previous expression with respect to p is $\omega + B_{-a}(u^{-1})'(p)(c_a g'(\omega) - p)$. Using the relation $(u^{-1})'(p) = 1/u'(u^{-1}(p))$, a necessary condition for a price p to be a best response is that

$$p + u'(u^{-1}(p))(u^{-1}(p) - 1/B_{-a}) = c_a g'(1 - B_{-a}u^{-1}(p)). \tag{3.3}$$

Notice that the right-hand side is positive and decreasing. We are going to prove that the left-hand side is increasing whenever it is positive, which will imply the existence of a unique best response. Indeed, using the monotone change of variable $y = u^{-1}(p)$, the left-hand side as a function of y is $h(y) := u(y) + u'(y)(y - 1/B_{-a})$. Taking the derivative,

$$h'(y) = 2u'(y) + u''(y)\left(y - \frac{1}{B_{-a}}\right) > \frac{2u'(y)}{u(y)}\left(u(y) + u'(y)\left(y - \frac{1}{B_{-a}}\right)\right)$$

where the inequality follows from $y - 1/B_{-a} \leqslant 0$ and $u(x)u''(x) < 2(u'(x))^2$. In conclusion, $h(y) \geqslant 0$ implies that $h'(y) \geqslant 0$, proving the claim. Transforming back to the original variables, we have that the best response for producer a equals $\beta_a = 1/u^{-1}(p^*) - B_{-a}$. □

An equilibrium in the space of simple supply functions is a vector $(\beta_a)_{a \in A}$ in which each producer plays a best response to the others' actions. Interestingly, we obtain a natural monotonicity property; namely, more efficient producers profit more at equilibrium and capture a larger market share. This is summarized in the following corollary.

Corollary 3.2 *If an equilibrium exists and producers i and j are such that $c_i < c_j$, then $\beta_i > \beta_j$ and $x_i^{\mathrm{NE}} > x_j^{\mathrm{NE}}$.*

Proof The proof of Proposition 3.1 implies that an equilibrium simultaneously satisfies (3.3) for all $a \in A$; hence, $c_a g'(1 - B_{-a}u^{-1}(p^*)) + u'(u^{-1}(p^*))/B_{-a}$ is constant across firms, where p^* is the market clearing price at equilibrium. Since this expression is decreasing with respect to B_{-a}, $c_i < c_j$ implies that $B_{-i} < B_{-j}$. From here we conclude that $\beta_i > \beta_j$ and the claim follows because $x_a^{\mathrm{NE}} = \beta_a u^{-1}(p^*)$. □

The two following results provide a sufficient and a necessary condition on the existence of equilibria. We establish that an equilibrium exists when the number of producers is larger than the maximum ratio of the marginal production cost to the unit production cost. Afterwards, we prove that equilibria cannot exist when the number of producers is smaller than the minimum of that ratio. Although these two conditions are not complementary, for the case of monomials that we consider in Sect. 4, the results provide a complete characterization of the existence of equilibria in the space of simple supply functions.

Theorem 3.3 *Assume the number of producers n participating in the market is strictly larger than* $\tilde{n} := \max_{x\geqslant 0}(xu'(x) + u(x))/u(x)$. *Then, the producers' game has an equilibrium where producers bid in the space of simple supply functions.*

Proof Let $\Gamma(\beta) = (\Gamma_1(\beta_{-1}), \ldots, \Gamma_n(\beta_{-n}))$ be the best response mapping in the space of simple supply functions, where $\Gamma_a(\beta_{-a})$ maps the supply functions of others to the best response of producer a. To prove that an equilibrium exists, we show that Γ maps a compact set into itself and use Brouwer's fixed point theorem. In particular, we prove that best responses are bounded.

Similarly to Proposition 3.1, consider $B_{-a} := \sum_{i\neq a}\beta_i$, the change of variable $p = u(1/(B_{-a} + \beta_a))$, and the profit function $P_a(p) = (1 - B_{-a}u^{-1}(p))(p - c_a u(1 - B_{-a}u^{-1}(p)))$. Using the definition of $\tilde{n}$ and doing similar calculations as before, a critical point p satisfies

$$p+u'(u^{-1}(p))(u^{-1}(p) - 1/B_{-a})=c_a g'(1-B_{-a}u^{-1}(p))\leqslant c_a\tilde{n}u(1-B_{-a}u^{-1}(p)).$$

Going back to variable β_a, and using that $\tilde{n}-1 \geqslant xu'(x)/u(x)$ applied to $1/(B_{-a}+\beta_a)$, we get that

$$\frac{u(\beta_a/(B_{-a} + \beta_a))}{u(1/(B_{-a} + \beta_a))} \geqslant \frac{B_{-a} - (\tilde{n} - 1)\beta_a}{B_{-a}c_a\tilde{n}}.$$

Furthermore, since the profit is nonnegative at equilibrium, $u\,(1/(B_{-a} + \beta_a)) - c_a u\,(\beta_a/(B_{-a} + \beta_a)) \geqslant 0$. Putting the two bounds together we conclude that

$$\frac{B_{-a} - (\tilde{n} - 1)\beta_a}{B_{-a}c_a\tilde{n}} \leqslant \frac{u(\beta_a/(B_{-a} + \beta_a))}{u(1/(B_{-a} + \beta_a))} \leqslant \frac{1}{c_a}. \tag{3.4}$$

We use (3.4) to prove that best responses are bounded. Let $\bar{c} = \max_{i\in A} c_i$ and $\underline{c} = \min_{i\in A} c_i$. Let us consider that for all $i \neq a$ the parameter β_i is bounded by

$$0 < \varepsilon \leqslant \beta_i \leqslant M := \frac{1}{u^{-1}(\underline{c}u(1/n))} < \infty\,,$$

where $\varepsilon > 0$ will be determined later. We have to prove that the best response β_a is bounded by the same constants. Let us first see the upper bound. From the second inequality in (3.4), the assumption $\beta_i \leqslant M$ for all $i \neq a$, and assuming that $\beta_a > M$, we have that

$$\frac{1}{\underline{c}} \geqslant \frac{1}{c_a} \geqslant \frac{u(\beta_a/(B_{-a} + \beta_a))}{u(1/(B_{-a} + \beta_a))} \geqslant \frac{u(\beta_a/((n-1)M + \beta_a))}{u(1/(B_{-a} + \beta_a))} \geqslant \frac{u(1/n)}{u(1/\beta_a)}.$$

This inequality says that $\beta_a \leqslant 1/u^{-1}(\underline{c}u(1/n)) = M$ which is a contradiction. To prove the lower bound, we first let a be such that β_a is the smallest and assume that $\beta_a < \varepsilon$, then using $B_a \geqslant (n-1)\varepsilon$ we can write

$$\frac{1}{\bar{c}\tilde{n}} - \frac{(\tilde{n}-1)\beta_a}{(n-1)\varepsilon\bar{c}\tilde{n}} \leqslant \frac{u(\beta_a/(B_{-a}+\beta_a))}{u(1/(B_{-a}+\beta_a))} \leqslant \frac{u(\varepsilon/(B_{-a}+\varepsilon))}{u(1/(B_{-a}+\varepsilon))}. \tag{3.5}$$

The remainder of the analysis is divided into two cases: when $B_{-a} > (n-1)\sqrt{\varepsilon}$ and when $B_{-a} \leqslant (n-1)\sqrt{\varepsilon}$. In the former case we have

$$\frac{u(\varepsilon/(B_{-a}+\varepsilon))}{u(1/(B_{-a}+\varepsilon))} \leqslant \frac{u(\varepsilon/((n-1)\sqrt{\varepsilon}+\varepsilon))}{u(1/((n-1)M+\varepsilon))} \leqslant \frac{u(\sqrt{\varepsilon}/(n-1))}{u(1/(nM))},$$

while in the latter

$$\frac{u(\varepsilon/(B_{-a}+\varepsilon))}{u(1/(B_{-a}+\varepsilon))} \leqslant \frac{u(1/n)}{u(1/(n\sqrt{\varepsilon}))}.$$

In both cases, since $u(0) = 0$ and $\lim_{x\to\infty} u(x) = \infty$ for $\varepsilon > 0$ small enough we have that (and this is how ε is defined)

$$\frac{u(\varepsilon/(B_{-a}+\varepsilon))}{u(1/(B_{-a}+\varepsilon))} \leqslant \frac{n-\tilde{n}}{\bar{c}\tilde{n}(n-1)}.$$

Putting this inequality back into (3.5) we obtain

$$\frac{1}{\bar{c}\tilde{n}} - \frac{(\tilde{n}-1)\beta_a}{(n-1)\varepsilon\bar{c}\tilde{n}} \leqslant \frac{n-\tilde{n}}{\bar{c}\tilde{n}(n-1)},$$

which is equivalent to $\beta_a \geqslant \varepsilon$, a contradiction.

Thus, we have proved that Γ is a continuous function that maps a compact set into itself. Brower's fixed point theorem implies that it has a fixed point, which is a Nash equilibrium. □

The following proposition looks at the opposite case and proves that an equilibrium does not exists if there are too few producers competing in the market.

Proposition 3.4 *If* $n \leqslant \min_{x>0}(xu'(x)+u(x))/u(x)$, *then the producers' game does not have an equilibrium where producers bid simple supply functions.*

Proof Differentiating (3.1), the best response condition $P_a'(\beta_a) = 0$ can be written as

$$\begin{aligned}\left(u\Big(\frac{1}{\sum_{i\in A}\beta_i}\Big) - c_a u\Big(\frac{\beta_a}{\sum_{i\in A}\beta_i}\Big)\right)\sum_{i\neq a}\beta_i \\ = \frac{\beta_a}{\sum_{i\in A}\beta_i}\left(u'\Big(\frac{1}{\sum_{i\in A}\beta_i}\Big) + c_a u'\Big(\frac{\beta_a}{\sum_{i\in A}\beta_i}\Big)\sum_{i\neq a}\beta_i\right).\end{aligned} \tag{3.6}$$

Using that $c_a > 0$, the previous equation implies that

$$u\left(\frac{1}{\sum_{i\in A}\beta_i}\right)\sum_{i\neq a}\beta_i > \frac{\beta_a}{\sum_{i\in A}\beta_i}u'\left(\frac{1}{\sum_{i\in A}\beta_i}\right),$$

which added together for all $a \in A$ leads to

$$(n-1)u\left(\frac{1}{\sum_{i\in A}\beta_i}\right)\sum_{a\in A}\beta_a > u'\left(\frac{1}{\sum_{i\in A}\beta_i}\right).$$

This contradicts the assumption that $xu'(x) + u(x) \geqslant nu(x)$ for $x = 1/\sum_{i\in A}\beta_i$. □

3.2 Robustness of equilibria

Having established conditions under which a simple equilibrium exists, we now show that this equilibrium concept is robust and extends to non-simple supply functions. The next theorem proves that for the equilibrium characterized in Sect. 3.1, it is in the best interest of producers to maintain their choices of simple supply functions, even when they are allowed to bid an arbitrary supply function.

Note that Klemperer and Meyer [26] gave a similar result for the special case of linear cost functions and two homogeneous producers but in the more general case of stochastic demand. They proved that among the many equilibria, there exists one in which producers bid linear supply functions (Actually, when the uncertainty has infinite support, the linear one is the unique equilibrium). Our result has the same flavor: it indicates that there is an equilibrium in which both cost and supply functions have the same shape. While it works in the setting of deterministic demand, it considers multiple and heterogenous producers with arbitrary cost functions that satisfy the assumptions of Sect. 2.

Theorem 3.5 *Consider an equilibrium in the space of simple supply functions where each producer $i \in A$ bids a simple supply function of the form $S_i(p) = \beta_i u^{-1}(p)$ for a $\beta_i > 0$ of their choice. For an arbitrary producer $a \in A$, S_a is also a best response function if producer a chooses supply functions from a strategy space consisting of all nondecreasing functions.*

Proof Let us consider that supply functions are fixed for producers $i \neq a$ and focus on producer $a \in A$. When computing a best response, producer a solves

$$\max_{S_a(\cdot)}\left\{S_a(p)(p - c_a u(S_a(p))) : \sum_{i\in A}S_i(p) = 1\right\}. \tag{3.7}$$

Since supply functions of others are fixed, the problem of producer a is equivalent to choosing the market-clearing price $p^* \in [0, \bar{p}_{-a}]$ at equilibrium, where $\bar{p}_{-a}$ is the market-clearing price when producer a does not participate, i.e., $\sum_{i\neq a}S_i(\bar{p}_{-a}) = 1$. Indeed, any p in that interval can be achieved, and given the supply functions of the other producers, the market share at the market-clearing price chosen by producer a is determined by $S_a(p^*) = 1 - \sum_{i \neq a} S_i(p^*)$. Proceeding as we did in Proposition 3.1, (3.7) is equivalent to

$$\max_p \left\{ \Big(1 - \sum_{i \neq a} S_i(p)\Big)\Big(p - c_a u\big(1 - \sum_{i \neq a} S_i(p)\big)\Big) : p \in [0, \bar{p}_{-a}] \right\}. \tag{3.8}$$

Denoting the objective function by $H_a(p)$, we compute

$$\begin{aligned} H_a'(p) = 1 + c_a & \left(\left(1 - \sum_{i \neq a} S_i(p) \right) u' \left(1 - \sum_{i \neq a} S_i(p) \right) + u \left(1 - \sum_{i \neq a} S_i(p) \right) \right) \sum_{i \neq a} S_i'(p) \\ & - \sum_{i \neq a} \big(S_i(p) + p S_i'(p) \big). \end{aligned} \tag{3.9}$$

Notice that $0 < p^* < \bar{p}_{-a}$ because $H_a(0) \leqslant 0$, $H_a(\bar{p}_{-a}) = 0$ and $H_a'(\bar{p}_{-a}) < 0$. Hence, the optimality of p^* implies that $H_a'(p^*) = 0$.

Since we assumed that the others' supply functions are simple, $H_a(p) = P_a(p)$, where $P_a(p)$ is defined as in (3.2). In particular, Proposition 3.1 implies that there is a unique global maximizer p^*. The space of simple supply functions is rich enough to achieve price p^* because the only condition needed is that the market share x_a at price p^* equals $1 - \sum_{i \neq a} S_i(p^*)$. Indeed, the original S_a optimizes β_a among all nonnegative values and hence is a best response to the others' supply functions, even when a is allowed to bid an arbitrary nondecreasing and differentiable supply function. □

We now extend this argument to any equilibrium and show that the shape of the supply function outside the market-clearing price is irrelevant. The next result proves that an equilibrium with arbitrary supply functions, can be restated using simple supply functions plus an additive constant.

Corollary 3.6 *Assume that each producer $i \in A$ bids a supply function $S_i(p)$ that is nondecreasing and differentiable. If $(S_i)_{i \in A}$ is at equilibrium in the space of all nondecreasing and differentiable supply functions, then this equilibrium is outcome-equivalent to another one where supply functions have the form $\tilde{S}_i(p) = \gamma_i + \beta_i u^{-1}(p)$ for γ_i and β_i chosen by each producer.*

Proof The first paragraph of proof of Theorem 3.5 implies that since the price p^* is optimal from the perspective of producer a, we must have that $H_a'(p^*) = 0$. Equation (3.9) hints that a producer just needs to know the values of $S_i(p^*)$ and $S_i'(p^*)$ for producers $i \neq a$ to know that p^* is the optimal choice of price; the values and derivatives of supply functions at other prices are irrelevant. Thus, producers need only two parameters to setup their supply functions and influence the decisions of others. Indeed, we can construct a new equilibrium based on supply functions of the form $\tilde{S}_i(p) = \gamma_i + \beta_i u^{-1}(p)$ for all $i \in A$, with values of $\beta_i = S_i'(p^*) u'(u^{-1}(p^*))$ and $\gamma_i = S_i(p^*) - \beta_i u^{-1}(p^*)$. Notice that the parameters are chosen such that the new supply functions and their derivatives are equal to the original ones at p^*. Proceeding in a similar way to Proposition 3.1, it is easy to observe that once all producers $i \neq a$ are bidding $\tilde{S}_i(p)$ then, $H_a'(p) = 0$ has a unique solution, which has to be p^*. Hence, p^* is the unique global maximum in (3.8). In conclusion, the new supply functions support an equilibrium that is outcome-equivalent to the original one because the market-clearing price and market shares are the same in both. $\square$

This type of equilibrium may well arise in practice as it occurs, for instance, in electricity and in bond markets. There, it is common to observe "hockey-stick" supply function bids that are flat up to the desired quantity and then sharply grow to infinity [37]. Since this situation is clearly undesirable, we focus on the case where producers only consider (or are allowed to bid) simple supply functions.

4 Monomial cost functions

In this section, we concentrate on monomial cost functions of the form $u(x) = x^q$, where $q > 0$ is a fixed real number. Hence, the total cost equals $\kappa_a(x_a) = c_a x_a^{1+q}$ which is a good first-order approximation to industries with increasing marginal costs. The simplification of cost functions allows us to further simplify the structure of supply functions because under monomial cost functions the producers' decisions can be viewed as selecting a fixed markup to be applied to the production cost. Furthermore, we can characterize equilibria in a sharper way, which we use to provide more structure, to prove the uniqueness of equilibria, to bound the supply functions as a function on the competitiveness among producers, and to construct an efficient procedure to compute equilibria in the case of linear cost functions.

First note that if u is a monomial, it satisfies all the assumptions required by the model in Sect. 2. Applying Theorem 3.5 to monomial cost functions implies that the simple supply functions introduced in Sect. 3.1 are $S_a(p) = \beta_a p^{1/q}$. Therefore, the corresponding price function is

$$p_a(x_a) = \frac{u(1/\beta_a)}{c_a} \cdot c_a u(x_a) = \alpha_a \cdot c_a x_a^q,$$

where we have separated the factor that encodes the producer's decision and denoted it by α_a. Notice that this factor α_a takes the form of a markup applied to the production cost that is independent of the production quantity. From now on, we will consider the markups $(\alpha_a)_{a \in A}$ to be the strategic variables, and the vector of markups corresponding to a supply function equilibrium will be referred to as a *markup equilibrium*.

Reinterpreting Corollary 3.2 in the setting of markups, we have that, at equilibrium, the ordering of producers $c_1 \leqslant \cdots \leqslant c_n$ implies that market shares satisfy $x_1^{\mathrm{NE}} \geqslant \cdots \geqslant x_n^{\mathrm{NE}}$. Furthermore, one can use (4.4) to prove that $\alpha_1 \geqslant \cdots \geqslant \alpha_n$. Intuitively, since the efficient producers know that consumers are going to buy regardless of price, they can increase the price to a level similar to the less efficient ones and still capture a bigger portion of the market.

4.1 Optimal assignment and best responses

With the structure in place, we can obtain explicit formulas for the optimal assignment and for the unique assignment corresponding to given markups. Indeed, $C(x)$ is a convex function and all producers are active under both assignments.

First, we define the competitiveness of an instance as a measure capturing the variability of the producers' efficiency, relying on the cost structure of the instance rather than on the equilibrium itself. As discussed in the introduction this measure has many natural properties and behaves in a desirable way. When it is close to one, there is a large gap between the most efficient and most inefficient producers, while when it is large we face a competitive economic environment.

Definition 2 The competitiveness of an instance is $\sigma := c_1(\sum_{i\in A}(1/c_i)^{1/q})^q \geqslant 1$, i.e., the $\ell_{1/q}$-norm of the vector whose components are c_1/c_i. In this case we say that the instance is σ-competitive.

With this definition, the first-order optimality conditions of the optimal assignment problem give that the market shares and cost of such assignment are:

$$x_a^{\mathrm{OPT}} = \frac{(1/c_a)^{1/q}}{\sum_{i\in A}(1/c_i)^{1/q}} = \left(\frac{c_1}{c_a\sigma}\right)^{1/q}, \quad C(x^{\mathrm{OPT}}) = \sum_{a\in A} x_a^{\mathrm{OPT}} c_a u(x_a^{\mathrm{OPT}}) = \frac{c_1}{\sigma}. \tag{4.1}$$

It also follows immediately that for arbitrary markups $\alpha_1, \ldots, \alpha_n$, the equilibrium allocation for the second-stage game[3] is

$$x_a^{\mathrm{NE}} = \frac{1/(\alpha_a c_a)^{1/q}}{\sum_{i\in A} 1/(\alpha_i c_i)^{1/q}}, \tag{4.2}$$

and under it, the total production cost equals

$$C(x^{\mathrm{NE}}) = \sum_{a\in A} x_a^{\mathrm{NE}} c_a u(x_a^{\mathrm{NE}}) = \left(\frac{1}{\sum_{i\in A} 1/(\alpha_i c_i)^{1/q}}\right)^{1+q} \left(\sum_{a\in A} \frac{1}{(c_a)^{1/q}(\alpha_a)^{1+1/q}}\right).$$

Equation (4.2) and the ordering of the market shares $x_1^{\mathrm{NE}} \geqslant \cdots \geqslant x_n^{\mathrm{NE}}$ imply that $\alpha_1 c_1 \leqslant \cdots \leqslant \alpha_n c_n$.

To obtain an optimal markup α_a, a producer needs to balance the tradeoff between charging high to increase revenue and charging low to increase sales. This is achieved by anticipating consumers decisions when maximizing their profit. Producer a finds α_a by solving $\max_{\alpha_a \geqslant 1} P_a(\alpha_a)$, where

[3] Note that because cost functions are all monomials of the same degree, the vector x^{NE} also represents the market shares that correspond to the solution that minimize the payments [16]. This solution would correspond to the case of a single buyer instead of the situation of perfect competition that we consider in this paper.

$$P_a(\alpha_a) := (\alpha_a - 1)\, x_a^{\mathrm{NE}} c_a u(x_a^{\mathrm{NE}}) = c_a(\alpha_a - 1) \cdot \left(1 + \sum_{i \neq a} \left(\frac{\alpha_a c_a}{\alpha_i c_i}\right)^{1/q}\right)^{-(1+q)},$$

while others' markups α_i for $i \neq a$ are fixed. By Proposition 3.1, the optimal markup is characterized by the first-order conditions of the optimization problem. Then, the best response $\Gamma_a(\alpha_{-a})$ for producer $a \in A$ to a given vector of markups α_{-a} is given by the unique solution to the equation:

$$\alpha_a = 1 + q + q \frac{\alpha_a}{(c_a \alpha_a)^{1/q} \sum_{i \neq a} 1/(\alpha_i c_i)^{1/q}}. \tag{4.3}$$

The marginal cost for producer a is $\kappa_a'(x_a) = (1+q) c_a x_a^q$, so (4.3) indicates that the optimal markup is to charge the marginal cost (the term equal to $1+q$) plus a term that depends on the competition in the market.

Unfortunately, we do not know how to solve the previous system and therefore we cannot compute the equilibrium directly. Then, to prove that a markup equilibrium $(\alpha_1, \ldots, \alpha_n)$ exists for general monomials, we look for a fixed point of the mapping $\Gamma : (\alpha_a)_{a \in A} \to (\Gamma_a(\alpha_{-a}))_{a \in A}$.

4.2 Characterization of equilibria

Specializing the results presented earlier to the case of markups, this section shows that an equilibrium on markups exists if and only if the number of producers exceeds $1 + q$. In addition, we establish the uniqueness of (simple) equilibria in the case of multiple and heterogenous producers with monomial cost functions of an arbitrary degree $q > 0$. Finally, we study some properties of the equilibrium.

We remark that early research on related models already showcased some situations when supply function equilibrium is unique. The details vary depending on the specifics of the model. For example, Turnbull [38] proved uniqueness in the space of linear supply functions in a model with two producers, linear cost functions and uncertain demand. Klemperer and Meyer [26] generalized the previous result by proving that the linear equilibrium is unique even in the space of general supply functions.

Proposition 4.1 *If $u(x) = x^q$ for $q > 0$, then there exists an equilibrium if and only if $n > 1 + q$. Furthermore, if the equilibrium exists, it is unique.*

Proof When $u(x) = x^q$, we have that $(xu'(x) + u(x))/u(x) = 1 + q$, making the conditions of Theorems 3.3 and 3.4 exact opposites. Hence, we have a necessary and sufficient condition for the existence of an equilibrium. We only need to show that whenever an equilibrium exist, it is unique.

Let $(\alpha_1, \ldots, \alpha_n)$ be a markup equilibrium. Observe that if we replace all costs by μc_a, for a scaling factor $\mu > 0$, $(\alpha_1, \ldots, \alpha_n)$ still solves (4.3) for all $a \in A$. Letting the factor μ be $\left(\sum_{a \in A} 1/(\alpha_a c_a)^{1/q}\right)^q$, we can express (4.3) simply by

$$\alpha_a \left(1 - \frac{q}{(c_a \mu \alpha_a)^{1/q} - 1}\right) = 1 + q. \tag{4.4}$$

Considering the function $w(\alpha_a) := 1 + q\alpha_a/(\alpha_a - 1 - q)$, we can rewrite (4.4) as $w(\alpha_a) = (c_a\mu\alpha_a)^{1/q}$. Hence, the solution α_a is achieved at the intersection between the curves w, which is decreasing on α_a, and $(c_a\mu\alpha_a)^{1/q}$, which is increasing on α_a. Thus, a single value of α_a satisfies the equation for fixed μ, and because (4.4) holds for all $a \in A$, a fixed μ generates a unique vector $(\alpha_1, \ldots, \alpha_n)$.

Now, if there were two different equilibria $(\alpha_1, \ldots, \alpha_n)$ and $(\alpha'_1, \ldots, \alpha'_n)$, their corresponding scaling factors μ and μ' had to be different. Assume, without loss of generality, that $\mu < \mu'$. Thus, as w is independent of μ, we have that $(c_a\mu\alpha_a)^{1/q} < (c_a\mu'\alpha'_a)^{1/q}$ for all $a \in A$. Summing across producers, we get that $\sum_{a\in A} 1/(c_a\mu\alpha_a)^{1/q} > \sum_{a\in A} 1/(c_a\mu'\alpha'_a)^{1/q}$, a contradiction since by the definition of μ and μ', both sides should equal one. □

The proof says that if we could guess the appropriate scaling factor μ for the costs, we could find the equilibrium using (4.4). Although this equation cannot be solved in closed form, we will use this idea to construct a procedure to compute the equilibrium for in the linear case.

Observe that $\sigma > (1+q)^q$ implies that $n > 1 + q$, from where $\sigma > (1+q)^q$ is a sufficient condition for an equilibrium to exist. On the contrary, if there are $1 + q$ producers with equal cost, we have that $\sigma = (1+q)^q$ and equilibria do not exist. Since we want to use the competitiveness σ of an instance to provide bounds on the markups and the inefficiency of the resulting equilibria, we are going to adopt the previous condition on σ, which is the tightest possible, to guarantee existence.

Proposition 4.2 *Assume that $q > 0$ and $\sigma > (1+q)^q$. A markup equilibrium $(\alpha_1, \ldots, \alpha_n)$ satisfies that $1 + q \leqslant \alpha_a \leqslant (1+q)(\sigma^{1/q} - 1)/(\sigma^{1/q} - 1 - q)$ for all $a \in A$.*

Proof The lower bound is immediate after (4.3). Let us bound the right-hand side of equation (4.3) for the producer applying the largest markup.

$$\alpha_a \leqslant \alpha_1 = 1 + q + q\frac{\alpha_1}{(\alpha_1 c_1)^{1/q}\sum_{i>1} 1/(\alpha_i c_i)^{1/q}} \leqslant 1 + q + q\frac{\alpha_1}{c_1^{1/q}\sum_{i\in A} 1/c_i^{1/q} - 1}.$$

The upper bound follows directly by recalling the definition of σ. □

For example, for linear cost functions, we have that $2 \leqslant \alpha_a \leqslant 2(\sigma - 1)/(\sigma - 2)$. In particular, if $\sigma = 4$, we know that α is between 2 and 3. This formula can also be used to find the minimum competitiveness that must be present to guarantee that markups will not exceed a given number.

4.3 Computation of equilibria with linear cost functions

We now provide a procedure to compute the unique markup equilibrium in the linear case (i.e., when $q = 1$). We assume that $n \geqslant 3$ because otherwise we know that an equilibrium does not exist. As before, we will choose the scaling factor μ that simplifies calculations. Assume that $(\alpha_1, \ldots, \alpha_n)$ is a markup equilibrium and let $\mu := \sum 1/(\alpha_a c_a) > 0$. Considering the instance with costs equal to μc_a and replacing

$\alpha_a \mu c_a$ by a variable w in (4.3), we get that w satisfies $w^2 - 2(\mu c_a + 1)w + 2\mu c_a = 0$. Solving this equation we get that $w = 1 + \mu c_a + \sqrt{1 + (\mu c_a)^2}$, or equivalently

$$\alpha_a = 1 + 1/(\mu c_a) + \sqrt{1 + 1/(\mu c_a)^2}. \tag{4.5}$$

Now that we know w, we can rewrite the definition of μ as

$$\sum_{a \in A} 1/(1 + \mu c_a + \sqrt{1 + (\mu c_a)^2}) = 1. \tag{4.6}$$

We have just shown that the existence of a markup equilibrium implies that it has to satisfy (4.5), where μ is defined by (4.6). Actually, this characterization can also be used to prove the existence of an equilibrium because (4.6) has a solution if and only if $n \geqslant 3$. More importantly, it can be used to provide a procedure to compute the equilibrium.

Proposition 4.3 *An ε-approximate markup equilibrium can be computed after evaluating the left-hand side of* (4.6) *up to* $\log(n/(\varepsilon \min_{a \in A} c_a))$ *iterations.*

Proof Let μ be such that (4.6) holds. Clearly, we can find an approximation $\tilde{\mu}$ using binary search on (4.6), and then use it to compute an approximate markup equilibrium $(\tilde{\alpha}_a)_{a \in A}$ using (4.5) and its allocation

$$\tilde{x}_a^{\mathrm{NE}} := \frac{1}{\tilde{\alpha}_a \tilde{\mu} c_a} = \frac{1}{1 + \tilde{\mu} c_a + \sqrt{1 + (\tilde{\mu} c_a)^2}}. \tag{4.7}$$

If $\tilde{\mu}$ approximates (4.6) within an additive $\varepsilon > 0$, then it is easy to see that $|\tilde{\mu} - \mu| \leqslant O(1)\varepsilon$ and $|\tilde{x}_a^{\mathrm{NE}} - x_a^{\mathrm{NE}}| \leqslant O(1)\varepsilon$. $\square$

5 Analysis of efficiency for monomial cost functions

In this section, we analyze the efficiency-loss at the markup equilibrium $((\alpha_a)_{a \in A}, x^{\mathrm{NE}})$ of a game, when compared to the optimal assignment x^{OPT}. In other words, we provide bounds on the price of anarchy for the game, which quantifies the loss generated by the lack of coordination in the system [27].

As we will see below, if we do not restrict the instances we consider, the assignment at equilibrium can be arbitrarily bad compared to the social optimum and the markups applied to costs can be arbitrarily large. This high inefficiency comes from instances where producers are extremely different. Suppose there is a very efficient producer and a very inefficient one. Although in an optimal assignment for the market most consumers buy from the efficient producer, at equilibrium the efficient producer will add a large markup to its cost to match the inefficient producer. Hence, as opposed to the optimal situation, the market shares at equilibrium will be comparable, making the social cost of an equilibrium much higher than that of an optimal assignment.

For this reason, we will parametrize all the bounds on the efficiency-loss at equilibrium with respect to the competitiveness of the market, measured by σ. Recall that, by Proposition 4.1, $\sigma > (1+q)^q$ guarantees that an equilibrium exists and is unique, while an equilibrium may not exist when $\sigma = (1+q)^q$. Thus, we concentrate on the former case.

First, we develop a general upper bound for the case when production costs are described by monomial functions. For $\sigma \to \infty$ the upper bound approaches 1, which proves that an equilibrium is almost optimal in a highly-competitive market. We then refine the upper bound in the case when unit costs are linear functions on the production quantity. The latter bound is almost tight, and allows us to show that the price of anarchy is exactly $3/2$ for arbitrary values of $\sigma \geqslant 2$. Parametrizing it with fixed values of σ, it decreases rapidly when σ increases.

5.1 An initial upper bound on the inefficiency

We start by analyzing instances with monomial cost functions of arbitrary degree $q > 0$ and provide bounds depending on the competitiveness in the instance σ. We highlight that our bounds only apply to markup equilibria.

As a warm-up exercise, we provide a simple upper bound on the worst-case production cost at equilibrium using the fact that markups at equilibrium cannot be arbitrarily large for large enough σ. Consider a markup equilibrium for a σ-competitive instance with monomial cost functions of degree $q > 0$ and $\sigma > (1+q)^q$. Then, an equilibrium x^{NE} exists and for any producer $a \in A$, (4.2) implies that $\sum_{i\in A} c_i\alpha_i(x_i^{\mathrm{NE}})^{1+q} = c_a\alpha_a(x_a^{\mathrm{NE}})^q$. Then, by Proposition 4.2,

$$
\begin{aligned}
(1+q)C(x^{\mathrm{NE}}) &\leqslant \sum_{i\in A} c_i\alpha_i(x_i^{\mathrm{NE}})^{1+q} = c_a\alpha_a(x_a^{\mathrm{NE}})^q = \Big(\sum_{i\in A}\frac{1}{(\alpha_i c_i)^{1/q}}\Big)^{-q} \\
&\leqslant \frac{(1+q)(\sigma^{1/q}-1)}{\sigma^{1/q}-1-q}C(x^{\mathrm{OPT}}),
\end{aligned}
$$

so that, $C(x^{\mathrm{NE}}) \leqslant (\sigma^{1/q}-1)/(\sigma^{1/q}-1-q)\cdot C(x^{\mathrm{OPT}})$. As an example, this bound evaluates to $(\sigma-1)/(\sigma-2)$ when production costs are linear and, in particular, to $3/2$ when $\sigma = 4$. In words, although there is no coordination and producers and consumers maximize their individual utilities, the inefficiency generated by competition cannot be extremely large.

Note that it is possible to find a set of instances for which the price of anarchy tends to infinity as σ approaches 1. Thus finding an explicit upper bound for the case of $\sigma < (1+q)^q$ is of less interest and we do not pursue it in this article.

The following theorem improves the previous bound on the efficiency-loss by working directly with the market shares x^{NE} at equilibrium. Although we do not know how to express x^{NE} in closed-form, we prove an upper and a lower bound on it and relax the equilibrium condition by considering a nonlinear programming problem that captures the essence of the calculation of the price of anarchy.

Theorem 5.1 *Consider a markup equilibrium for a σ-competitive instance with monomial cost functions of degree $q > 0$. Assume that ℓ and u are two positive* *numbers such that* $\ell(c_1/c_a)^{1/q} \leqslant x_a^{\mathrm{NE}} \leqslant u(c_1/c_a)^{1/q}$. *If* $\sigma > (1+q)^q$, *then the price of anarchy is bounded by*

$$(\ell\sigma^{1/q})^{1+q} + \sigma(1-\ell\sigma^{1/q})\,\frac{u^{1+q}-\ell^{1+q}}{u-\ell}.$$

Proof Since we do not know how to characterize a markup equilibrium exactly, we will relax the requirement that market shares are at equilibrium and consider an arbitrary market share vector that satisfies the box constraints $\ell(c_1/c_a)^{1/q} \leqslant x_a^{\mathrm{NE}} \leqslant u(c_1/c_a)^{1/q}$. To find an upper bound on the worst-case inefficiency of an equilibrium $C(x^{\mathrm{NE}})/C(x^{\mathrm{OPT}})$, we solve the following nonlinear programming problem:

$$\max\left\{\frac{\sigma}{c_1}\sum_{a\in A} c_a x_a^{1+q},\ \text{subject to} \sum_{a\in A} x_a = 1,\ \ell\left(\frac{c_1}{c_a}\right)^{1/q} \leqslant x_a \leqslant u\left(\frac{c_1}{c_a}\right)^{1/q} \text{ for } a\in A\right\}$$

Considering slack variables z_a from the lower bound, any feasible solution can be written as $x_a = \ell(c_1/c_a)^{1/q} + z_a$, and the first constraint is equivalent to $\sum_a z_a = 1 - \ell\sigma^{1/q}$. Since $c_1 \leqslant \cdots \leqslant c_n$, an optimal solution satisfies that

$$z_a = \begin{cases} 0 & \text{for } 1 \leqslant a < k \\ 1-\ell\sigma^{1/q} - \sum_{i\neq a} z_i & \text{for } a = k \\ (u-\ell)(c_1/c_a)^{1/q} & \text{for } k < a \leqslant n. \end{cases}$$

Here, k is determined so that x_k satisfies the box constraints. Evaluating the optimal objective value and using Newton's generalized binomial theorem [18, p. 162], we get a bound on the price of anarchy. In the following derivation, we also use the upper bound for z_a and the expression for their sum.

$$\begin{aligned}
\frac{\sigma}{c_1}\sum_{a\in A} c_a x_a^{1+q} &= \sigma \sum_{a\in A}\sum_{k=0}^{\infty}\binom{1+q}{k}\left(\frac{c_1}{c_a}\right)^{(1-k)/q} \ell^{1+q-k} z_a^k \\
&= \sigma\ell^{1+q}\sum_{a\in A}\left(\frac{c_1}{c_a}\right)^{1/q} + \sigma\sum_{a\in A} z_a \sum_{k=1}^{\infty}\binom{1+q}{k}\left(\frac{c_1}{c_a}\right)^{(1-k)/q} \ell^{1+q-k} z_a^{k-1} \\
&\leqslant (\ell\sigma^{1/q})^{1+q} + \sigma\sum_{a\in A} z_a \sum_{k=1}^{\infty}\binom{1+q}{k} \ell^{1+q-k}(u-\ell)^{k-1} \\
&= (\ell\sigma^{1/q})^{1+q} + \sigma(1-\ell\sigma^{1/q})\frac{u^{1+q}-\ell^{1+q}}{u-\ell}.
\end{aligned}$$

Since $\ell \leqslant u$ and $\ell\sigma^{1/q} \leqslant 1$, the bound is well defined. □

Interestingly, the upper bound converges to 1 for $\sigma \to \infty$, for any fixed value of q. This says that when competition is high, then equilibria are almost efficient. Unfortunately, for small σ (i.e., $\sigma \approx (1+q)^q$) the bound becomes rather loose; actually, it approaches infinity for $\sigma \to (1+q)^q$.

As an example, let us see that the framework put forward by the previous theorem can be used to provide a meaningful bound on the price of anarchy. Consider the case of a monomial cost function $u(x) = x^q$ for $q \geqslant 1$. To get lower and upper bounds on x_a^{NE}, we use Proposition 4.2. Applying (4.2) to the denominator of (4.3), we have that

$$x_a^{\mathrm{NE}} = \frac{1}{1+q+\frac{q(1+q)}{\alpha_a-1-q}} = \frac{1}{1+q+\frac{(1+q)(\alpha_a c_a)^{1/q}}{\alpha_a}\sum_{i\neq a}\frac{1}{(c_i\alpha_i)^{1/q}}}.$$

Since $q \geqslant 1$, the previous expression is nondecreasing as a function of α_i for all $i \in A$. Thus, using that $\alpha_a \geqslant 1+q$, we can provide the lower bound $x_a^{\mathrm{NE}} \geqslant 1/(q + (\sigma c_a/c_1)^{1/q}) \geqslant (c_1/c_a)^{1/q}/(q+\sigma^{1/q})$.

For the upper bound, (4.2) together with the bounds for α_a imply that

$$x_a^{\mathrm{NE}} \leqslant \left(1+\frac{q}{\sigma^{1/q}-1-q}\right)^{1/q}\frac{1}{\sum_{i\in A}(c_a/c_i)^{1/q}} = \left(\frac{\sigma^{1/q}-1}{\sigma(\sigma^{1/q}-1-q)}\right)^{1/q}\left(\frac{c_1}{c_a}\right)^{1/q}.$$

Putting it all together, we can set $\ell = 1/(q+\sigma^{1/q})$ and $u = ((\sigma^{1/q-1}-1/\sigma)/(\sigma^{1/q}-1-q))^{1/q}$ in the previous theorem to get a bound on the price of anarchy. In particular, when production costs are linear the bound is $(\sigma^2-\sigma-1)/(\sigma^2-\sigma-2)$, which evaluates to $11/10$ when $\sigma = 4$.

Finally, note that better upper and lower bounds on the market shares at equilibrium could be given if one iterates the best responses further. Instead of continuing in that direction, we shall focus on linear cost functions, and use the characterization of equilibria proposed previously to provide an almost tight bound. As a benchmark to evaluate the previously-cited value of $11/10$, the exact price of anarchy for $\sigma = 4$—which we compute in the next section—is approximately 1.027. Moreover, the price of anarchy is exactly $3/2$ if one considers all instances with $\sigma \geqslant 2$.

5.2 Tight bounds using a nonlinear programming formulation

We now compute a tight bound on the price of anarchy for linear cost functions. Our goal is to come up with the worst-case example among instances that are σ-competitive using mathematical programming. In the context of computing the worst-case inefficiency of equilibria, this approach was pioneered by Johari and Tsitsiklis [23]. More precisely we proceed as follows. We characterize the price of anarchy exactly for a fixed $\sigma \geqslant 2$ as a nonconvex nonlinear optimization problem having all c_a's as decision variables. After transforming this problem slightly, we use its optimality conditions to cast it as a nonlinear mixed integer programming problem with few decision variables, an idea that might be useful elsewhere. Then, we relax the integrality constraints an explicitly solve the remaining problem, which provides an almost tight bound for the price of anarchy.

Using the characterization of equilibria developed earlier, the following mathematical program finds a worst-case σ-competitive instance, represented as the vector $(c_1, \ldots, c_n)$, by maximizing the gap between the total production cost at equilibrium to that of an optimal solution.

$$POA(\sigma) := \sup_{n\geqslant 3}\left\{\max \left(\sum_{a=1}^{n}\frac{1}{c_a}\right)\left(\sum_{a=1}^{n}\frac{c_a}{(1+c_a+\sqrt{1+c_a^2})^2}\right)\right\} \quad (5.1a)$$

$$\text{s.t.} \sum_{a=1}^{n}\frac{1}{1+c_a+\sqrt{1+c_a^2}} = 1 \quad (5.1b)$$

$$c_1\sum_{a=1}^{n}\frac{1}{c_a} = \sigma \quad (5.1c)$$

$$0 < c_1 \leqslant c_a \quad \forall a \in \{2,\ldots,n\}. \quad (5.1d)$$

Here, (5.1b) guarantees that we can use the characterization given previously to compute equilibria, (5.1c) imposes that the instance is σ-competitive, and the objective function, which equals $C(x^{\mathrm{NE}})/C(x^{\mathrm{OPT}})$, computes the inefficiency of the markup equilibrium of the instance represented by the feasible solution. Notice that this allows us to compute the inefficiency of an equilibrium without explicitly computing the equilibrium. For a fixed $n \geqslant 3$, the maximum in (5.1) is attained, and we will see that it grows as n goes to infinity.

The previous problem has a nonlinear nonconvex objective function and constraints. To get around this, note that the constraint $c_1 > 0$ can be relaxed to $c_1 \geqslant 0$, because $c_1 = 0$ implies $C(x^{\mathrm{OPT}}) = C(x^{\mathrm{NE}}) = 0$. Then, we can consider variables $0 \leqslant y_a \leqslant 1$, for all $a \in A$, defined by

$$y_a := 1 - \frac{2}{1+c_a+\sqrt{1+c_a^2}}. \quad (5.2)$$

The inverse transformation is $c_a = 2y_a/(1-y_a^2)$, so problem (5.1), is reformulated as

$$\begin{aligned} POA(\sigma) = \sup_{n\geqslant 3}\ \max\ & \frac{\sigma}{4}\left(\frac{1}{y_1}-y_1\right)\left(n+2-2\sum_{a=1}^{n}\frac{1}{1+y_a}\right) \\ \text{s.t.} & \sum_{a=1}^{n} y_a = n-2 \\ & \sum_{a=1}^{n}\left(\frac{1}{y_a}-y_a\right) = \sigma\left(\frac{1}{y_1}-y_1\right) \\ & 0 \leqslant y_1 \leqslant y_a \leqslant 1 \quad \forall a \in \{2,\ldots,n\}. \end{aligned} \quad (5.3)$$

Lemma 5.2 *The maximum in the subproblem of* (5.3) *is increasing with* n.

Proof Increasing n to $n' > n$ increases the objective. Indeed, given a solution with n components, setting the new $n' - n$ variables to 1 achieves the same objective value since all constraints are satisfied and the objective value does not change. □

The argument in the proof also implies that variables y_a taking a value of 1 are not useful. To see this, notice that if for a given value of n, it is optimal to set some variables to 1, we can remove those producers without affecting the feasibility of the solution, nor its objective value.

Let us consider the following problem in which we take $n \geqslant 3$ and $0 < y_1 < 1$ fixed, and optimize over $y_2, \ldots, y_n$. We characterize the structure of the optimal solution to this problem and optimize over n and y_1 afterwards to get the solution to (5.3).

$$\begin{aligned} &\min \sum_{a=2}^{n} \frac{1}{1+y_a} \\ &\text{s.t.} \sum_{a=2}^{n} y_a = n-2-y_1 \\ &\sum_{a=2}^{n} \frac{1}{y_a} = \sigma\left(\frac{1}{y_1} - y_1\right) + n - 2 - \frac{1}{y_1} \\ &y_1 \leqslant y_a \leqslant 1 \quad \forall a = 2, \ldots, n. \end{aligned} \tag{5.4}$$

We denote the dual variables of this subproblem by λ, μ, ℓ_a and u_a, in the order in which they appear in the formulation above. We use the standard KKT conditions to characterize the optimal solution to this problem (see, e.g., [32, p. 321]). Indeed, if a vector $y = (y_2, \ldots, y_n)$ is optimal, then when the gradients of the active constraints at y are linearly independent, y verifies the KKT conditions

$$\frac{-1}{(1+y_a)^2} + \lambda - \frac{\mu}{y_a^2} + u_a - \ell_a = 0 \quad \forall a = 2, \ldots, n. \tag{5.5}$$

With the previous conditions we conclude that an optimal solution has a well-defined structure.

Lemma 5.3 *An optimal solution* $(y_2, \ldots, y_n)$ *to* (5.4) *satisfies that there are two numbers* $\bar{y}$ *and* $\bar{\bar{y}}$ *in the open interval* $(y_1, 1)$ *such that* $y_a \in \{y_1, \bar{y}, \bar{\bar{y}}\}$ *for all* $a \in A$.

Proof Assume vector $y = (y_2, \ldots, y_n) \in [y_1, 1)^{n-1}$ is an optimal solution to (5.4). We can assume that there are at least three different values larger than y_1 because otherwise the claim holds. We refer to three of those values with y_i, y_j and y_k, ordered from low to high. The KKT conditions hold for y because the gradients of the active constraints at y are linearly independent. Indeed, the gradients of the two equality constraints are $(1, \ldots, 1)$ and $(-1/y_2^2, \ldots, -1/y_n^2)$, and those of the variables with values equal to y_1 are $(0, \ldots, 0, -1, 0, \ldots, 0)$ with the -1 in the position corresponding to the variable. The constraints of variables with value different from y_1 are not active so they are not considered. Considering the variables y_i and y_j only, the restricted gradients are $(1, 1)$, $(-1/y_i^2, -1/y_j^2)$, and $(0, 0)$. Because the first two are linearly independent, all the vectors are linearly independent and, thus, y_i, y_j and y_k satisfy (5.5). Using the complementary slackness property for these variables and solving for λ, we have that

$$\lambda = \frac{1}{(1+y_i)^2} + \frac{\mu}{y_i^2} = \frac{1}{(1+y_j)^2} + \frac{\mu}{y_j^2} = \frac{1}{(1+y_k)^2} + \frac{\mu}{y_k^2}.$$

We focus first on y_i and y_k. Solving for μ and plugging the result back in, we get that

$$\lambda = \frac{\frac{1}{(1+y_i)^2} - \frac{1}{(1+y_k)^2}}{1 - \left(\frac{y_k}{y_i}\right)^2} + \frac{1}{(1+y_k)^2} = \frac{1}{y_i^2 - y_k^2}\left[\left(\frac{y_i}{1+y_i}\right)^2 - \left(\frac{y_k}{1+y_k}\right)^2\right].$$

Taking the derivative with respect to y_k, we see that the right-hand side of the previous equation is increasing on y_k. After some algebra, the derivative is positive if and only if

$$\left(\frac{y_k}{1+y_k}\right)^2 - \left(\frac{y_i}{1+y_i}\right)^2 > \frac{y_k^2 - y_i^2}{(1+y_k)^3},$$

which holds because $y_i < y_k$. Now, doing the same calculation with y_i and y_j provides us with another λ. But that cannot happen because y_j and y_k satisfy the KKT conditions. □

The previous result implies that an optimal solution to (5.4) has the structure $(y_1, \ldots, y_1, y_i, \ldots, y_i, y_j, \ldots, y_j)$, where y_1 is repeated k_1 times and y_i is repeated k_i times and y_j is repeated $k_j = n-1-k_1-k_i$ times. Hence, (5.4) can be reformulated as:

$$\begin{aligned}
\min\ & \frac{k_1}{1+y_1} + \frac{k_i}{1+y_i} + \frac{k_j}{1+y_j} \\
\text{s.t. } & k_1 y_1 + k_i y_i + k_j y_j = n - 2 - y_1 \\
& \frac{k_1}{y_1} + \frac{k_i}{y_i} + \frac{k_j}{y_j} = \sigma\left(\frac{1}{y_1} - y_1\right) + n - 2 - \frac{1}{y_1} \\
& k_1 + k_i + k_j = n - 1 \\
& y_1 \leqslant y_i \leqslant y_j \leqslant 1, k_1, k_i, k_j \in \mathbb{N},
\end{aligned} \tag{5.6}$$

where y_i, y_j, k_1, k_i, and k_j are the variables, and y_1 and n are fixed. We conclude the following.

Lemma 5.4 *The price of anarchy for a given value of $\sigma > 1$ is given by:*

$$\begin{aligned}
POA(\sigma) = \sup\ & \frac{\sigma}{4}\left(\frac{1-y_1^2}{y_1}\right)\left(n + 2 - \frac{2(k_1+1)}{1+y_1} - \frac{2k_i}{1+y_i} - \frac{2k_j}{1+y_j}\right) \\
& s.t.\ (y_i, y_j, k_1, k_i, k_j) \textit{ solves problem } (5.6) \textit{ for } y_1 \textit{ and } n. \\
& 0 \leqslant y_1 \leqslant 1, n \geqslant 3, n \in \mathbb{N}
\end{aligned} \tag{5.7}$$

Observe that although the previous problem only has seven variables, it is a nonconvex integer programming problem. Nevertheless, we can compute almost tight lower and upper bounds to this problem, which we express in closed-form as a function of $\sigma \geqslant 2$. The lower bound comes from feasible solutions and the upper bound follows from a relaxation of (5.6).

Note that relaxing the integrality constraints on k_1, k_i and k_j in problem (5.6) leads to an optimal solution in which $y_i = y_j$. Indeed, consider the relaxation of a subproblem of (5.6) in which k_1 is fixed, k_i and k_j are variables, and the three are nonnegative reals, and assume that an optimal solution satisfies that $y_1 < y_i < y_j$ with k_i and k_j strictly positive. The gradients of the constraints at the solution are linearly independent so we can use the KKT conditions. We denote the dual variables of the equality constraints in (5.6) with λ, μ and η, respectively. The KKT conditions for k_i and k_j imply that $f(y) := 1/(1+y)+\lambda y+\mu/y+\eta = 0$ when we evaluate in $y = y_i$ and in $y = y_j$, and thus there is a point $y_i < \gamma < y_j$ where the derivative $f'(\gamma)$ vanishes. In addition, the KKT conditions for y_i and y_j say that $f'(y_i) = f'(y_j) = 0$. This is a contradiction since $-f'(y) = 1/(1+y)^2 + \mu/y^2 - \lambda$ is unimodal in (0,1) so it can vanish at most twice.

The observation in the previous paragraph implies that the following reformulation of (5.7) with values of y equal to y_1 or y_i provides a closed-form upper bound on the price of anarchy.

$$
\begin{aligned}
&\sup \frac{\sigma}{4}\,\frac{1-y_1^2}{y_1}\left(2-\frac{2(k_1+1)(y_i-y_1)}{(1+y_1)(1+y_i)}-n\frac{1-y_i}{1+y_i}\right)\\
&\text{s.t. } (k_1+1)(y_i-y_1)+n(1-y_i)=2\\
&\qquad (k_1+1)(y_i-y_1)\frac{1+y_1y_i}{y_1y_i}+n(1-y_i)\frac{1+y_i}{y_i}=\sigma\left(\frac{1}{y_1}-y_1\right)\\
&\qquad 0\leqslant y_1\leqslant y_i\leqslant 1,\, k_1\geqslant 0,\, n\geqslant 3.
\end{aligned}
\tag{5.8}
$$

To solve this problem, we first solve the linear system for k_1 and n given by the equality constraints. Then we plug the result back into the objective function, allowing us to rewrite (5.8) as:

$$
\max_{0\leqslant y_1\leqslant y_i\leqslant 1} \frac{\sigma}{4}\,\frac{1-y_1}{y_1}\left(\frac{2y_1+y_i(2+4y_1-\sigma(1-y_1^2))}{1+y_i}\right).
$$

Observe that the objective is a rational function of y_i, and a simple calculation shows that it is increasing if and only if $y_1 \geqslant 1-2/\sigma$. If the rational function were decreasing in an optimal solution, that would imply that $y_i = y_1$, which would evaluate to 1 in the maximum above. Since that cannot be the case, the rational function has to be increasing. Therefore, $y_i = 1$ in an optimal solution, thus making the previous maximum equal to

$$
\max_{1-2/\sigma\leqslant y_1\leqslant 1} \frac{\sigma}{4}\,\frac{1-y_1}{y_1}\left(1+3y_1-\frac{\sigma(1-y_1^2)}{2}\right),
$$

which is strictly greater than 1. To maximize the previous expression, we set its derivative to zero, and find the roots of $-2\sigma y_1^3 + (\sigma-6)y_1^2 + (\sigma-2)$. For $\sigma \geqslant 2$, the largest real root $y(\sigma)$ is in the interval $[1-2/\sigma, 1]$ (actually, there is only one root when $2 \leqslant \sigma \leqslant 2.33462\ldots$). Solving the cubic equation in closed form, we conclude that the optimal solution for y_1 is

$$y(\sigma) := (v + \sqrt{v^2 - r^6})^{1/3} + (v - \sqrt{v^2 - r^6})^{1/3} + r\,, \tag{5.9}$$

where $v = 55/216 - 7/(12\sigma) + 1/(2\sigma^2) - 1/(\sigma^3)$ and $r = (\sigma - 6)/(6\sigma)$. Plugging the value back into the objective, we get the following theorem.

Theorem 5.5 *If* $\sigma \geqslant 2$, *then*

$$POA(\sigma) \leqslant \frac{\sigma}{4}\left(\frac{1 - y(\sigma)}{y(\sigma)}\right)\left(1 + 3y(\sigma) - \frac{\sigma(1 - y(\sigma)^2)}{2}\right),$$

where $y(\sigma)$ *is given by* (5.9). *This bound is tight infinitely often for* $\sigma \to \infty$. *In particular, the value is exactly* $3/2$ *when* $\sigma = 2$.

Proof We have already proved the upper bound, we still need to show that the bound is tight for infinitely many values of σ. We specifically show that the bound is tight whenever

$$k_1(\sigma) = \sigma\frac{1 + y(\sigma)}{1 - y(\sigma)} - \left(\frac{1 + y(\sigma)}{1 - y(\sigma)}\right)^2 \tag{5.10}$$

is integral. Indeed, in this case we can evaluate (5.8) with $k_1 = k_1(\sigma)$, $y_1 = y(\sigma)$, a large enough n, and the appropriate value of y_i (which will be close to 1). It is not hard to see that the objective value approaches that in the claim of this theorem when $n \to \infty$. Since in this situation k_1 and n are integral, we can construct a sequence of instances whose inefficiency asymptotically equal our bound. Observe that because $y(\sigma) = 1 - 2/\sigma + o(1/\sigma)$, $k_1(\sigma)$ increases to infinity, so that it is integral for infinitely many values of σ. In particular for $\sigma = 2$, $y(\sigma) = 0$ and then $k_1(\sigma) = 1$. □

An almost matching lower bound for other values of σ (different from those leading to an integer $k_1(\sigma)$) is obtained by restricting solutions to have only two values. For fixed σ, and fixed integers n and k_1, we can find the best possible solution to (5.7) for which $y_i = y_j$. This is done by solving (5.10) with those parameters fixed, which is easy since the constraints amount to explicitly evaluate y_1 and y_i (and y_1 turns out to be the solution of a cubic equation similar to (5.9)). In the limit when $n \to \infty$ we obtain the following result.

Theorem 5.6 *If* $\sigma > 2$, *then*

$$POA(\sigma) \geqslant \sigma\frac{2(1 - y_1^2) - (k_1 + 1)(1 - y_1)^3}{8y_1},$$

where k_1 *is a nonnegative integer and* $y_1 = (k_1 - 1 + \sqrt{\sigma^2 - 4k_1})/(\sigma + k_1 + 1)$.

The two natural candidates for k_1 in Theorem 5.6 are the integers closest to $k_1(\sigma)$ given by (5.10). More precisely we evaluate the lower bound for $\lfloor k_1(\sigma)\rfloor$ and $\lceil k_1(\sigma)\rceil$, and take the maximum of the two values. For the special case of $\sigma = 2$, this bound is not well defined because $y_1 = 0$. Nevertheless, the worst case instance is easy to

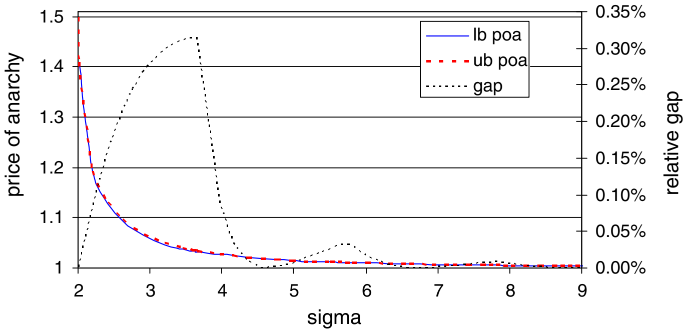

Fig. 1 Lower bound and upper bounds for the inefficiency of Nash equilibria. The *vertical axis* on the *right* displays the relative distance between the two bounds

construct in this case by taking $n = 3$ and three different values of y. A calculation shows that, in the limit, this worst-case instance has a price of anarchy of 3/2, matching the upper bound of Theorem 5.5.

Notice that the analysis proposed in this section provided us with significantly improved bounds compared to those in the previous section (which were arbitrarily loose for $\sigma \to 2$). Our new bound is tight in that case, and also when $\sigma \to \infty$, it fails to be tight in general because of the integrality gap arising in Lemma 5.4. Moreover, we have established that the bound is tight whenever the resulting k_1 is integral, which happens for infinite values of σ. Figure 1 depicts both bounds as a function of σ. The figure also includes the relative gap $(ub(\sigma) - lb(\sigma))/ub(\sigma)$ between the two bounds on the secondary vertical axis. As an example, the upper bound evaluates to approximately 1.02717 when $\sigma = 4$ while the lower bound evaluates to 1.02642. The worst relative gap between the lower and upper bounds is 0.316 % for a value of $\sigma \approx 3.65$.

6 General market structures

In this section, we characterize markup equilibria for markets that contain some complementary goods. We model unit demand for a "bundle", which may in fact be purchased from a number of separate producers, each selling some particular component. For linear cost functions (i.e. $u(x) = x$), and markets taking a Series-Parallel (SP) network structure, we extend the existence result of Proposition 4.1. In Sect. 6.4, we show that the result may no longer hold for more general market structures.

Supply function equilibrium is a natural modeling choice in the bundled setting, as scheduled quantity-dependent price adjustments remove the ambiguity around revenue-splitting that would result in a Cournot-type model of complementary producers. Supply function models yield a structure where bundle-level purchase quantities (i.e., path flows) uniquely determine both the producer-level purchase quantities (link flows) and the market price of each producer's output. In contrast, a pure quantity-commitment model lacks a mechanism for setting individual prices.

Practically, bundling is an important consideration in many industries. In freight shipping, for example, point to point routes often involve multiple carriers, each servicing a distinct geography and/or mode of transport. The model also applies to decentralized assembly supply chains, where a manufacturer contracts separately to purchase components from any number of suppliers. Such outsourcing typically requires a modular product structure that is amenable to series-parallel representation. Taking the assembler as a monopsonistic buyer, one could employ our model to study competition among individual component suppliers (e.g. producers of processors, hard disks, displays, etc. in a computer system supply chain).

Let $\mathcal{B} := \{B_1 \ldots B_m\}$ represent a set of bundles, all equivalent in the eyes of costumers, that may each be used to satisfy the demand. We propose a network model to represent potential mappings of producers to bundles. That is, we model the set of available purchase combinations as paths from the source s to the sink t of a directed network, G, comprising a set of n links A_G. Each link $a \in A_G$ represents a producer, and each path through G a bundle in $\mathcal{B}$. As before, customers choose the lowest-priced complete bundle, with a set of path flows f_i denoting the proportion of customers choosing bundle B_i. We say that $a \in B_i$ if link a appears in bundle B_i, and set $x_a = \sum_{B_i \ni a} f_i$. Thus, the vector $f \in \mathbb{R}^m$ describes consumption decisions, while taking $x \in \mathbb{R}^n$ to describe production quantities, and the full set of possible production-consumption pairs is given by

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

For a given set of markups α, the second-stage production quantities are given by an assignment $(x(\alpha), f(\alpha)) \in \mathcal{F}$ satisfying

$$\sum_{a \in B_i} \alpha_a c_a u(x_a(\alpha)) \leqslant \sum_{b \in B_j} \alpha_b c_b u(x_b(\alpha)) \tag{6.1}$$

for all $B_i, B_j \in \mathcal{B}$ such that $f_i(\alpha) > 0$. The assignment $x(\alpha)$ is unique for any α because the function $u(\cdot)$ is strictly increasing [10]. There may, however, be multiple consumption allocations that give rise to x.

6.1 Series-parallel networks

The SP restriction that we employ is suited for markets with both complements and substitutes. The class of SP networks are exactly those that can be constructed recursively through series and parallel compositions. To formalize, we define the composition operations $S(\cdot)$ and $P(\cdot)$, each of which takes as input a set $\mathcal{G}$ of SP networks, and returns a single SP network. In the case of $S(\mathcal{G})$, 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 $P(\mathcal{G})$, the input networks are composed in parallel so that all share a common source and sink. Given this recursive construction, we use the notation $g \subseteq G$ to refer to those *submarkets* that have been nested within G. Note that the allocation of customers to paths in a submarket g satisfies a condition equivalent to (6.1), but g serves only a portion of the total demand for G. A submarket g can be characterized as either a *series submarket*, indicating that $g = S(\mathcal{G})$ for some set $\mathcal{G}$ of submarkets, or a *parallel submarket*, composed as $g = P(\mathcal{G})$.

We introduce notation to describe the structure of nesting within G. Let $\psi(g)$ be the set of component markets comprising 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. Now, for submarkets g', g with $A_{g'} \subseteq A_g$, let $\psi_{g'}(g)$ return only the component market of g that contains g', instead of the full set of components. We let $\nu_g := (G, \psi_g(G), \psi_g^2(G), \ldots, \psi_g^{h_g}(G) = g)$ denote the unique sequence of submarkets starting with G within which g is nested, where h_g is the depth at which g is nested. Finally, let $\nu_{g,P} = (g_1, g_2 \ldots g_d)$ (alternatively, $\nu_{g,S}$) be the subsequence of odd or even elements of ν_g restricted to only parallel (series) submarkets. The sequence $\nu_{g,P}$ provides the increasingly specific decisions that a customer must make before purchasing from g.

Recall from Sect. 3 that the market-clearing price for a market g of parallel links and unit demand solves $\sum_{i \in A_g} \beta_i u^{-1}(p_g^*) = 1$. Replacing the right-hand side with an arbitrary demand d_g, and writing it in terms of markups gives $p_g^* = d_g(\sum_{i \in A_g} 1/(\alpha_i c_i))^{-1}$. This expression decomposes into $p_g^* = d_g R_g(\alpha)$, where we denote the response of the network by a single *network price multiplier* $R_g(\alpha)$. As $R_g(\alpha)$ does not depend on d_g, the multiplier defines a linear price function for the market g as a whole. Furthermore, beginning with $R_a(\alpha) = \alpha_a c_a$ for an individual producer a, all submarket price multipliers can be constructed recursively according to:[4]

$$R_{S(\mathcal{G})}(\alpha) = \sum_{g \in \mathcal{G}} R_g(\alpha), \quad \text{and} \quad R_{P(\mathcal{G})}(\alpha) = \left(\sum_{g \in \mathcal{G}} 1/R_g(\alpha) \right)^{-1}. \tag{6.2}$$

Lastly, for notation, given g, g' with $A_g \subseteq A_{g'}$, we use $g' \backslash g$ to denote the market in g' with producers from g removed, and α_{-g} to denote markups of producers in $G \backslash g$.

6.2 Second stage analysis

In this section we present a precise functional form of the second-stage assignment $x(\alpha)$, where α is an arbitrarily fixed vector of markups. As α is fixed, we will suppress dependence of assignments and price multipliers on α where possible to simplify notation. The assignment in (4.2) holds within each parallel submarket, and accounting for nested component choices yields $x_a = \prod_{g \in \nu_{a,P}} R_g / R_{\psi_a(g)}$. However, this formula

[4] Equation (6.2) matches that used for electrical circuits to compute the equivalent resistance when placing resistors in series and parallel. Ohm's law, *Voltage* $=$ *Current* $\cdot$ *Resistance*, is analogous to the price function $p_a = x_a R_a$. Although the equations describing both systems are identical, the difference is that we impose a nonnegativity restriction on flows, whereas in electricity networks this is not needed. It is precisely those restrictions that complicate the analysis of a general network as we will discuss in Sect. 6.4.

is hard to manipulate directly and provides little insight into how producers will set their markups. We thus provide an alternative formula that is more amenable to analysis. In particular, we now express x_a in terms of aggregate measures of competitiveness for a's substitute and complementary producers, respectively.

Its own markup aside, each firm's production increases with the markups of substitute products, while abating in response to those of complementary competitors. If producer *a spans* the market, that is its link a connects s and t directly, 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 requires a complement, the residual demand for product a is shifted downwards as the markups on complementary items increase. Both effects can occur simultaneously for a competitor $b \neq a$, and so the impact of α_b on x_a is not clear *a priori*.

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, a spans the pivoted market and all competition with a is transformed to a substitute. 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 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 can be 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 the pivoted network, the *substitute network* for producer a and denote it by $G \ominus a$. Figure 2 demonstrates the construction of the substitute network. Note that the example in (c) contains complements and so requires pivoting.

The uniqueness of the niche that producer a fills will determine the multitude of paths in $G \ominus a$, and plays a key role in determining market power. A measure of this market power can be encoded by $R_{\ominus a}$ as a function of the prevailing markups of others α_{-a}, where from now on for brevity we omit G from $G \ominus a$ in subindices. In

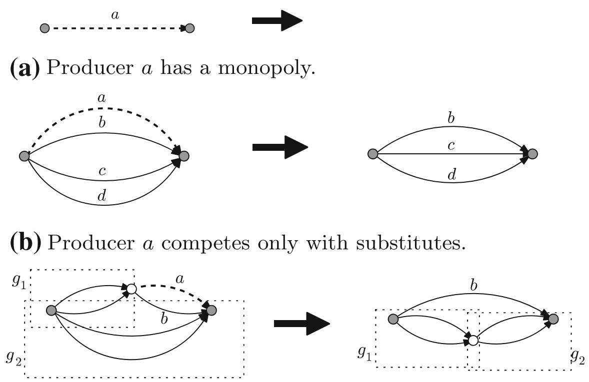

(a) Producer a has a monopoly.

(b) Producer a competes only with substitutes.

(c) Producer a competes with substitute and complementary producers. Links in g_1 are reversed to form $G \ominus a$.

Fig. 2 The producer's substitute network $G \ominus a$

general, $R_{\ominus a}$ measures the substitutability of producer a in equilibrium, with a higher multiplier indicating a relative absence of attractive alternatives to a. As shown in Fig. 2, $G \ominus a$ is empty in the extreme case of a monopolist, so $R_{\ominus a} = \infty$.

We show that the effect of a producer's markup on its own profit is captured succinctly through the ratio of R_a to $R_{\ominus a}$. To isolate this effect, we express demand as the product of two factors. One factor depends entirely on this ratio, and the other is a scaling factor, also independent of α_a, that measures the demand for producer a when $\alpha_a = 0$. Indeed, factor μ_a, which depends only on α_{-a}, captures the competitiveness of complementary, rather than substitute, producers for a. In doing so, it accounts for the position of a in the market, prior to pivoting. Both the multiplier $R_{\ominus g}$ and scaling factor μ_g are defined analogously for any submarket $g \subseteq G$.

The scaling factor μ_a is expressed:

$$\mu_a = \prod_{g \in \nu_{a,S}} \frac{R_{\ominus g}}{R_{\ominus g} + R_{g \backslash \psi_a(g)}}. \tag{6.3}$$

To illustrate, in Fig. 2c, $\nu_{a,S}$ contains a single element, $S(g_1, P(a, b))$, and $\mu_a = R_{g_2}/(R_{g_2} + R_{g_1})$. Note that when a competes with substitute producers only, $\mu_a = 1$ for any value of α. In the case with complements, the factor is strictly less than one, and decreasing in the markups demanded for complements of a. In general, the factor μ_a may be increasing, decreasing, or unaffected by R_b, depending on whether b is largely a substitute or a complement of a.

To summarize, for a fixed vector α_{-a}, the parameters μ_a and $R_{\ominus a}$ measure, respectively, the complementary and substitute competition facing producer a. In Proposition 6.1, we express x_a in terms of these two quantities and producer a's own multiplier. In this formulation, μ_a determines the intercept of producer a's residual demand, and $R_{\ominus a}$ determines the slope with respect to α_a.

Proposition 6.1 *For a market G with price functions fixed according to α, and for any producer a, the equilibrium assignment $x(\alpha)$ takes the form*

$$x_a = \mu_a \left[\frac{R_{\ominus a}}{R_{\ominus a} + R_a} \right] = \mu_a \left[\frac{R_{\ominus a}}{R_{\ominus a} + \alpha_a c_a} \right]. \tag{6.4}$$

Proof We will extend this property inductively to all submarkets, including individual producers, beginning with the full market G. We have that $x_G = \mu_G R_{\ominus G}/(R_{\ominus G} + R_G)$ because demand x_G is inelastic and equal to 1, the factor $\mu_G = 1$, R_G is finite, and $R_{\ominus G} = \infty$. Now we assume that (6.4) holds for a submarket g and prove it for an arbitrary component $g' \in \psi(g)$. If g is composed in series, then $g = S(g', g'')$ where $g'' = g \backslash g'$. Then $R_g = R_{g'} + R_{g''}$ and $R_{\ominus g'} = R_{g''} + R_{\ominus g}$. Because g is series, we adjust the scaling factor so that $\mu_{g'} = \mu_g R_{\ominus g}/(R_{\ominus g} + R_{g''})$. Then,

$$\begin{aligned} x_{g'} = x_g = \mu_g \left[1 + R_g/R_{\ominus g}\right]^{-1} &= \mu_g \left[1 + R_{g'}/R_{\ominus g} + R_{g''}/R_{\ominus g}\right]^{-1} \\ = \mu_{g'} \left[1 + R_{g'}/R_{\ominus g'}\right]^{-1}. \end{aligned}$$

If g is composed in parallel, then $g = P(g', g'')$ where $g'' = g \backslash g'$. Then $R_g = [1/R_{g'} + 1/R_{g''}]^{-1}$ and $R_{\ominus g'} = [1/R_{\ominus g} + 1/R_{g''}]^{-1}$. Since g is parallel, $\mu_{g'}$ is exactly μ_g and $x_{g'} = x_g R_g/R_{g'}$. So:

$$\begin{aligned} x_{g'} &= (R_g/R_{g'})\mu_g[1 + R_g/R_{\ominus g}]^{-1} = (R_g/R_{g'})\mu_{g'}[1 + (R^{-1}_{\ominus g'} - R^{-1}_{g''})R_g]^{-1} \\ &= \mu_{g'}[1 + R_{g'}/R_{\ominus g'}]^{-1}. \end{aligned}$$

□

Setting the derivative of the profit function with respect to α_a to zero, we can characterize the best response of any producer. The size of producer a's markup in the first stage will depend on its market power; i.e. the substitutability metric $R_{\ominus a}$.

Proposition 6.2 *The best-response markup of any producer a to its competitors' markups is*

$$\Gamma_a(\alpha_{-a}) = 2 + R_{\ominus a}(\alpha_{-a})/c_a. \tag{6.5}$$

Proof Having shown that the solution to the second-stage game is $x_a(\alpha) = \mu_a \left[R_{\ominus a}/(R_{\ominus a} + \alpha_a c_a)\right]$, in the first stage, producer a chooses $\alpha_a \geqslant 1$ that maximizes $P_a(\alpha_a) = (\alpha_a - 1)c_a(x_a(\alpha_a, \alpha_{-a}))^2$. Evidently, the profit-maximizing markup is interior in $(1, \infty)$ because $P_a(\alpha_a) \to 0$ as $\alpha_a \to 1$ and $P_a(\alpha_a) = O\left(\alpha_a^{-1}\right) \to 0$ as $\alpha_a \to \infty$. From the first-order optimality conditions,

$$\alpha_a = 1 - [x_a(\alpha)]\left[\partial x_a(\alpha)/\partial \alpha_a\right]^{-1}/(2c_a).$$

Since $R_{\ominus a}$ and μ_a do not depend on α_a, it is straightforward to differentiate $x_a(\alpha)$, getting $\partial x_a(\alpha)/\partial \alpha_a = -x_a(\alpha)/(R_{\ominus a} + \alpha_a c_a)$. Note that this term is nonzero for any finite α. Substituting into the above gives: $\alpha_a = 1 + \frac{1}{2}\left(R_{\ominus a}/c_a + \alpha_a\right)$. □

In terms of $R_{\ominus a}$, the per-unit price that producer a will charge in equilibrium is $p_a(x_a) = R_a x_a = 2c_a x_a + R_{\ominus a} x_a$. 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 $R_{\ominus a} x_a$ that depends on the substitutability of a. Note for a market without complements that $R_{\ominus a} = (\sum_{i \neq a} 1/(\alpha_i c_i))^{-1}$, and (6.5) generalizes from (4.3).

6.3 Graph connectivity and existence of equilibria

We now explore the existence of markup equilibria, which requires we establish some upper bound on the markups α. We will see that the critical property in establishing a bound is the degree of connectivity of the network structure. A set of links whose removal disconnects the graph is a *cut*, and a graph is *k-edge-connected* if there are no cuts containing less than k links [4]. The connectivity between two nodes is the maximum number of disjoint paths connecting them, and the connectivity of a network is the minimum over an arbitrary pair of nodes. In the above, a producer's demand elasticity is shown to depend directly on the number (and ultimately, price) of alternative paths available for joining the nodes that the producer connects in G. As such, a high degree of connectivity should translate to some bound on the markups of any individual producer. In this section we formalize this idea.

For a submarket g, the connectivity $Q(g)$ is the largest k for which g is k-edge-connected. The orientation of the graph does not factor into $Q(g)$. In contrast, an *internal cut* is one that does not separate the source from the sink, so that producers in the cut belong to some common bundle. Redefining connectivity in terms of internal cuts alone gives the internal connectivity $I(g)$. Composing g with additional producers in parallel may increase $Q(g)$. The internal connectivity $I(g)$ provides an upper bound on the connectivity of any market within which g is nested.

Looking at (6.5), it is clear with only two producers that the combined sensitivity of the responses leads to an infinitely increasing sequence of markups. This applies as well to any network with $Q(G) < 3$. Essentially, stability requires that the substitute network for any producer is 2-edge-connected. When $G \ominus a$ is not 2-edge-connected, it is producer a and the producer that disconnects $G \ominus a$ that combine to drive instability. The key to establishing existence in G is that producers are arranged into submarkets in such a way that their sensitivity, in the aggregate, to competitors' markups diminishes with the size of those markups. We now show that when the graph is 3-edge-connected, there is enough competition to ensure that markups are bounded.

Theorem 6.3 *A markup equilibrium exists in G if and only if the network is 3-edge-connected. When it exists, this equilibrium is unique.*

Proof For a submarket g and fixed price multiplier $R_{\ominus g}$, a *partial markup equilibrium* on g results if all producers in g choose markups optimally, keeping those in $G \ominus g$ fixed. The mapping $\phi_{g'|g}(R_{\ominus g})$, denoted simply as $\phi_g(\cdot)$ when $g' = g$, returns the multiplier $R_{g'|g}$ that results on g' from this partial markup equilibrium. A vector α is a markup equilibrium on G if and only if $R_a(\alpha)$ is a fixed point of the function $h_a : h_a(R_a) \to \phi_a(\phi_{\ominus a}(R_a))$ for all $a \in A_G$. We prove that $h_a(R_a)$ has a unique fixed point if and only if $Q(G \ominus a) \geqslant 2$. The forward direction follows because when $Q(G \ominus a) = 1$, (6.5) ensures unbounded markups, preventing the existence of an equilibrium. For the reverse direction, we show inductively that the following holds for each subnetwork $g \subseteq G$:

(i) If $R_{\ominus g} < \infty$, g admits a unique partial markup equilibrium.
(ii) $0 < \phi_g'(R_{\ominus g}) < \phi_g(R_{\ominus g})/R_{\ominus g}$.
(iii) If $Q(g) \geqslant 2$ then $\lim_{R_{\ominus g} \to \infty} \phi_g(R_{\ominus g})/R_{\ominus g} = 0$. Otherwise, $\lim_{R_{\ominus g} \to \infty} \phi_g(R_{\ominus g})/R_{\ominus g} = 1$.

Note that (i) makes ϕ_g well-defined, and since the response function in (6.5) is C^1, we invoke the Implicit Function Theorem to show ϕ_g' is well-defined. The basic case of the induction consists on producers. For them, the three properties follow from writing (6.5) as $R_a = 2c_a + R_{\ominus a}$. For the induction, see that each subnetwork g is built up from producers through the $S(\cdot)$ and $P(\cdot)$ operations. As $Q(G) \geqslant 3$, we have

$I(g) \geqslant 3$ for all $g \subseteq G$. We thus address only those operations preserving $I(g) \geqslant 3$ (when $I(g')$ and $I(g'') \geqslant 3$):

Case 1: g = P{g', g''}: Observe that $I(g) \geqslant 3$ and $Q(g) \geqslant 2$. We address the partial equilibrium on g with $R_{\ominus g}$ fixed to r. Here, we establish a fixed point $R_{g'|g}$ of $h_{g'|g} : R \to \phi_{g'}([1/\phi_{g''}([1/R + 1/r]^{-1}) + 1/r]^{-1})$, whose form is derived from the relation $R_{\ominus g'} = [1/R_{g''} + 1/R_{\ominus g}]^{-1}$. See that the argument of $\phi_{g'}(\cdot)$ in this expression is bounded above by r. We can thus restrict $R_{g'|g}$ to fall within $[0, \phi_{g'}(r)]$. A fixed point must then exist, by Brouwer's fixed point theorem. Furthermore, see that

$$\frac{\partial h_{g'|g}(R)}{\partial R} = \phi'_{g'}\Big(\Big[\frac{1}{\phi_{g''}([\frac{1}{R} + \frac{1}{r}]^{-1})} + \frac{1}{r}\Big]^{-1}\Big)\Big[\frac{1/\phi_{g''}([\frac{1}{R} + \frac{1}{r}]^{-1})}{1/\phi_{g''}([\frac{1}{R} + \frac{1}{r}]^{-1}) + \frac{1}{r}}\Big]^2 \phi'_{g''}$$
$$\Big(\Big[\frac{1}{R} + \frac{1}{r}\Big]^{-1}\Big)\Big[\frac{\frac{1}{R}}{\frac{1}{R} + \frac{1}{r}}\Big]^2.$$

From the inductive assumption of (ii), this can be strictly upper bounded by

$$\frac{1}{R}\phi_{g'}\Big(\Big[\frac{1}{\phi_{g''}([\frac{1}{R} + \frac{1}{r}]^{-1})} + \frac{1}{r}\Big]^{-1}\Big)\Big(\frac{1/\phi_{g''}([\frac{1}{R} + \frac{1}{r}]^{-1})}{1/\phi_{g''}([\frac{1}{R} + \frac{1}{r}]^{-1}) + \frac{1}{r}}\Big)\Big(\frac{\frac{1}{R}}{\frac{1}{R} + \frac{1}{r}}\Big).$$

The last expression is at most 1 at a fixed point, and is decreasing in R because of (ii). For these reasons, there exists exactly one fixed point. We compute $\phi'_{g'|g}(r)$ and $\phi'_{g''|g}(r)$ using the implicit functions arising from the fixed point equations that characterize $R_{g'}$ and $R_{g''}$ for each r. Applying the lower and upper bounds from (ii) gives $0 < \phi'_{g'|g}(r) < \phi_{g'|g}(r)/r$ and $0 < \phi'_{g''|g}(r) < \phi_{g''|g}(r)/r$. Claim (ii) now follows as $\phi'_g(r) = \phi_g(r)^2[\phi'_{g'|g}(r)/\phi^2_{g''|g}(r) + \phi'_{g''|g}(r)/\phi^2_{g'|g}(r)] < \phi_g(r)/r$. For claim (iii), it is sufficient that both $L_1 := \lim_{r\to\infty} \phi_{g'|g}(r)/r = 0$ and $L_2 := \lim_{r\to\infty} \phi_{g''|g}(r)/r = 0$. Note that $\lim_{r\to\infty} \phi_{g'|g}(r)$ is bounded if and only if $\lim_{r\to\infty} \phi_{g''|g}(r)$ is bounded. If both are bounded, then L_1 and L_2 are 0 as required. Otherwise, $\lim_{R_{\ominus g'}\to\infty} \phi_{g'}(R_{\ominus g'})/R_{\ominus g'} = \lim_{r\to\infty} \phi_{g'|g}(r)(1/r + 1/\phi_{g''|g}(r)) = L_1 + L_1/L_2$, which is less than or equal to 1 by inductive assumption (iii). Similarly, $\lim_{R_{\ominus g''}\to\infty} \phi_{g''}(R_{\ominus g''})/R_{\ominus g''} = L_2 + L_2/L_1 \leqslant 1$. The only possibility to satisfy both inequalities simultaneously is for L_1 and L_2 to be zero (with the limits converging at the same rate).

Case 2: g = S {g', g'' }, $\max\{Q(g'), Q(g'')\} \geqslant 2$: Observe that $I(g) \geqslant 3$ and $Q(g) = \min\{Q(g'), Q(g'')\}$. Without loss of generality, assume $Q(g'') \geqslant 2$, As before, we address the partial equilibrium on g with $R_{\ominus g}$ fixed to r, and establish a fixed point $R_{g'|g}$ of $h_{g'|g} : R \to \phi_{g'}(\phi_{g''}(R + r) + r)$. See that $\partial h_{g'|g}(R)/\partial R = \phi'_{g'}(r + \phi_{g''}(r + R))\phi'_{g''}(r + R) \leqslant [\phi_{g'}(r + \phi_{g''}(r + R))/(r + \phi_{g''}(r + R))][\phi_{g''}(r + R)/(r + R)]$. Both terms in the upper bound are decreasing in R by (ii), and the second tends to zero when $R \to \infty$ by (iii). Hence, at some finite $\hat{R}$, $h'_{g'|g}(R) < 1\ \forall R > \hat{R}$, guaranteeing a fixed point. Uniqueness follows as $\partial h_{g'}(R)/\partial R$ and its upper bound cannot be larger than 1 at and to the right of any fixed point. As in the parallel case, we obtain $0 < \phi'_{g'|g}(r) < \phi_{g'|g}(r)/r$ and $0 < \phi'_{g''|g}(r) < \phi_{g''|g}(r)/r$, which provides $\phi'_g(r) =$

$\phi'_{g'|\ominus g}(r) + \phi'_{g''|\ominus g}(r) < \phi_g(r)/r$. To prove (iii), we define $L_1 := \lim_{r\to\infty} \phi_{g'|g}(r)/r$ and $L_2 := \lim_{r\to\infty} \phi_{g''|g}(r)/r$. If $Q(g) = 1$, we must show $L_1 + L_2 = 1$. From inductive assumption of (iii), $\lim_{R_{\ominus g'}\to\infty} \phi_{g'}(R_{\ominus g'})/R_{\ominus g'} = \lim_{r\to\infty} \phi_{g'|g}(r)/(r + \phi_{g''|g}(r)) = [L_2/L_1 + 1/L_1]^{-1} = 1$ and, similarly, $\lim_{R_{\ominus g''}\to\infty} \phi_{g''}(R_{\ominus g''})/R_{\ominus g''} = [L_1/L_2 + 1/L_2]^{-1} = 0$. Starting with the latter, we must have that either $L_2 = 0$ or $L_1 = \infty$. If $L_2 = 0$, the former implies that $L_1 = 1$, proving that the sum is 1 as needed. Otherwise, $L_1 = \infty$, which is a contradiction because L_2 would also be unbounded and $0 \neq 1^{-1}$. In the remaining case, we must show $L_1 + L_2 = 0$ when $Q(g) \geqslant 2$. Indeed, when $[L_2/L_1 + 1/L_1]^{-1} = [L_1/L_2 + 1/L_2]^{-1} = 0$, both limits must be zero, as required.

Finally, note that the inductive hypothesis applies to G itself. Furthermore, for any producer a, $h'_a(R_a) = \phi'_a(\phi_{\ominus a}(R_a))\phi'_{\ominus a}(R_a) < [\phi_a(\phi_{\ominus a}(R_a))/\phi_{\ominus a}(R_a)][\phi_{\ominus a}(R_a)/R_a]$. Both fractions in the upper bound are decreasing in R_a, and their product is equal to 1 at a fixed point of $h_a(\cdot)$. Thus, $h'_a(R_a) < 1$ at, and to the right of, any fixed point. Therefore, $h_a(R_a) = R_a$ can be satisfied by at most one point. □

Define $\tilde{\Gamma}(\alpha) : \mathbb{R}^n \to \mathbb{R}^n$ such that $\tilde{\Gamma}_a(\alpha) = \Gamma_a(\alpha_{-a})$. By bounding α, we restrict the image of $\tilde{\Gamma}(\alpha)$ to a compact set, assuring the existence of a markup equilibrium. We observe further that $\Gamma_a(\alpha_{-a})$ is increasing in α_b for all $b \neq a$. As a result, any sequence $\{\alpha^\tau\}$ with $\alpha^\tau = \tilde{\Gamma}(\alpha^{\tau-1})$ will be increasing element-wise. Starting at $\alpha^0_a = 2$ 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 a markup equilibrium in this way for any game that satisfies the 3-edge-connectivity condition.

6.4 General networks

As the following example demonstrates, Theorem 6.3 does not immediately generalize to networks that are not series-parallel. We present a very simple network structure that is 3-edge-connected but not Series-Parallel for which no markup equilibrium exists.

Critically, when the network is not SP, we cannot guarantee that all producers are active in equilibrium. In the network depicted to the left of Fig. 3, 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.

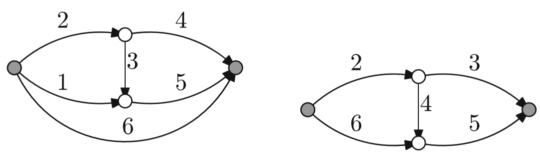

Fig. 3 *Left* a 3-edge-connected network that is not series-parallel. *Right* substitute network for producer 1

When this is 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 it from the analysis entirely, with no affect on the equilibrium. The remaining producers then constitute a series-parallel network that is not 3-edge-connected. There is no equilibrium in such a network, so producer 3 must be active in any equilibrium.

For producer 3 to be active, the price of the bundle $B_4 = \{2, 3, 5\}$, must be equal to that of $B_1 = \{1, 5\}$, $B_2 = \{2, 4\}$, and $B_3 = \{6\}$. For a given set of markups α, the consumption assignment f satisfies:

$$\begin{cases} f_1 R_1 + (f_1 + f_4) R_5 = (f_2 + f_4) R_2 + f_2 R_4 = f_3 R_6 \\ \quad = (f_2 + f_4) R_2 + f_4 R_3 + (f_1 + f_4) R_5, \\ f_1 + f_2 + f_3 + f_4 = 1. \end{cases} \tag{6.6}$$

Solving this system for f yields the consumption and production assignments for a second-stage equilibrium. After constructing the profit functions for each producer, we find that each producer's optimal markup is again of the form, $\Gamma_a(\alpha_{-a}) = 2 + R_{\ominus a}/c_a$, where $R_{\ominus a}$ is the price of an equilibrium assignment in a substitute network.

Structurally, the network to the left of Fig. 3 is entirely symmetric, in the sense that $G \ominus a$ has the same structure for any choice of a. The graph of $G \ominus 1$ is shown to the right of the figure, and the logic to follow will apply symmetrically to each producer's markup. For a given set of markups α_{-1}, we have that

$$R_{\ominus 1} = \begin{cases} \hat{f}_1 R_2 + \hat{f}_1 R_3 = \hat{f}_2 R_5 + \hat{f}_2 R_6 & \text{if } \hat{f}_3 = 0 \\ (\hat{f}_1 + \hat{f}_3) R_2 + \hat{f}_3 R_4 + (\hat{f}_2 + \hat{f}_3) R_5 & \text{if } \hat{f}_3 > 0, \end{cases}$$

where $\hat{f}$ is a consumption assignment satisfying $\hat{f}_1 + \hat{f}_2 + \hat{f}_3 = 1$. If $\hat{f}_3 = 0$, then $R_{\ominus 1} = ((R_2 + R_3)\hat{f}_1 + (R_5 + R_6)\hat{f}_2)/2 \geqslant \min\{R_2, R_3, R_5, R_6\}$. If $\hat{f}_3 > 0$, then $R_{\ominus 1} \geqslant R_2 \hat{f}_1 + R_4 \hat{f}_3 + R_5 \hat{f}_2 \geqslant \min\{R_2, R_4, R_5\}$. Employing the symmetric arguments, $R_a = \Gamma_a(\alpha_{-a}) c_a > \min_{b \in A_G}\{R_b\}$ for all $a \in A_G$, which is a contradiction. It follows that there are no markup equilibria for which producer 3 is active, and consequently, no markup equilibria in the market represented by the 3-edge-connected network, G.

7 Conclusion

There are several possible extensions that would be interesting to explore. These include more general costs structures, other forms of supply functions, more general demand structures, and further analysis of the networked case. In what follows we describe some of these pointing out specific problems of interest.

One direction is to allow for fixed costs and other types of cost functions, especially those usually found in electricity markets and other relevant applications. This seems particularly challenging when producers are asymmetric, for instance, when they have different fixed costs.

Although in this paper we have assumed inelastic demand, some of our results extend to the case of elastic demand. In particular, for series-parallel networks and linear per-unit cost functions, Lederman [29] studies necessary and sufficient conditions for the existence of equilibria. In particular he shows that that an equilibrium exists if and only if adding two parallel links between the source node and the sink node results in a 3-edge-connected network. Unfortunately for nonlinear costs we do not know conditions for existence of equilibria in series-parallel networks, because solving the second stage game in closed form is challenging.

In light of Corollary 3.6, it would be very interesting to bound the price of anarchy when producers are allowed to bid supply functions with a fixed constant term. Although we expect that in general this would lead to larger inefficiencies, it would apply to arbitrary supply function equilibria. On the other hand, we do not know the extent of the inefficiency of equilibria in the case of elastic demands with general cost functions—studied by Chau and Sim [13] for the second stage game. Another interesting avenue is to bound the price of anarchy in series-parallel networks. It is easy to extend our initial bound in Sect. 5 to the network case, for linear per-unit cost functions, however the tighter bounds do not seem to readily apply.

More general demand structures such as oligopsonies to model different consumers' market power, is another line of research that may prove fruitful. If we follow the analysis of Sect. 5 for the second-stage game, after the supply functions are fixed, both an equilibrium and a centralized assignment minimizing the total price paid coincide. This happens because the negative externalities are proportional to the marginal costs [16]. At this point it is natural to consider other market structures leading to potentially suboptimal assignments. This more general version of our model will have to consider two sources of inefficiency: the chosen supply functions that may distort costs; and the consumers that may be assigned suboptimally. In follow-up work, Kuleshov and Wilfong [28] study a SFE mechanism for an oligopoly where players simultaneously optimize the demand as well as the bundles they choose. Considering series-parallel markets and linear per-unit supply functions, they establish that the inefficiency induced by such mechanism is not too large.

Acknowledgments This research was partially funded by the Millenium Nucleus Information and Coordination in Networks ICM/FIC P10-024F, the Instituto Sistemas Complejos de Ingeniería at Universidad de Chile, by Fondecyt Chile grant 1130671, by the Center for International Business Education and Research at Columbia University, and by Conicet Argentina grant Resolución 4541/12. The authors wish to thank the associate editor and an anonymous referee for their constructive comments, Roberto Cominetti, Andy Philpott, and Gabriel Weintraub for stimulating discussions, and the participants of seminars at Columbia Business School, the School of Industrial and Systems Engineering of Georgia Tech, Ecole Polytechnique Fédéral de Lausanne, IESE Business School, and Universidad de Chile, for their questions, comments and suggestions.

References

1. Acemoglu, D., Bimpikis, K., Ozdaglar, A.: Price and capacity competition. Games Econ. Behav. **66**(1), 1–26 (2009)
2. Acemoglu, D., Ozdaglar, A.: Competition and efficiency in congested markets. Math. Oper. Res. **32**(1), 1–31 (2007)

3. Acemoglu, D., Ozdaglar, A.: Competition in parallel-serial networks. IEEE J. Sel. Areas Commun. **25**(6), 1180–1192 (2007)
4. Ahuja, R.K., Magnanti, T.L., Orlin, J.B.: Network Flows: Theory, Algorithms, and Applications. Prentice-Hall, Englewood Cliffs (1993)
5. Akgün, U.: Mergers with supply function. J. Ind. Econ. **52**(4), 535–545 (2004)
6. Anderson, E.J., Hu, X.: Finding supply function equilibria with asymmetric firms. Oper. Res. **56**(3), 697–711 (2008)
7. Anderson, E.J., Philpott, A.B.: Using supply functions for offering generation into an electricity market. Oper. Res. **50**(3), 477–489 (2002)
8. Baldick, R., Grant, R., Kahn, E.: Theory and application of linear supply function equilibrium in electricity markets. J. Regul. Econ. **25**(2), 143–167 (2004)
9. Baldick, R., Hogan, W.: Stability of supply function equilibria: Implications for daily versus hourly bids in a poolco market. J. Regul. Econ. **30**(2), 119–139 (2006)
10. Beckmann, M.J., McGuire, C.B., Winsten, C.B.: Studies in the Economics of Transportation. Yale University Press, New Haven (1956)
11. Bowley, A.: The Mathematical Groundwork of the Economics. Oxford University Press, Oxford (1924)
12. Casadesus-Masanell, R., Nalebuff, B., Yoffie, D.B.: Competing complements. Harvard Business School Working Paper, No. 09-009 (2010)
13. Chau, C.K., Sim, K.M.: The price of anarchy for non-atomic congestion games with symmetric cost maps and elastic demands. Oper. Res. Lett. **31**(5), 327–334 (2003)
14. Chawla, S., Roughgarden, T.: Bertrand competition in networks. In: Monien, B., Schroeder, U. (eds.) Algorithmic Game Theory, First International Symposium (SAGT 2008), volume 4997 of Lecture Notes in Computer Science. Springer (2008)
15. Cournot, A.: Recherches sur les Principes Mathematiques de la Theorie des Richesses. Hachette, Paris (1838)
16. Dafermos, S.C., Sparrow, F.T.: The traffic assignment problem for a general network. J. Res. U.S Natl. Bureau Stand. **73B**, 91–118 (1969)
17. Engel, E., Fischer, R., Galetovic, A.: Toll competition among congested roads. Top. Econ. Anal. Policy **4**(1). Article 4 (2004)
18. Graham, R.L., Knuth, D.E., Patashnik, O.: Concrete Mathematics: A Foundation for Computer Science, 2nd edn. Addison-Wesley, Reading, MA (1994)
19. Granot, D., Yin, S.: Competition and cooperation in decentralized push and pull assembly systems. Manage. Sci. **54**(3), 733–747 (2008)
20. Grossman, S.: Nash equilibrium and the industrial organisation of markets with large fixed costs. Econometrica **49**, 1149–1172 (1981)
21. Hendricks, K., McAfee, R.P.: A theory of bilateral oligopoly. Econ. Inq. **48**(2), 391–414 (2010)
22. Jiang, L., Wang, Y.: Supplier competition in decentralized assembly systems with price-sensitive and uncertain demand. Manuf. Serv. Oper. Manage. **12**(1), 93–101 (2010)
23. Johari, R., Tsitsiklis, J.N.: Network resource allocation and a congestion game. Math. Oper. Res. **29**(3), 407–435 (2004)
24. Johari, R., Tsitsiklis, J.N.: Parameterized supply function bidding: equilibrium and welfare. Oper. Res. **59**(5), 1079–1089 (2011)
25. Johari, R., Weintraub, G.Y., Van Roy, B.: Investment and market structure in industries with congestion. Oper. Res. **58**(5), 1303–1317 (2010)
26. Klemperer, P.D., Meyer, M.: Supply function equilibria in oligopoly under uncertainty. Econometrica **57**(6), 1243–1277 (1989)
27. Koutsoupias, E., Papadimitriou, C.H.: Worst-case equilibria. In: Meinel, C., Tison, S. (eds.) Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science (STACS), Trier, Germany, Volume 1563 of Lecture Notes in Computer Science, pp. 404–413. Springer, Heidelberg (1999)
28. Kuleshov, V., Wilfong, G.: Cournot and bertrand competition in two-sided markets. Preprint (2011)
29. Lederman, R.: Strategic models in supply network design. Ph. D. thesis, Graduate School of Business, Columbia University, New York, NY (2012)
30. Mas-Colell, A., Whinston, M.D., Green, J.R.: Microeconomic Theory. Oxford University Press, New York (1995)
31. McAfee, R.P., Williams, M.A.: Horizontal mergers and antitrust policy. J. Ind. Econ. **40**(2), 181–187 (1992)
32. Nocedal, J., Wright, S.J.: Numerical Optimization. Springer, Berlin (2006)

33. Papadimitriou, C., Valiant, G.: A new look at selfish routing. In: Chi-Chih Yao, A. (ed.) Innovations in Computer Science (ICS). Tsinghua University Press, Beijing, China (2010)
34. Perry, M.K., Porter, R.H.: Oligopoly and the incentive for horizontal merger. Am. Econ. Rev. **75**(1), 219–225 (1985)
35. Robson, A.: Implicit oligopolistic collusion is destroyed by uncertainty. Econ. Lett. **7**, 144–148 (1981)
36. Rudkevich, A., Duckworth, M., Rosen, R.: Modeling electricity pricing in a deregulated generation industry: the potential for oligopoly pricing in a poolco. Energy J. **19**(3), 19–48 (1998)
37. Sioshansi, R., Oren, S.: How good are supply function equilibrium models: an empirical analysis of the ERCOT balancing market. J. Regul. Econ. **31**, 1–35 (2007)
38. Turnbull, S.: Choosing duopoly solutions by consistent conjectures and by uncertainty. Econ. Lett. **13**, 253–258 (1983)
39. Vives, X.: Strategic supply function competition with private information. Econometrica **79**(6), 1919–1966 (2011)
40. Wang, Y.: Joint pricing-production decisions in supply chains of complementary products with uncertain demand. Oper. Res. **54**(6), 1110–1127 (2006)
41. Wichiensin, M., Bell, M.G.H., Yang, H.: Impact of congestion charging on the transit market: an inter-modal equilibrium model. Transp. Res. **41A**(7), 703–713 (2007)
42. Wilson, R.: Auctions of shares. Q. J. Econ. **94**, 675–689 (1979)
43. Xiao, F., Yang, H., Han, D.: Competition and efficiency of private toll roads. Transp. Res. **41B**(3), 292–308 (2007)
44. Yang, S., Hajek, B.: Revenue and stability of a mechanism for efficient allocation of a divisible good. Working paper (2006)