

TAXES AND MARKET POWER: A PRINCIPAL COMPONENTS APPROACH

ANDREA GALEOTTI, BENJAMIN GOLUB, SANJEEV GOYAL,
EDUARD TALAMÀS, AND OMER TAMUZ

ABSTRACT. Suppliers of differentiated goods make simultaneous pricing decisions, which are strategically linked. Because of market power, the equilibrium is inefficient. We study how a policymaker should target a budget-balanced tax-and-subsidy policy to increase welfare. A key tool is a certain basis for the goods space, determined by the network of interactions among suppliers. It consists of *eigenbundles*—orthogonal in the sense that a tax on any eigenbundle passes through *only* to its own price—with pass-through coefficients determined by associated eigenvalues. Our basis permits a simple characterization of optimal interventions. A planner maximizing consumer surplus should tax eigenbundles with low pass-through and subsidize ones with high pass-through. The *Pigouvian leverage* of the system—the gain in consumer surplus achievable by an optimal tax scheme—depends only on the dispersion of the eigenvalues of the matrix of strategic interactions. We interpret these results in terms of the network structure of the market.

1. INTRODUCTION

Imperfect competition has attracted renewed attention in recent years, as evidence accumulates of the rising importance of oligopolistic industries (De Loecker et al., 2020; Azar and Vives, 2021). As firms in such industries have market power, the market outcome is typically inefficient.¹ This creates the scope for targeted interventions that tax some suppliers/sectors and subsidize others to further a social objective. Surprisingly, very little is known about this classical Pigouvian problem for oligopolies. This paper addresses it through the lens of networks.

In our model, firms price their products strategically (in Bertrand competition), and demand arises from the behavior of a representative consumer. Some products may be substitutes and others may be complements. The structure of such relationships determines the strategic interactions among firms. This, in turn, determines the Nash equilibrium prices, how surplus is allocated among producers, and consumer surplus.

The implications of taxing or subsidizing the suppliers in these markets are complicated, since changes to firms' costs affect the prices and quantities of other firms' through the strategic competition firms face. We summarize the strategic interactions among suppliers in a *spillover matrix* D where D_{ij} specifies how the demand of firm i changes when the price of firm j changes. We think of D as describing a *network*

Date Printed. April 21, 2022.

Galeotti is at the London Business School, Golub is at Northwestern University, Goyal is at University of Cambridge, Talamàs is at IESE Business School, and Tamuz is at Caltech. This work was supported by the ERC grant 724356 (Galeotti), the Pershing Square Fund for Research on the Foundations of Human Behavior (Golub), the National Science Foundation (SES-1658940, SES-1629446, Golub & DMS-1944153, Tamuz), the Sloan Foundation (Tamuz) and the US-Israel Binational Science Foundation (2018397, Tamuz).

¹For recent empirical evidence on the size of the welfare loss see Pellegrino (2021) and Ederer and Pellegrino (2021).

structure on the firms, following a literature in network theory (e.g., Candogan et al. (2012) and Bloch and Querou (2013)), and analyze certain properties of this network to facilitate the study of optimal tax policies.

The spillover matrix induces a basis of eigenvectors (also called principal components) of the space of goods produced by the strategic firms. We call the vectors in this basis *eigenbundles*. This basis has three special properties. First, when the cost of one eigenbundle is changed (e.g., by taxation), the effect is to change equilibrium prices *only of that eigenbundle*. These eigenbundles therefore identify independent, or orthogonal, dimensions of the market, such that the costs of one eigenbundle do not affect the prices of others. Second, the price pass-throughs (i.e., the fraction of a cost increase that is passed to the consumer) associated with various eigenbundles can be calculated in terms of corresponding eigenvalues of the spillover matrix. Third, the market induces a ranking of pass-throughs: the eigenbundles with larger eigenvalues—which are more representative (in a precise sense) of the demand interdependence across products offered in the market—are those with smaller pass-throughs.

These properties facilitate simple expressions of the effect of tax-subsidy schemes on prices and welfare. That, in turn, permits intuitive characterization of interventions that maximize consumer surplus. We focus on optimal *small* changes to taxes. Theoretically, the analysis of optimal small interventions requires fewer assumptions about the structure of demand and technology at allocations far from current conditions than a global analysis would, and comes down to an analysis of marginal effects. Practically, planners may prefer small adjustments to the status quo rather than implementing a large change all at once—for example, due to political constraints, or because of risk aversion.

Our study of the small-intervention problem builds on a large literature on the so-called “tax reform approach” initiated by Feldstein (1976) and Diewert (1978), Dixit (1979) and Tirole and Guesnerie (1981). Our first contribution here is to formalize the incentives making the planner choose small interventions. We posit that reform is implemented with some noise and the policymaker is risk-averse over the welfare criterion of interest (e.g., consumer surplus). As the risk-aversion increases, the planner prefers smaller interventions. We formalize this problem and characterize optimal tax interventions.

We obtain the following economic insights. First, through the lens of our decomposition, the optimal intervention may be described as follows: High pass-through eigenbundles are subsidized, with revenue collected from low pass-through eigenbundles. Thus, taxes levied on the low pass-through eigenbundles are mostly paid by the firms, while the subsidies offered to the high pass-through eigenbundles are mostly passed on to the consumer. This results in increased consumer welfare. As we have mentioned, the ranking of eigenbundles by pass-throughs is the same order of the corresponding eigenvalues of the spillover matrix.

Second, we characterize the gain in surplus that a planner can achieve through a small tax intervention. This depends on the scale of the intervention that a planner is willing to implement, but beyond this it also depends on the network: some networks permit the same planner to achieve a much greater gain, offering a bigger “bang for the buck.” The name we give to this “bang for the buck” measure is the *Pigouvian leverage*.

Our third contribution is to show that the **Pigouvian leverage** depends in a simple way on the eigenvalues of the spillover matrix: It is proportional to the sample

variance of the eigenvalues of this matrix. An equivalent characterization is that the **Pigouvian leverage** is proportional to the sum of the squares of the off-diagonal elements of the spillover matrix, which is a natural measure of the intensity of strategic effects. The “overall magnitude” of spillovers thus determines how much scope there is for improving consumer surplus. Notably, the *signs* of strategic interactions do not matter for this measure. Whether firms’ prices are strategic complements or substitutes, it is only an aggregate measure of the intensity of strategic interactions that determines the Pigouvian leverage. This gives an economically substantive characterization of when planners can achieve large benefits for consumers, but also illustrates the value of the spectral approach in generating such economic insights.

We conclude by exploring **non-small optimal interventions**. To do this, we focus on the classical case of **linear demand**. The characterization is again in terms of the spectral decomposition of demand spillovers, and echoes the result in the small intervention regime. The optimal policy still collects tax revenue from the eigenbundles with low pass-through, where the impact on prices and output is relatively small, and allocates them—via subsidies—to the eigenbundles with high pass-through, where the impact on prices and output is relatively large. The fact that the deadweight loss from taxation increases as a quadratic function of its size determines the optimal size of this tax program.

1.1. Related literature. Our paper contributes to a literature on the structure and theoretical properties of market power. For an early theoretical paper see [Dixit \(1986\)](#); more recent studies include, for example, [Vives \(1999\)](#) and [Azar and Vives \(2021\)](#). A recent literature in macroeconomics and industrial organization uses network models of differentiated oligopoly, with models similar to the one studied here, to provide empirical estimates of welfare losses due to market power (see e.g., [Pellegrino \(2021\)](#) and [Ederer and Pellegrino \(2021\)](#)).²

Our paper makes two contributions to this literature. First, taking a network perspective, we provide a geometric approach to analyzing pass-through in the pricing game with market power. This builds on work emphasizing the value of pass-through as a conceptual tool (e.g., [Weyl and Fabinger \(2013\)](#) and [Miklos-Thal and Shaffer \(2021\)](#)), and shows how it can be described tractably and intuitively in markets with very rich heterogeneity. Second, motivated by the empirical research on welfare losses due to market power, we apply our decomposition to characterize policy interventions that maximize consumer surplus and shed light on the economic forces that make them most effective.

In taking a network approach to inefficiencies in market interactions among firms we also connect to work in macroeconomics and trade on production networks (e.g., [Baqae \(2018\)](#), [Grassi \(2017\)](#) and [Liu \(2019\)](#) and [Grassi and Sauvagnat \(2019\)](#)). This work highlights the importance of the interaction between market structure and production networks for a variety of outcomes, especially the amplification of productivity shocks. One important contrast is that those models take pricing distortions as given, rather than endogenizing prices as we do in a Bertrand model.

²See also [Elliott and Galeotti \(2019\)](#) for related arguments on how network methods can be useful for competition authorities in developing antitrust investigations.

More generally, our paper contributes to the theory of network interventions; prominent early contribution to this theory include [Borgatti \(2006\)](#) and [Ballester, Calvó-Armengol, and Zenou \(2006\)](#).³ In a recent paper, [Galeotti, Golub, and Goyal \(2020\)](#) study intervention in a certain canonical class of quadratic network games. They use the singular value decomposition of the interaction matrix for the study of (costly) interventions that alter the stand-alone marginal benefits of individual activity. Their intervention problem takes a stylized form that cannot (even approximately) accommodate the notion of a budget-balanced tax scheme. As a result, the characterizations of optimal policy in our setting are distinctive, for example the finding that the planner subsidizes some eigenbundles and taxes others. At a conceptual level, we introduce the idea of Pigouvian leverage—how much the planner can achieve with budget-balanced interventions. This notion, and the characterizations that we provide of how it relates to the market structure, are one of the main contributions of the present paper.

2. PRICE PASS-THROUGH AND OLIGOPOLY COMPETITION

Symbols denoting vectors and matrices are in bold. For any matrix M , the symbol m_{ij} stands for its element in the i th row and j th column, and M^T denotes its transpose. The symbol $\langle \mathbf{a}, \mathbf{b} \rangle$ denotes the dot product of \mathbf{a} and \mathbf{b} .

2.1. Market structure. The set of *goods* is $\mathcal{N} = \{1, 2, \dots, N\}$. Good i is produced by firm i . The firms engage in a Bertrand pricing game, simultaneously choosing prices $\mathbf{p} = (p_i)_{i \in \mathcal{N}}$. Firm i has profit function

$$\pi_i(\mathbf{p}) = q_i(p_i - c_i),$$

where q_i is the quantity demanded of good i , and c_i is the (constant) marginal cost of production.

There is a representative consumer with a quasi-linear utility function that depends on goods, described by a quantity vector \mathbf{q} , and a separate commodity called the numeraire, whose quantity is denoted by m :

$$\mathcal{U}(\mathbf{q}, m) = \mathcal{V}(\mathbf{q}) + m.$$

Here \mathcal{V} is a twice-differentiable, concave utility function.

Example 1. For a concrete example, the reader may keep in mind the classical case in which the utility of the representative consumer is quadratic in consumption:

$$\mathcal{V}(\mathbf{q}) = \boldsymbol{\beta}^T \mathbf{q} - \frac{1}{2} \mathbf{q}^T \mathbf{B} \mathbf{q},$$

where $\boldsymbol{\beta}$ is a vector with positive entries and \mathbf{B} is a symmetric, positive definite matrix.

Returning to the general case, the consumer chooses an optimal bundle, paying price p_i for commodity i in terms of the numeraire. We assume that $\mathcal{V}(\cdot)$ is such that, for

³The literature on this subject is very large. Other contributions of network intervention in models of information diffusion, advertising, and pricing include [Banerjee, Chandrasekhar, Duflo, and Jackson \(2013\)](#), [Belhaj and Deroian \(2017\)](#), [Bloch and Querou \(2013\)](#), [Candogan, Bimpikis, and Ozdaglar \(2012\)](#), [Demange \(2017\)](#), [Fainmesser and Galeotti \(2017\)](#), [Galeotti and Goyal \(2009\)](#), [Galeotti and Rogers \(2013\)](#), and [Leduc, Jackson, and Johari \(2017\)](#).

any price vector \mathbf{p} , there is a unique, interior demand profile solving the consumer's problem

$$\max_{\mathbf{q}} \mathcal{V}(\mathbf{q}) - \langle \mathbf{q}, \mathbf{p} \rangle. \quad (1)$$

Let $q_i(\mathbf{p})$ be the quantity demanded of good i at the profile of prices \mathbf{p} .

2.2. Equilibrium. We focus on an (interior pure-strategy Nash) equilibrium \mathbf{p}^* of the price-setting game between the firms. The first-order conditions that characterize an interior equilibrium are:

$$q_i(\mathbf{p}^*) + \frac{\partial q_i(\mathbf{p}^*)}{\partial p_i} (p_i^* - c_i) = 0 \quad \text{for all } i \in \mathcal{N}. \quad (2)$$

We take the functions q_i as fixed and take the parameters of the game to be the marginal costs \mathbf{c} (which will be modified by tax/subsidy interventions). We make the following assumptions for technical convenience.

Assumption 1. There is an open set \mathcal{C} of costs such that, for all $\mathbf{c} \in \mathcal{C}$,

- (1) an interior equilibrium \mathbf{p}^* exists and is unique;
- (2) the demand function $\mathbf{q}(\mathbf{p})$ is locally linear in a neighborhood of equilibrium prices \mathbf{p}^* ;
- (3) each $q_i(\mathbf{p})$ is strictly decreasing in each p_i .

We assume throughout the analysis that the status quo is at some $\mathbf{c} \in \mathcal{C}$, and we then perturb the market relative to that starting point. This model is a standard differentiated oligopoly (e.g., [Vives \(1999\)](#) and [Choné and Linnemer \(2020\)](#)). Firms produce products that can be complements and substitutes in consumption and, in turn, this determines strategic interactions across the firms.

2.3. Pass-through. Slightly abusing terminology, we often denote equilibrium prices \mathbf{p}^* and quantities $\mathbf{q}(\mathbf{p}^*)$ by \mathbf{p} and \mathbf{q} , respectively.

We are interested in how a change in production costs (e.g., because of taxation) passes through to equilibrium prices, and how it changes welfare. To this end, we introduce a dummy parameter ξ and posit that costs vary according to given functions $\mathbf{c}(\xi)$. We let \dot{x} denote $\frac{dx}{d\xi}$ for any variable x in the model. The vector

$$\dot{\mathbf{c}} = (\dot{c}_1, \dot{c}_2, \dots, \dot{c}_N),$$

which will be specified exogenously, determines the first-order response of the economy. To analyze this response, we will calculate the derivatives in ξ of other variables in the model.

Totally differentiating (2) around the equilibrium \mathbf{p} yields, using the local linearity assumption:

$$\sum_{j \in \mathcal{N}} \frac{\partial q_i(\mathbf{p})}{\partial p_j} \dot{p}_j + \frac{\partial q_i(\mathbf{p})}{\partial p_i} (\dot{p}_i - \dot{c}_i) = 0. \quad (3)$$

To write this equation in vector notation, we define

$$D_{ij}(\mathbf{p}) = \frac{\partial q_i(\mathbf{p})}{\partial p_j} \quad \forall i, j \in \mathcal{N},$$

which is element ij of the the Jacobian of the vector \mathbf{q} in the prices \mathbf{p} ; let \mathbf{D} be the respective $N \times N$ matrix.

Because demand arises from the preferences of a representative consumer (with twice-differentiable, concave utility), the matrix D is symmetric and negative semidefinite (Nocke and Schutz, 2017). We also introduce a normalization that is useful in the rest of our analysis; Appendix A shows that this normalization is without loss of generality, as it can be achieved by changing the units of quantity of each good. We summarize this in the following property:

Property A. D is a symmetric negative semidefinite matrix and it satisfies the normalization $D_{ii} = -1$ for each $i \in \mathcal{N}$.

Equation 3 becomes

$$\dot{p}_i = \dot{c}_i + \sum_{j \in \mathcal{N}} D_{ij} \dot{p}_j, \quad (4)$$

which illustrates how the the strategic relations between any two suppliers i and j are captured by the sign of D_{ij} . We shall say that two distinct goods i and j are *strategic substitutes* if D_{ij} is negative and *strategic complements* if D_{ij} is positive. (Note that the strategic spillover is “opposite” to the interaction of the goods in the consumer’s utility: if two goods are complements in consumption, they are strategic substitutes in the pricing game.) The prices pass-through can be expressed in matrix form as follows:

$$[I - D]\dot{\mathbf{p}} = \dot{\mathbf{c}}. \quad (5)$$

Note that when demand is linear (the utility of the representative consumer is quadratic in consumption as in Example 1), the pass-through equation 5 holds exactly when $\dot{\mathbf{c}}$ is understood as an arbitrary change.

2.4. Pass-through in terms of eigenbundles. We now express the pass-throughs of cost changes in a compact way by changing to a convenient basis of *eigenbundles*.

Property A implies that D is orthogonally diagonalizable. That is, there exists an $N \times N$ orthonormal matrix U such that

$$D = U\Sigma U^T,$$

where Σ is an $N \times N$ diagonal matrix whose ℓ^{th} diagonal element is the ℓ^{th} -largest eigenvalue of D , called σ_ℓ ; it is nonpositive because D is negative semidefinite. The ℓ^{th} column \mathbf{u}^ℓ is the eigenvector of D corresponding to σ_ℓ . We call this the ℓ^{th} *eigenbundle* of D . These vectors have norm 1 and are orthogonal to each other.⁴

For any $\mathbf{x} \in \mathbb{R}^N$, let $\underline{\mathbf{x}}$ denote $U^T \mathbf{x}$; that is, $\underline{\mathbf{x}}$ is the vector \mathbf{x} expressed in the basis U . We choose the signs of \mathbf{u}^ℓ so that the equilibrium quantities of eigenbundles are all positive, i.e., $q_\ell \geq 0$ for each $\ell \in \mathcal{N}$.

The *consumer surplus* associated to an price profile \mathbf{p} is defined to be

$$\mathcal{C} = \mathcal{V}(\mathbf{q}(\mathbf{p})) - \langle \mathbf{q}(\mathbf{p}), \mathbf{p} \rangle.$$

This is the net utility to the consumer after paying for the goods.

⁴Eigenbundles may have entries with both positive and negative signs and these signs can be interpreted as combining long holdings of goods (those with positive signs) and short holdings of goods (those with negative signs).

Proposition 1. Fixing any \dot{c} , the derivative of the equilibrium price of the ℓ^{th} eigenbundle has the following form:

$$\dot{p}_\ell = \lambda_\ell \dot{c}_\ell, \quad (6)$$

where

$$\lambda_\ell = \frac{1}{1 - \sigma_\ell}$$

is increasing in ℓ . The associated pass-through to consumer surplus is

$$\dot{C} = - \sum_\ell \lambda_\ell q_\ell \dot{c}_\ell.$$

Proof. From (5) we get $(\mathbf{I} - \mathbf{U}\Sigma\mathbf{U}^\top) \dot{\mathbf{p}} = \dot{\mathbf{c}}$. Multiplying both sides by \mathbf{U}^\top we get

$$\mathbf{U}^\top (\mathbf{I} - \mathbf{U}\Sigma\mathbf{U}^\top) \mathbf{U} \mathbf{U}^\top \dot{\mathbf{p}} = \mathbf{U}^\top \dot{\mathbf{c}},$$

that is, $\dot{\mathbf{p}} = (\mathbf{I} - \Sigma)^{-1} \dot{\mathbf{c}}$.

Turning to consumer surplus, using Equation 1 we have that

$$\begin{aligned} \dot{C} &= \sum_i \left(\frac{\partial \mathcal{V}(\mathbf{q})}{\partial q_i} \dot{q}_i - q_i \dot{p}_i - \dot{q}_i p_i \right) \\ &= \sum_i (p_i \dot{q}_i - q_i \dot{p}_i - \dot{q}_i p_i) \\ &= -\langle \dot{\mathbf{p}}, \mathbf{q} \rangle, \end{aligned}$$

where the second equality follows by the envelope theorem applied to Equation 1 (which implies that $\frac{\partial \mathcal{V}(\mathbf{q})}{\partial q_i} = p_i$). Changing to the basis of \mathbf{U} gives the result. \square

The usefulness of the eigenbundles is that—as Proposition 1 states—a change in the cost of one of them affects only its own prices with a certain coefficient λ_ℓ , called the *price pass-through*. The price pass-throughs of the eigenbundles are ordered according to their corresponding eigenvalues: the larger the ℓ^{th} eigenvalue σ_ℓ , the higher the pass-through from changes in the cost of the ℓ^{th} eigenbundle to its equilibrium price.

The following example illustrates the proposition.

2.4.1. Example. Consider an oligopoly model with three products. Product 1 and 2 are independent; product 3 is a complement to product 1 and to product 2. For example, product 1 is a video game, product 2 is a word processor, while product 3 is a computer. This is represented by the following matrix (which is normalized):

$$\mathbf{D} = - \begin{bmatrix} 1 & 0 & 1/\sqrt{2} \\ 0 & 1 & 1/\sqrt{2} \\ 1/\sqrt{2} & 1/\sqrt{2} & 1 \end{bmatrix}$$

Figure 1 depicts the network associated with \mathbf{D} (omitting self-loops). Note that even if product 1's demand is independent of product 2's price and vice-versa, a shock that affects, say, the price of product 1, will change product 3's price and, consequently, will affect the pricing strategy of the producer of product 2.

Proposition 1 tells us that we can take any shock in the economy and decompose it into three shocks—each proportional to one of the three eigenbundles of \mathbf{D} —and study the pass-through within each eigenbundle independently. We do this next. We will then compute the overall price pass-through of the shock to each of the three products as a combination of the price pass-throughs within each eigenbundle.




FIGURE 1. A three-product oligopoly.

The first eigenbundle of \mathbf{D} is $\mathbf{u}^1 = (1/2, 1/2, -1/\sqrt{2})$. A cost shock proportional to \mathbf{u}^1 changes the prices of the two independent goods in the same direction and changes the price of good 3 in the opposite direction. The strategic substitutability between these products amplifies this divergence in the price of the two independent products viz. product 3. Initial shocks are amplified. In fact, since the eigenvalue $\sigma_1 = 0$, we have that $\lambda_1 = 1$ and so there is a 100% price pass-through along this dimension.

The second eigenbundle of \mathbf{D} is $\mathbf{u}^2 = (1/\sqrt{2}, -1/\sqrt{2}, 0)$. A cost shock proportional to \mathbf{u}^2 moves prices of the two independent goods in opposite directions—without affecting the price of product 3. Moreover, as product 3 is complementary to both 1 and 2 (in a symmetric way) the opposite price changes of 1 and 2 keeps the price of producer 3 constant. We obtain that pass-through in this eigenbundle is like the (normalized) monopoly pass-through of $1/2$. Indeed, the second eigenvalue is $\sigma_2 = -1$, and so $\lambda_2 = 1/2$.

Finally, the third eigenbundle of \mathbf{D} is $\mathbf{u}^3 = (1/2, 1/2, 1/\sqrt{2})$. It ranks products in terms of their eigenvector centrality. A cost shock proportional to \mathbf{u}^3 will therefore comove the price of the three producers in the same direction. These effects are dampened by their strategic reaction, so the resulting pass-through to prices is relatively small. Indeed, the eigenvalue is $\sigma_3 = -2$ and so only $\lambda_3 = 1/3$ of the shock passes through to the price of the third eigenbundle.

We now derive the price implications of a specific marginal shock to costs—as well as the overall effect of this shock on consumer surplus. Consider a shock that marginally increases the cost of product 3, i.e., $\dot{\mathbf{c}} = \{0, 0, \dot{c}_3\}$. Then we have that $\dot{\underline{\mathbf{c}}} = \{-\dot{c}_3/\sqrt{2}, 0, \dot{c}_3/\sqrt{2}\}$ and $\dot{\underline{\mathbf{p}}} = \{-\dot{c}_3/\sqrt{2}, 0, \dot{c}_3/(3\sqrt{2})\}$.

The final price pass-throughs to the three products can be obtained by combining the different eigenbundles' pass-throughs with how represented each of the products is in them, i.e., $\dot{p}_i = \sum_{\ell} \mathbf{u}_i^{\ell} \dot{\underline{\mathbf{p}}}_{\ell}$. That is

$$\begin{aligned} \dot{p}_1 = \dot{p}_2 &= \frac{\dot{c}_3}{2} \left(\frac{1}{3\sqrt{2}} - \frac{1}{\sqrt{2}} \right) = -\frac{\dot{c}_3}{3\sqrt{2}} \\ \dot{p}_3 &= \frac{\dot{c}_3}{\sqrt{2}} \left(\frac{1}{3\sqrt{2}} + \frac{1}{\sqrt{2}} \right) = \frac{2\dot{c}_3}{3} \end{aligned}$$

Finally, the resulting change in consumer surplus is approximately the sum of the price change within each eigenbundle times the equilibrium demand of this eigenbundle. The first eigenbundle is shocked negatively, $\dot{\underline{c}}_1 = -\dot{c}_3/\sqrt{2}$, so the prices of the two independent goods decrease and the price of the complement good 3 increases, increasing consumer surplus by $\dot{c}_3 q_1/\sqrt{2}$. The second eigenbundle is not shocked, i.e., $\dot{\underline{c}}_2 = 0$, and so there is no effect on consumer surplus. The cost shock to the third eigenbundle $\dot{\underline{c}}_3 = \dot{c}_3/\sqrt{2}$ is positive and it increases prices and so it has a negative effect on

consumer surplus ($-\dot{c}_3 q_3 / (3\sqrt{2})$). Aggregating these three effects we obtain

$$\dot{c} = \frac{\dot{c}_3}{\sqrt{2}} \left(q_1 - \frac{1}{3} q_3 \right) = -\frac{\dot{c}_3}{3\sqrt{2}} \left(\frac{4}{\sqrt{2}} q_3 - (q_1 + q_2) \right)$$

If the initial equilibrium demand for the complementary product 3 is at least (roughly) 35% of the total quantity demanded of product 1 and product 2, then the increase in the cost of product 3 has a negative effect on consumer surplus; otherwise the effect is positive. In particular, when the demand for products 1 and 2 is relatively large, an increase in the cost of product 3 can—via the associated strategic response of the producers of products 1 and 2, who decrease prices—increase consumer welfare.

2.4.2. *Comments on the meaning of eigenbundles.* We now develop the intuition for the decomposition of pass-through into eigenbundles, building on the ideas shown in the example. For concreteness, we consider a case where all products are strategic substitutes ($D_{ij} \leq 0 \forall i, j$, corresponding to complements in consumption), and we focus on the first and the last eigenbundles.

The last eigenbundle u^n (the one associated with the highest eigenvalue in absolute value) has all entries positive. It corresponds to the eigenvector centralities in $-D$ of various products. Producer i is highly central—that is, u_i^n is high—when the good that she produces is highly complementary to other goods which are also highly complementary to other goods, etc.

Consider a shock that changes the producers' marginal costs proportionally to their eigenvector centralities. The first effect is to move the price of all producers in the same direction. Since producers sell complementary products, the underlying pricing game is one of strategic substitutes. Hence, the initial co-movement in price will be attenuated by producers' strategic responses. Since the first eigenbundle captures the extent to which each product is complementary to others taking into account the whole structure of feedback, this is the eigenbundle in which strategic responses most strongly dampen the price pass-through of a shock.

In contrast, the first eigenbundle u^1 reflects a very different aspect of the structure of D . Its entries have *opposite* signs for complementary goods insofar as possible.⁵ A shock that is proportional to u^1 typically changes the prices of complementary goods in opposite directions. Since prices of such goods are strategic *substitutes*, these initial price effects are amplified; second-order effects go in the same direction as first-order effects, reinforcing them further. So the corresponding pass-through to prices is large.

3. PIGUVIAN INTERVENTION: SMALL TAX REFORMS

Because firms have market power, the equilibrium will typically be inefficient. Thus there is scope for a planner to suitably choose taxes and subsidies at the firm level to improve welfare. We consider Pigouvian interventions that attempt to increase consumer surplus, but the analysis can be easily extended to other surpluses. We also focus on interventions that are:

⁵See Bramoullé et al. (2014) and Galeotti et al. (2020) for extended discussions of the structure of such vectors.

- **Budget neutral.** A planner with a loose budget constraint can subsidize all firms at a high rate and increase consumer surplus substantially. However, under a tight budget constraint, the planner will need to impose positive taxes on some firms in order to subsidise others. To understand these trade-offs we focus on the case that the planner must run a balanced budget; allowing the planner to run a small deficit will not change the qualitative nature of the results.
- **Small.** There are two reasons to focus on small interventions. One is realism, as an administrative authority may be unwilling to make large changes. Secondly, if changes are constrained to be small, only properties of the market local to the equilibrium outcome are relevant, and these are much more feasible to assess. In fact, by focusing on small interventions we can derive economic insights which do not depend on the specific parametrization of the demand function of the representative consumer. As we shall see these results capture properties of global interventions that can be derived in special cases; see Section 4 where we develop a non-small optimal intervention for the special case of linear demand.

We begin by giving a simple foundation under which a planner would indeed choose to enact small changes and then we characterise the optimal intervention.

Let $\boldsymbol{\tau} = (\tau_1, \dots, \tau_n)$ be the profile of per-unit taxes that are assessed. This means that if q_i units are produced of input i , then the planner collects $\tau_i q_i$ units of revenue, which may be negative if $\tau_i < 0$. The planner does not have perfect control over the actual taxes assessed. What the planner chooses is a *target* tax $\bar{\tau}_i$ for each i , and the actual tax assessed is

$$\tau_i = \bar{\tau}_i \eta_i,$$

where η_i is an independent random variable with mean 1, variance ν^2 , and support bounded by $S\nu^2$ for some large S .⁶ On average, the planner gets what she wants, but there are shocks that perturb the implemented tax some amount from its intended target. This τ_i represents *implementation noise*, which for technical reasons enters in a multiplicative way.

The timing is that the planner sets the profile $(\bar{\tau}_i)_i$, the implementation noise η_i is realized for every i , then equilibrium outcomes are determined, and payoffs are received.

We assume that the planner seeks to increase consumer surplus, but is risk-averse over this outcome. Let $\hat{\mathcal{C}}$ denote the change in consumer surplus, which is a random variable because the implemented tax is random. The simple specification we use is that the planner's value function depends on the mean and variance of $\hat{\mathcal{C}}$ as follows:

$$\mathcal{W} = \mathbb{E} [\hat{\mathcal{C}}] - \frac{a}{2} \text{Var} [\hat{\mathcal{C}}]. \quad (7)$$

Thus, the planner cares about improving consumer surplus but is risk-averse and dislikes changes that lead to a highly variable consumer surplus. The parameter a measures the extent of the planner's risk-aversion.

As we will see, when a is large, then the planner chooses a vector of taxes and subsidies whose norm is small. The risk-aversion in the planner's payoff function can be seen as capturing, for example, political considerations: if ex post the penalty for a harmful reform is greater than the reward for a successful one, then it is intuitive that

⁶The assumption of bounded support is a technical one that can be relaxed at the cost of some additional details to be dealt with in proofs.

a planner will be cautious. We will see that this means that the planner will pursue first-order gains that small tax schemes allow, while limiting their size to avoid exposure to risk. This feature also contributes to the tractability of the model, allowing us to capture the main economic insights in a closed-form solution of an optimization problem.

Denote by R the net revenue of the tax to the government. We require that the planner runs a balanced budget in expectation. Thus, the planner's problem, for a given level a of risk-aversion, is

$$\text{choose } \bar{\tau} \text{ to } \max \mathcal{W} \text{ subject to } \mathbb{E}[R] = 0. \quad (\text{T(a)})$$

Our result on the form of the optimal intervention involves a statistic of the spillover matrix, which we now define. Let $\text{var}[\boldsymbol{\sigma}]$ denote the sample variance of the eigenvalues $(\sigma_1, \dots, \sigma_N)$, i.e.,

$$\text{var}[\boldsymbol{\sigma}] = \frac{1}{N} \sum_{i=1}^N \sigma_i^2 - \left(\frac{1}{N} \sum_{i=1}^N \sigma_i \right)^2.$$

Since D is symmetric with $D_{ii} = -1$, it is easy to verify that

$$\text{var}[\boldsymbol{\sigma}] = \frac{1}{N} \sum_{i \neq j} D_{ij}^2.$$

Thus, $\text{var}[\boldsymbol{\sigma}]$ captures the magnitude of the total strategic interaction—either substitutes or complements—of the different producers.

As we now show, the eigenvalue variance plays a key role in characterizing both the form of the optimal intervention and its efficacy.

Theorem 1. For large enough a , the optimal policy taxes each eigenbundle ℓ with $\lambda_\ell < z$ and subsidizes each eigenbundle ℓ with $\lambda_\ell > z$, where z is the shadow price of the budget-balance constraint and is equal to:

$$z = \left[2 + \frac{1}{2} \text{var}[\boldsymbol{\sigma}] \right]^{-1}.$$

Furthermore, in the limit $a \rightarrow \infty$ of small interventions, the optimal policy sets target taxes $\bar{\tau}_\ell$ so that

- the tax raised from the ℓ^{th} eigenbundle satisfies

$$\lim_{a \rightarrow \infty} a q_\ell \bar{\tau}_\ell = (z - \lambda_\ell)(1 + \sigma_\ell)^2,$$

- the change in consumer surplus satisfies

$$\lim_{a \rightarrow \infty} a \widehat{C}^* = N \frac{\text{var}[\boldsymbol{\sigma}]}{4 + \text{var}[\boldsymbol{\sigma}]} \quad (8)$$

The proof appears in Appendix B. The following example illustrates the theorem. We then comment on the broad implications.

Example 2. Consider a generalization of the Example 3.2 in which product 1 and product 2 are independent, but where we set $D_{31} = D_{32} = g$ with $g \in [-1/\sqrt{2}, 1/\sqrt{2}]$. When $g = 0$, the three products are independent. When $g < 0$, product 3 is a complement to both products 1 and 2, as we discussed earlier in Example 2; when $g > 0$ product 3

is a substitute to product 1 and 2. (For example, we may take product 1 to be a music player, product 2 to be a digital camera, and product 3 to be a smartphone.) The property that D is negative semi-definite means that $|g| \leq 1/\sqrt{2}$.

In this case, we have the following expressions for the eigenvalues of D :

$$\sigma_1 = -1 + \sqrt{2}|g| \quad \sigma_2 = -1 \quad \sigma_3 = -1 - \sqrt{2}|g|.$$

The variance of the eigenvalues is $\text{var}[\sigma] = 4g^2/3$. Furthermore, the shadow price of the budget balance constraint is

$$z(g) = \frac{3}{6 + 2g^2},$$

and at the limit $a \rightarrow \infty$ of small intervention, the increase in consumer surplus is proportional to

$$\frac{g^2}{3 + g^2}.$$

When all three products are independent, $g = 0$, each eigenvalue is equal to -1 , which means that the pass-through for each eigenbundle is the normalized monopoly price pass-through, $1/2$. In this case, budget neutral interventions are ineffective because any gains in consumer surplus due to subsidizing some eigenbundles can only be achieved by taxing some eigenbundles and that will create an equal loss in consumer surplus. Because budget-neutral interventions are ineffective, the shadow price of the budget neutral constraint is large, $z(0) = 1/2$. Indeed, if we give a dollar to the planner, the planner will be able to pass through half of it to consumers.

As $|g|$ increases, the first eigenvalue becomes larger than -1 , the second eigenvalue remains constant and the third becomes lower than -1 . Now, the eigenbundles are heterogeneous with respect to their price pass-throughs: the price pass-through of the first eigenbundle becomes greater than $1/2$, $\lambda_1 = 1/(2 + \sqrt{2}|g|)$ and the price pass-through of the last eigenbundle becomes lower than $1/2$, $\lambda_3 = 1/(2 - \sqrt{2}|g|)$. The planner can now increase consumer surplus by taxing the third eigenbundle, which creates little reduction in consumer surplus, and use the tax revenue to subsidise the first eigenbundle, which creates large gains in consumer surplus. Indeed, the optimal intervention becomes more and more effective in increasing consumer surplus as $|g|$ increases and it reaches a (roughly) 15% increase when $|g| = 1/\sqrt{2}$. As the planner can do more and more by targeting taxes and subsidises, the shadow price of the budget neutral constraint is decreasing in $|g|$.

We now draw out some general implications. First, the eigenvalue variance $\text{var}[\sigma]$ is a key statistic that captures the extent to which targeted tax/subsidy interventions are effective in an oligopoly setting. As this number grows, the consumer surplus gain achieved by a planner with the same level of risk aversion increases. At the same time, as the planner can increase the objective without running deficits, the shadow price of the budget neutral constraint z decreases.

The basic intuition is that an increase in the eigenvalue variance of D means that the price pass-through of the different eigenbundles becomes more heterogeneous. The planner can then tax eigenbundles with low price pass-through, thereby create a small decrease in consumer surplus, and then redistribute these resources to eigenbundles with high pass-throughs, which leads to a large increase in consumer surplus.

Second, recall that we can write the eigenvalue variance of \mathbf{D} as $\text{var}[\boldsymbol{\sigma}] = \frac{1}{N} \sum_{i \neq j} D_{ij}^2$. When products are independent then \mathbf{D} is the identity matrix and so the eigenvalue variance is 0 and so the Pigouvian leverage is zero. In contrast, complement and substitute relations across products, manifested in non-zero off-diagonal entries in the matrix \mathbf{D} , lead to strategic interaction in firms pricing behavior, and create scope for planner intervention. In this regard, the Pigouvian leverage does not depend on whether products are complements or substitutes, or a mix. It is only the intensity of demand interdependence of each pair of products that matters. A planner can achieve the same amount of improvement to consumer surplus no matter how the signs of various D_{ij} are changed.

4. GLOBAL PIGOUVIAN POLICY: THE CASE OF LINEAR DEMAND

The aim of this section is to show that the qualitative insights on optimal interventions are robust and extend beyond the setting of small tax reforms. To do so, we consider the case where interventions are not noisy, i.e., the variance of the implementation noise parameters η_i is zero. In this case, the consumer surplus determined by a specific intervention is deterministic and the planner may very well choose large interventions.

We characterize the optimal intervention under the classical case of linear-quadratic specification of utility that gives rise to linear demands (Example 1). The consumer has (gross) utility for consuming a bundle \mathbf{q} of goods given by

$$\mathcal{V}(\mathbf{q}) = \boldsymbol{\beta}^\top \mathbf{q} - \frac{1}{2} \mathbf{q}^\top \mathbf{B} \mathbf{q}, \quad (9)$$

where $\boldsymbol{\beta} \in \mathbb{R}_+^N$ and \mathbf{B} is a given positive-definite matrix (Amir, Erickson, and Jin (2017), Choné and Linnemer (2020), and Vives (1999)). Under the assumption that the consumer has sufficient income, her optimal consumption choice induces linear demands $\mathbf{q} = \mathbf{B}^{-1}[\boldsymbol{\beta} - \mathbf{p}]$ with:

$$\partial q_i / \partial p_j = -(\mathbf{B}^{-1})_{ij} \quad \text{and} \quad \frac{\partial^2 q_i}{\partial p_i \partial p_j} = 0 \quad \text{for any } i, j \in \mathcal{N}. \quad (10)$$

In this case, the matrix $\mathbf{D} = -\mathbf{B}^{-1}$ and the normalization in Property A means that $D_{ii} = -1$ for all $i \in \mathcal{N}$. A standard calculation also implies that, under the consumer's optimal consumption, the consumer surplus is

$$\mathcal{C} = \frac{1}{2} \mathbf{q}^\top \mathbf{B} \mathbf{q}. \quad (11)$$

Finally, the firms' pricing game under linear demand leads to the following equilibrium prices:

$$[\mathbf{I} - \mathbf{D}]\mathbf{p} = -\mathbf{D}\boldsymbol{\beta} + \mathbf{c}.$$

Let $\boldsymbol{\tau} = (\tau_1, \dots, \tau_n)$ be the profile of per-unit taxes on products introduced by the planner. We seek to solve the following problem:

$$\begin{array}{lll} \text{Maximize}_{\boldsymbol{\tau}} & \frac{1}{2} \mathbf{q}^\top \mathbf{B} \mathbf{q} & \text{(Consumer surplus)} \\ \text{subject to} & \mathbf{q} = -\mathbf{D}(\boldsymbol{\beta} - \mathbf{p}) & \text{(Optimal consumption)} \\ & [\mathbf{I} - \mathbf{D}]\mathbf{p} = -\mathbf{D}\boldsymbol{\beta} + \mathbf{c} + \boldsymbol{\tau} & \text{(Equilibrium prices)} \\ & \langle \boldsymbol{\tau}, \mathbf{q} \rangle = 0 & \text{(Budget balance)} \end{array} \quad (12)$$

Proposition 2 summarises the solution to this problem.

Proposition 2. The optimal policy taxes each eigenbundle ℓ with $\lambda_\ell < z$, and subsidizes each eigenbundle ℓ with $\lambda_\ell > z$, where z is the shadow price of the budget balance constraint and it is the unique positive solution of the following equation:

$$\sum_{\ell} \left(\underline{q}_{\ell}^0 \right)^2 \frac{(z - \lambda_{\ell})}{(1 - \lambda_{\ell})(2z - \lambda_{\ell})^2} = 0. \quad (13)$$

For generic \underline{q}^0 , it satisfies $\max\{\lambda_1, \lambda_n/2\} \leq z \leq \lambda_n$. Finally, the optimal taxes/subsidizes are:

$$\tau_{\ell} = \frac{\underline{q}_{\ell}^0(z - \lambda_{\ell})}{(1 - \lambda_{\ell})(2z - \lambda_{\ell})} \quad \forall \ell \in \mathcal{N}. \quad (14)$$

The proof appears in Appendix C.

Consistently with the optimal small intervention solution, in order to maximize consumer surplus, the planner taxes the eigenbundles with a low price pass-through (high eigenvalues) and uses the revenue generated to subsidize the eigenbundles with high price pass-through (low eigenvalues). The threshold that distinguishes eigenbundles which are taxed from the ones that are subsidised is the shadow price z . Generically, the shadow price must be larger than $\max(\lambda_1, \lambda_n/2)$ and it must be smaller than λ_n . In other words, the largest and smallest eigenvalues of the spillover matrix \mathbf{D} determine an upper bound and a lower bound of the shadow price of public funds. This result is reminiscent of the role, in the small intervention case, of the sample variance of the eigenvalues in determining the value of z .

In relation to this, note also that, for the intervention to create an improvement in consumer surplus, products must have some degree of complementarity or substitutability. If all products are independent, i.e., $\mathbf{D} = \mathbf{I}$, then eigenbundles have equal price pass-through and, therefore, the optimal policy is to maintain the status quo, i.e., $\tau_{\ell} = 0$ for all $\ell \in \mathcal{N}$.

We now illustrate the optimal policy within the example with three products, returning to the setting of Section 2.4.1.

4.1. Example. We assume that product 1 and product 2 are approximately independent, i.e., $D_{12} = D_{21} \approx 0$.⁷ We also assume that product 3 is equally complementary to product 1 and 2, $D_{13} = D_{12} = g > 0$. In this case, the eigenbundles are independent of g and are the same as the ones reported in Example 2.4.1. Furthermore, the eigenvalues are as in Example 2:

$$\sigma_1 = 1 + \sqrt{2}|g| \quad \sigma_2 = 1 \quad \sigma_3 = 1 - \sqrt{2}|g|.$$

We set the marginal costs for the three products to be zero, the demand parameters in Equation 9 to be $\beta = (10, 10, 15)$, and the degree of complementarity to be $g = 0.5$. In this case,

- the initial demand for the three products is (1.51, 1.51, 4.2),
- the price pass-throughs of the three eigenbundles are (0.37, 0.5, 0.78) and
- the initial quantities of the eigenbundles are (4.48, 0, 1.44).

⁷For \mathbf{D} to be negative definite we need that $D_{12} = D_{21} \neq 0$. We report the results for the limit case in which $D_{12} = D_{21} \rightarrow 0$.

The shadow price of the budget constraint is $z = 0.59$ and, therefore, the optimal policy taxes eigenbundle 1, is neutral with respect to eigenbundle 2 (because the initial quantity is zero) and subsidizes eigenbundle 3. In particular, the tax revenue from eigenbundle 1 (which is the same as the subsidy expenditure for eigenbundle 3) is 6.3. The value of this intervention can be measured by comparing the obtained increase in consumer surplus relative to the status quo. We obtain that consumer surplus goes up from 9.5 to 10.8.

Consider a technological change that increases the complementarity between product 3 and products 1 and 2; for concreteness, suppose that g increases from 0.5 to 0.55. The difference between the largest and lowest eigenvalues increases and therefore the difference between the price pass-through of the two extreme eigenbundles increases as well: the price pass-through of the first eigenvalue decreases from 0.37 to 0.36, whereas the price pass-through of the third eigenvalue increases from 0.78 to 0.82. This implies that now taxing the first eigenbundle induces less distortion and subsidizing the third eigenbundle creates greater benefits. Indeed, we obtain that the technological change increases the shadow price z of public funds from 0.59 to 0.7 (meaning that the intervention is more valuable). These qualitative insights carry over when product 3 is a substitute to products 1 and 2 (e.g., $g < 0$).⁸

Translated back to the original product space, the logic of the optimal policy in this example is rather intricate, as we now discuss. When we are in the complements case, as in Section 2.4.1, the optimal policy strongly taxes products 1 and 2 (the software applications) and more lightly subsidizes product 3 (the computer). The purpose of such optimal policies is to leverage the strategic interactions between producers to induce a reduction in the price of the high-volume product 3. Taxing the software applications *indirectly* incentivizes the computer producer—via the resulting price increase in the prices of complements—to reduce its price. (Recall prices are strategic substitutes in this case.) This reinforces the direct effect of the subsidy to computers. The computer subsidy must be kept fairly low, as otherwise (via the strategic substitutability in prices) we would end up *lowering* the price of the applications, which would be counterproductive because of the aforementioned strategic substitutes.

Consider now the case where the products are substitutes, (e.g., $g = -0.5$): the music player (product 1), the camera (product 2), and the smartphone (product 3). Smartphones remain the highest-volume product under our parameterization. Now, the optimal policy taxes smartphones and subsidizes music players and cameras. This is counterintuitive, because one might think that the planner would want to subsidize the product that consumers buy the most of, to improve consumer surplus. The logic behind the optimal policy can again be understood via thinking about the strategic spillovers. The subsidies to the specialized products *indirectly* incentivize the smartphone producer—via the resulting price decrease in the specialized substitute products—to reduce its price. (The prices are strategic complements in this case.) This more than compensates for the effect of the tax directly levied on smartphones, which is used to raise revenue. Indeed, this scheme turns out to be better than subsidizing smartphones directly: if the planner were to implement such subsidies, they would have to be offered on the whole (large) quantity of smartphones sold, and end up

⁸However, the fact that prices are now strategic complements implies that the first eigenbundle (i.e., in which costs shocks are dampened the most) is now what before was the third eigenbundle $(0.5, 0.5, -1/\sqrt{2})$, and the third one (i.e., in which costs shocks are dampened the least) is now what before was the first eigenbundle $(0.5, 0.5, 1/\sqrt{2})$.

being quite expensive per unit of welfare generated compared to the optimal intervention.

In the principal component decomposition, these intricate policies both correspond to the natural policy of subsidizing the high pass-through eigenbundles and taxing the low pass-through ones.

5. CONCLUSION

The paper studies firms interacting strategically in an environment with rich heterogeneity. We bring a network perspective to the interactions between the different firms. The key contribution is to decompose any cost change into components that have a very convenient form: for these components, there is an unambiguous ranking of pass-throughs from cost changes to equilibrium prices. We show how a policy-maker can use this ranking to design a tax policy that maximizes consumer surplus: The optimal policy leverages the strategic interactions among producers (in particular, that taxes on some dimensions have a higher pass-through to equilibrium prices) to tax some producers in order to use the tax revenue to subsidize other producers. The effectiveness of such interventions is greater in markets with strong demand interdependence across products, as summarised by the eigenvalue variance of the matrix of spillovers D .

REFERENCES

- Rabah Amir, Philip Erickson, and Jim Jin. 2017. On the microeconomic foundations of linear demand for differentiated products. *Journal of Economic Theory* 169 (2017), 641–665.
- J. Azar and X. Vives. 2021. General equilibrium oligopoly and ownership structure. *Econometrica* 89, 3 (2021), 999–1048.
- C. Ballester, A. Calvó-Armengol, and Y. Zenou. 2006. Who's who in Networks. Wanted: the Key Player. *Econometrica* 74, 5 (2006), 1403–1417.
- Abhijit Banerjee, Arun G Chandrasekhar, Esther Duflo, and Matthew O Jackson. 2013. The diffusion of microfinance. *Science* 341, 6144 (2013), 1236498.
- David Reza Baqaee. 2018. Cascading Behavior in Production Networks. *Econometrica* 86, 5 (2018), 1819–1838.
- Mohamed Belhaj and Frederic Deroian. 2017. Contracting on Networks. (2017). Mimeo.
- Francis Bloch and Nicolas Querou. 2013. Pricing in Social Networks. *Games and Economic Behavior* 80, 4 (2013), 263–281.
- Steven Borgatti. 2006. Identifying sets of key players in a social network. *Computational and Mathematical Organization Theory* 12 (2006), 21–34.
- Yann Bramoullé, Rachel Kranton, and Martin D'amours. 2014. Strategic interaction and networks. *The American Economic Review* 104, 3 (2014), 898–930.
- Ozan Candogan, Kostas Bimpikis, and Asuman Ozdaglar. 2012. Optimal Pricing in Networks with Externalities. *Operations Research* 60, 4 (2012), 883–905.
- Philippe Choné and Laurent Linnemer. 2020. Linear demand systems for differentiated goods: Overview and user's guide. *International Journal of Industrial Organization* (2020), 102663.
- Jan De Loecker, Jan Eeckhout, and Gabriel Unger. 2020. The rise of market power and the macroeconomic implications. *The Quarterly Journal of Economics* 135, 2 (2020), 561–644.

- Gabrielle Demange. 2017. Optimal Targeting Strategies in a Network under Complementarities. *Games and Economic Behaviour* 105 (2017), 84–103.
- W Erwin Diewert. 1978. Optimal tax perturbations. *Journal of Public Economics* 10, 2 (1978), 139–177.
- Avinash Dixit. 1979. Price changes and optimum taxation in a many-consumer economy. *Journal of Public Economics* 11, 2 (1979), 143–157.
- Avinash Dixit. 1986. Comparative statics for oligopoly. *International economic review* (1986), 107–122.
- Florian Ederer and Bruno Pellegrino. 2021. A Tale of Two Networks: Common Ownership and Product Market Rivalry. *Working Paper, University of Maryland* (2021).
- Matthew Elliott and Andrea Galeotti. 2019. The role of networks in antitrust investigations. *Oxford Review of Economic Policy* 35, 4 (2019), 614–637.
- Itay Fainmesser and Andrea Galeotti. 2017. Pricing network effects. *Review of Economic Studies* 83 (2017), 165–198. Issue 1.
- Martin Feldstein. 1976. On the theory of tax reform. *Journal of Public Economics* 6, 1-2 (1976), 77–104.
- Andrea Galeotti, Benjamin Golub, and Sanjeev Goyal. 2020. Targeting interventions in networks. *Econometrica* 88, 6 (2020), 2445–2471.
- Andrea Galeotti and Sanjeev Goyal. 2009. Influencing the Influencers: a theory of strategic diffusion. *The Rand Journal of Economics* 40, 3 (2009), 509–532.
- A. Galeotti and B. W. Rogers. 2013. Strategic Immunization and Group Structure. *American Economic Journal: Microeconomics* 5, 2 (2013), 1–32.
- Basile Grassi. 2017. IO in I-O: Size, Industrial Organization, and the Input-Output Network Make a Firm Structurally Important. *Working Paper, Bocconi University and IGIER* (2017).
- Basile Grassi and Julien Sauvagnat. 2019. Production Networks and Economic Policy. *Oxford Review of Economic Policy* 35, 4 (2019), 638–677.
- Matt V Leduc, Matthew O Jackson, and Ramesh Johari. 2017. Pricing and referrals in diffusion on networks. *Games and Economic Behavior* 104 (2017), 568–594.
- Ernest Liu. 2019. Industrial Policies in Production Networks. *Quarterly Journal of Economics* 134, 4 (2019), 1883–1948.
- Jeanine Miklos-Thal and Greg Shaffer. 2021. Pass-through as an economic tool: on exogenous competition, social incidence, and price discrimination. *Journal of Political Economy* 129, 1 (2021), 323–335.
- Volker Nocke and Nicolas Schutz. 2017. Quasi-linear integrability. *Journal of Economic Theory* 169 (2017), 603–628.
- Bruno Pellegrino. 2021. Product Differentiation and Oligopoly: a Network Approach. *Working Paper, University of Maryland* (2021).
- Jean Tirole and Roger Guesnerie. 1981. Tax reform from the gradient projection viewpoint. *Journal of Public Economics* 15, 3 (1981), 275–293.
- Xavier Vives. 1999. *Oligopoly pricing: old ideas and new tools*. MIT press.
- E Glen Weyl and Michal Fabinger. 2013. Pass-through as an economic tool: Principles of incidence under imperfect competition. *Journal of Political Economy* 121, 3 (2013), 528–583.

Appendix

APPENDIX A. NORMALIZATION OF SPILLOVER MATRIX

For any function $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$, let $\mathcal{D}f(\mathbf{x})$ be the Jacobian matrix of the function evaluated at $\mathbf{x} \in \mathbb{R}^n$, whose (i, j) entry is $\partial f_i / \partial x_j$, where f_i denotes coordinate i of the function.

Here we will be explicit about distinguishing quantity variables \mathbf{q} from the corresponding demand function; to this end, we will write the function as q .

Consider the change of coordinates for quantities given by $\tilde{q}_i = r_i q_i$. Keeping units of money fixed, **the corresponding prices are** $\tilde{p}_i = p_i / r_i$. Let \mathbf{R} be the diagonal matrix whose (i, i) entry is r_i . With these new units, we can define a function

$$\tilde{q}(\tilde{\mathbf{p}}) = \mathbf{R}q(\mathbf{R}\tilde{\mathbf{p}})$$

and by the chain rule we have that

$$\mathcal{D}\tilde{q}(\tilde{\mathbf{p}}) = \mathbf{R}[\mathcal{D}q(\mathbf{R}\tilde{\mathbf{p}})]\mathbf{R}.$$

For a given demand function $q : \mathbb{R}^N \rightarrow \mathbb{R}^N$, recall \mathbf{D} is defined to be $\mathcal{D}q(\mathbf{p}^*)$, where \mathbf{p}^* are equilibrium prices, uniquely determined under our maintained assumptions. Here we will be explicit about the demand function and write \mathbf{D}^q for $\mathcal{D}q(\mathbf{p}^*)$. It follows from this and the above paragraph that

$$\mathbf{D}^{\tilde{q}} = \mathbf{R}\mathbf{D}^q\mathbf{R}.$$

Now set $r_i = 1/\sqrt{|D_{ii}^q|}$. It is clear from the above formula that $\mathbf{D}^{\tilde{q}}$ has -1 on the diagonal.

Thus, under a suitable choice of units, the matrix \mathbf{D} may be assumed to have diagonal -1 .

APPENDIX B. PROOFS FOR THE SMALL-INTERVENTION ANALYSIS

Denote by \hat{x} the change in any endogenous variable x relative to the initial equilibrium for a given realization of random taxes.

We will state a useful lemma:

Lemma 1. The pass-through from the marginal tax profile $\boldsymbol{\tau}$ to consumer surplus, in the principal component basis, is

$$\hat{\mathcal{C}} = -\sum_{\ell} \lambda_{\ell} \underline{q}_{\ell} \underline{\tau}_{\ell} + \mathcal{O}(\|\boldsymbol{\tau}\|^2). \quad (15)$$

Proof of Lemma 1. From Equation 1, and using the envelope theorem, the change in the utility of the representative consumer is

$$\hat{\mathcal{C}} = -\langle \hat{\mathbf{p}}, \mathbf{q} \rangle + \mathcal{O}(\|\hat{\mathbf{p}}\|^2).$$

Since \mathbf{U} is orthonormal, we have that $\hat{\mathcal{C}} = -\langle \hat{\mathbf{p}}, \mathbf{q} \rangle + \mathcal{O}(\|\hat{\mathbf{p}}\|^2)$. The conclusion follows as in Proposition 1. We use the bounded-support assumption on η to ensure that as $\|\hat{\mathbf{p}}\|^2$ becomes small, the comparative statics of the prices are characterized exactly by Proposition 1. \square

Recall that the planner's problem is

$$\text{choose } \bar{\tau} \text{ to } \max \mathcal{W} \text{ subject to } \mathbb{E}[R] = 0 \quad (\mathbf{I}(a))$$

Consider an arbitrary small tax profile τ and let \tilde{q} be the new quantity and \tilde{p} be the new price. Now, by equation (2), under the normalization assumed by Property A, we have

$$\tilde{q} = \tilde{p} - c - \tau.$$

Then the budget-balance condition can be transformed as follows:

$$\begin{aligned} 0 &= \mathbb{E} \left[\sum_i \tau_i \tilde{q}_i \right] \\ &= \mathbb{E} \left[\sum_i \tau_i (q_i + \dot{p}_i - \tau_i) \right] \\ &= \mathbb{E} \left[\sum_i (\tau_i q_i + \tau_i \dot{p}_i - \tau_i^2) \right] \end{aligned}$$

Transforming to the principal component basis, we get a clean formula in terms of the principal components for the budget-balance constraint:

$$\mathbb{E} \left[\sum_{\ell} (\tau_{\ell} \underline{q}_{\ell} + \lambda_{\ell} \tau_{\ell}^2 - \tau_{\ell}^2) \right] = 0 \quad (16)$$

Now, in order to write the maximand of the planner's problem in a nice way, we will recall two things. First, [Lemma 1](#) permits us to write

$$\hat{C} = - \underbrace{\sum_{\ell} \lambda_{\ell} \underline{q}_{\ell} \tau_{\ell}}_{\hat{C}_{\text{FO}}} + \mathcal{O}(\|\tau\|^2).$$

Here we have defined \hat{C}_{FO} to be the first-order term in the change. Second, recall our assumption $\tau_i = \bar{\tau}_i \eta_i$, where η_i has a mean of 1. Assume without loss of generality⁹ that the variance of implemented taxes is $\nu^2 = 1$. We may write

$$\text{Var}[\hat{C}_{\text{FO}}] = \sum_{\ell} \lambda_{\ell}^2 \underline{q}_{\ell}^2 \bar{\tau}_{\ell}^2.$$

We will solve an auxiliary problem of maximizing the planner's utility as if the change in consumer surplus were only the first-order change. That is, we will maximize $\mathcal{W}(\hat{C}_{\text{FO}})$ subject to the budget constraint in expectation. At the end, we will show the approximation of replacing \hat{C} by \hat{C}_{FO} does not change the characterization of the optimal intervention as $a \rightarrow \infty$.

Letting x be the Lagrange multiplier on the constraint that $\mathbb{E}[R] = 0$, the Lagrangian of the planner's problem ($\mathbf{I}(a)$) can be written as¹⁰

$$\begin{aligned} \mathcal{L} &= \mathcal{W}(\hat{C}_{\text{FO}}) + z \mathbb{E}[R] \\ &= - \sum_{\ell} \left[\lambda_{\ell} \underline{q}_{\ell} \bar{\tau}_{\ell} + \frac{a}{2} \lambda_{\ell}^2 \underline{q}_{\ell}^2 \bar{\tau}_{\ell}^2 - z \underline{q}_{\ell} \bar{\tau}_{\ell} \right] \end{aligned}$$

⁹Since we are taking a to be large, the limit does not depend on ν^2 .

¹⁰Recall the planner's utility function from [Equation 7](#).

Differentiating with respect to each $\bar{\tau}_\ell$ gives

$$q_\ell \bar{\tau}_\ell = \frac{z - \lambda_\ell}{a \lambda_\ell^2}. \quad (17)$$

B.1. Solution for the Lagrange multiplier z . Plugging these into the formula (16) above yields the following sequence of equivalent expressions

$$\begin{aligned} \sum_\ell \bar{\tau}_\ell q_\ell &= 0 \\ \sum_\ell \left(\frac{z - \lambda_\ell}{\lambda_\ell^2} \right) &= 0 \\ z \sum_\ell \lambda_\ell^{-2} - \sum_\ell \lambda_\ell^{-1} &= 0. \end{aligned}$$

Recalling the definition $\lambda_\ell := (1 - \sigma_\ell)^{-1}$, we can deduce that

$$z = \frac{\frac{1}{N} \sum_\ell (1 - \sigma_\ell)}{\frac{1}{N} \sum_\ell (1 - \sigma_\ell)^2} \quad (18)$$

This gives an explicit form for the Lagrange multiplier.

To understand this formula better, define $\mu_\ell = \lambda_\ell^{-1} = 1 - \sigma_\ell$. Let μ be a random variable equal to μ_ℓ , with ℓ selected uniformly at random. Then we can rewrite

$$z = \frac{\mathbb{E}[\mu]}{\mathbb{E}[\mu^2]} = \frac{\mathbb{E}[\mu]}{\mathbb{E}[\mu]^2 + \text{Var}[\mu]}.$$

Simplifying this requires a preliminary lemma.

Lemma 2. $\sum_\ell \mu_\ell = 2N$.

Proof. Recall that σ_ℓ is the ℓ th eigenvalue of \mathbf{D} . By Property A, the trace of \mathbf{D} is $-N$. This implies that the sum of its eigenvalues, $\sum_\ell \sigma_\ell$, is equal to $-N$. Thus the sum of $\mu_\ell = 1 - \sigma_\ell$ is equal to $2N$. \square

We may then characterize the Lagrange multiplier as

$$z = \left[2 + \frac{1}{2} \text{var}[\boldsymbol{\sigma}] \right]^{-1}. \quad (19)$$

B.2. A formula for the change in social surplus. We can also plug in our solution for the optimal intervention to obtain an asymptotically correct formula for $\hat{\mathcal{C}}$.

First, recall from Lemma 1 that

$$\hat{\mathcal{C}}_{\text{FO}} = - \sum_\ell \lambda_\ell q_\ell \bar{\tau}_\ell.$$

Plugging in (17), we get

$$E(\hat{\mathcal{C}}_{\text{FO}}) = - \sum_\ell \frac{z - \lambda_\ell}{a \lambda_\ell}.$$

Thus

$$aE(\hat{\mathcal{C}}_{\text{FO}}) = - \sum_\ell \left(\frac{z}{\lambda_\ell} - 1 \right)$$

$$\begin{aligned}
&= N - \sum_{\ell} \frac{\mu_{\ell}}{2 + \frac{1}{2} \text{Var}[\sigma]} \\
&= N - N \frac{2}{2 + \frac{1}{2} \text{Var}[\sigma]} \quad \text{using Lemma 2.}
\end{aligned}$$

From this we conclude that

$$\frac{a}{N} E(\widehat{\mathcal{C}}_{\text{FO}}) = 1 - \frac{1}{1 + \frac{1}{4} \text{Var}[\sigma]}.$$

B.3. Dispensing with the approximate consumer surplus. It remains to show that neglecting higher-order terms in the optimization problem—solving the problem with $\widehat{\mathcal{C}}_{\text{FO}}$ replacing $\widehat{\mathcal{C}}$ —does not change the asymptotic characterization of the optimal intervention.

We will prove the result by studying an equivalent optimization problem using Berge's Theorem of the Maximum. Let $\check{\tau} = a\bar{\tau}$. We will now define a rescaled version of the problem, $\check{\mathbb{I}}(a)$.

$$\begin{aligned}
&\text{choose } \check{\tau} \text{ to } \max a\mathcal{W}(\check{\mathcal{C}}(a^{-1}\check{\tau})) && (\check{\mathbb{I}}(a)) \\
&\text{subject to } \mathbb{E}[R] = 0
\end{aligned}$$

Here $\check{\mathcal{C}}$ is the change consumer surplus when the tax profile is represented in the diagonal basis. This is clearly equivalent to the original problem. Let $\check{\tau}^*(a)$ be the (possibly set-valued) solution for risk-aversion level $a > 0$.

The problem $\check{\mathbb{I}}(a)$ is not yet defined at $a = \infty$, but we now define it there. Let the objective at $a = \infty$ be the limit of $a\mathcal{W}(a^{-1}\check{\tau})$ as $a \rightarrow \infty$; it is easy to check that this limit is actually

$$F(\check{\tau}) = - \sum_{\ell} \left[\lambda_{\ell} \underline{q}_{\ell} \check{\tau}_{\ell} + \frac{1}{2} \lambda_{\ell}^2 \underline{q}_{\ell}^2 \check{\tau}_{\ell}^2 \right].$$

Let the constraint be

$$\sum_{\ell} \check{\tau}_{\ell} \underline{q}_{\ell} = 0, \quad (20)$$

which is (16) with the second-order terms dropped.

When we restrict the optimization problem $\check{\mathbb{I}}(a)$ to any compact set \mathcal{K} containing the origin, the of Berge's Theorem of the Maximum are satisfied: The constraint correspondence is continuous at $a = \infty$, while the objective function is jointly continuous in a and $\check{\tau}$. The Theorem of the Maximum therefore implies that the maximized value is continuous at $a = \infty$. Because the convergence of the objective is actually uniform on \mathcal{K} , this is possible if and only if $\check{\tau}$ approaches a solution of the problem

$$\begin{aligned}
&\max_{\check{\tau}} F(\check{\tau}) \\
&\text{s.t. } \sum_{\ell} \check{\tau}_{\ell} \underline{q}_{\ell} = 0.
\end{aligned}$$

Taking \mathcal{K} to contain the interior solutions of this problem, this shows that the conclusions claimed hold without the first-order approximation.

B.4. Another way to write the variance of σ . Here is another way to write the variance of the eigenvalues, for what it's worth: Note that the average of the σ_ℓ is -1 , so that

$$\begin{aligned}
\text{var}(\boldsymbol{\sigma}) &= \frac{1}{N} \sum_{\ell} (\sigma_{\ell} + 1)^2 \\
&= \frac{1}{N} \text{trace}(\mathbf{I} + \boldsymbol{\Sigma})^2 \\
&= \frac{1}{N} \text{trace}[(\mathbf{I} + \mathbf{D})^2] && \text{trace invariant under similarity} \\
&= \frac{1}{N} \text{trace}(\mathbf{I} + 2\mathbf{D} + \mathbf{D}^2) \\
&= \frac{1}{N} (N - 2N + \text{trace}(\mathbf{D}^2)) && \text{trace of } \mathbf{D} \text{ is } N \\
&= -1 + \frac{1}{N} \sum_{i,j} D_{ij}^2 = \frac{1}{N} \sum_{i \neq j} D_{ij}^2
\end{aligned}$$

APPENDIX C. ANALYSIS OF THE LINEAR MODEL

The representative consumer chooses a bundle \mathbf{q} to maximize

$$\mathcal{C} = \boldsymbol{\beta}^T \mathbf{q} - \frac{1}{2} \mathbf{q}^T \mathbf{B} \mathbf{q} - \langle \mathbf{p}, \mathbf{q} \rangle.$$

Recall that we are assuming that the representative consumer has enough income to buy the optimal consumption bundle. The first-order conditions for optimal consumption imply that, at an interior solution,

$$\mathbf{q} = \mathbf{B}^{-1} (\boldsymbol{\beta} - \mathbf{p}) \tag{21}$$

The consumer's utility under optimal consumption (i.e., the consumer surplus) is therefore

$$\mathcal{C} = \frac{1}{2} \mathbf{q}^T \mathbf{B} \mathbf{q}.$$

The matrix of spillovers \mathbf{D} defined in the main text is $\mathbf{D} = -\mathbf{B}^{-1}$ and the normalization in Property A sets $D_{ii} = -(\mathbf{B}^{-1})_{ii} = -1$. Furthermore, given linear demand, the equilibrium in prices (under the assumption of uniqueness and interiority) is

$$\mathbf{p} = \mathbf{q} + \mathbf{c}$$

Using equilibrium demand 21 we have

$$\mathbf{p} = -\mathbf{D} (\boldsymbol{\beta} - \mathbf{p}) + \mathbf{c}$$

or, after defining $\mathbf{a} := -\mathbf{D}\boldsymbol{\beta}$,

$$[\mathbf{I} - \mathbf{D}]\mathbf{p} = \mathbf{a} + \mathbf{c}$$

The intervention problem specified in 12 is analogous to:

$$\begin{aligned}
&\underset{\boldsymbol{\tau}}{\text{Maximize}} && \frac{1}{2} \mathbf{q}^T \mathbf{B} \mathbf{q} \\
&\text{subject to} && \mathbf{q} = \mathbf{a} + \mathbf{D} \mathbf{p}, \\
& && [\mathbf{I} - \mathbf{D}]\mathbf{p} = \mathbf{a} + \mathbf{c} + \boldsymbol{\tau} \\
& && \langle \boldsymbol{\tau}, \mathbf{q} \rangle \geq 0
\end{aligned} \tag{22}$$

Proof of Proposition 2. Using the decomposition $D = U\Sigma U^\top$, we can rewrite the problem as:

$$\begin{aligned} & \underset{\underline{\tau}}{\text{Maximize}} && -\frac{1}{2}\underline{q}^\top \Sigma^{-1} \underline{q} \\ & \text{subject to} && \underline{q} = \underline{a} + \Sigma \underline{p}, \\ & && [\mathbf{I} - \Sigma] \underline{p} = \underline{a} + \underline{c} + \underline{\tau} \\ & && \langle \underline{\tau}, \underline{q} \rangle \geq 0 \end{aligned} \quad (23)$$

Combining the demand equation with the equilibrium price equation (i.e., the first two constrains), and defining the quantity consumed in equilibrium pre-intervention by $\underline{q}^0 = \underline{a} + \Sigma [\mathbf{I} - \Sigma]^{-1} (\underline{a} + \underline{c})$, we obtain that the quantities consumed after intervention $\underline{\tau}$ are

$$\underline{q} = \underline{q}^0 + \Sigma [\mathbf{I} - \Sigma]^{-1} \underline{\tau}.$$

Therefore the taxation problem is

$$\begin{aligned} & \underset{\underline{\tau}}{\text{Maximize}} && \frac{1}{2}\underline{q}^\top \Sigma^{-1} \underline{q} \\ & \text{subject to} && \underline{q} = \underline{q}^0 - \Sigma [\mathbf{I} + \Sigma]^{-1} \underline{\tau} \\ & && \langle \underline{\tau}, \underline{q} \rangle \geq 0 \end{aligned} \quad (24)$$

which can be rewritten as follows:

$$\begin{aligned} & \underset{\underline{\tau}}{\text{Maximize}} && -\frac{1}{2} \sum_{\ell} \frac{1}{\sigma_{\ell}} \left(\underline{q}_{\ell}^0 + \frac{\sigma_{\ell}}{1 - \sigma_{\ell}} \underline{\tau}_{\ell} \right)^2 \\ & \text{subject to} && -\sum_{\ell} \underline{\tau}_{\ell} \left(\underline{q}_{\ell}^0 + \frac{\sigma_{\ell}}{1 - \sigma_{\ell}} \underline{\tau}_{\ell} \right) \leq 0 \end{aligned} \quad (25)$$

Letting $z \geq 0$ be the multiplier on the constraint (the shadow price of public funds), we have the Lagrangian¹¹

$$\mathcal{L} = -\frac{1}{2} \sum_{\ell} \frac{1}{\sigma_{\ell}} \left(\underline{q}_{\ell}^0 + \frac{\sigma_{\ell}}{1 - \sigma_{\ell}} \underline{\tau}_{\ell} \right)^2 + z \sum_{\ell} \underline{\tau}_{\ell} \left(\underline{q}_{\ell}^0 + \frac{\sigma_{\ell}}{1 - \sigma_{\ell}} \underline{\tau}_{\ell} \right).$$

We can choose the signs of the eigenbundles so that $\underline{q}_{\ell}^0 \geq 0$ for every $\ell \in \mathcal{N}$. Let us also order the eigenvalues so that $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_n$. Recall that $\lambda_{\ell} = 1/(1 - \sigma_{\ell})$.

Lemma 3. The necessary conditions for an optimum, $\frac{\partial \mathcal{L}}{\partial \underline{\tau}_{\ell}} = 0$ for all $\ell \in \mathcal{N}$, imply the optimal taxes are characterized by

$$\underline{\tau}_{\ell} = \frac{\underline{q}_{\ell}^0}{1 - \lambda_{\ell}} \frac{z - \lambda_{\ell}}{2z - \lambda_{\ell}}.$$

Proof. Writing out $\frac{\partial \mathcal{L}}{\partial \underline{\tau}_{\ell}} = 0$ we have

$$-\left(\underline{q}_{\ell}^0 + \frac{\sigma_{\ell}}{1 - \sigma_{\ell}} \underline{\tau}_{\ell} \right) \frac{1}{1 - \sigma_{\ell}} + z \left(\underline{q}_{\ell}^0 + 2 \frac{\sigma_{\ell}}{1 - \sigma_{\ell}} \underline{\tau}_{\ell} \right) = 0$$

or

$$\underline{\tau}_{\ell} \frac{-\sigma_{\ell}}{1 - \sigma_{\ell}} \left(\frac{1}{1 - \sigma_{\ell}} - 2z \right) = \underline{q}_{\ell}^0 \left(\frac{1}{1 - \sigma_{\ell}} - z \right)$$

or

$$\underline{\tau}_{\ell} \frac{-\sigma_{\ell}}{1 - \sigma_{\ell}} = \underline{q}_{\ell}^0 \frac{z(1 - \sigma_{\ell}) - 1}{2z(1 - \sigma_{\ell}) - 1}.$$

Solving for $\underline{\tau}_{\ell}$ and rewriting in terms of λ_{ℓ} yields the claim. \square

¹¹We first write $-z$ as the coefficient on the derivative of the constraint, and then cancel the minus sign on the constraint in (25) with the minus sign on z .

Lemma 4. The quantity of each eigenbundle at the planner's optimum is

$$\underline{q}_\ell = \underline{q}_\ell^0 \frac{z}{2z - \lambda_\ell}.$$

Proof. Using the formula

$$\underline{\mathbf{q}} = \underline{\mathbf{q}}^0 + \Sigma [\mathbf{I} - \Sigma]^{-1} \underline{\boldsymbol{\tau}}$$

in the principal component basis, along with the expression for optimal taxes from Lemma 3,

$$\begin{aligned} \underline{q}_\ell &= \underline{q}_\ell^0 + \frac{\sigma_\ell}{1 - \sigma_\ell} \underline{q}_\ell^0 \frac{1}{1 + \lambda_\ell} \frac{z - \lambda_\ell}{2z - \lambda_\ell} \\ &= \underline{q}_\ell^0 \left(1 + \frac{\sigma_\ell}{1 - \sigma_\ell} \frac{1}{1 - \lambda_\ell} \frac{z - \lambda_\ell}{2z - \lambda_\ell} \right) \\ &= \underline{q}_\ell^0 \left(1 - \frac{z - \lambda_\ell}{2z - \lambda_\ell} \right) \\ &= \underline{q}_\ell^0 \cdot \frac{z}{2z - \lambda_\ell}. \end{aligned}$$

□

Lemma 5. The tax revenue raised from the ℓ^{th} eigenbundle is:

$$\underline{q}_\ell \underline{\tau}_\ell = (\underline{q}_\ell^0)^2 \frac{1}{1 - \lambda_\ell} \frac{z(z - \lambda_\ell)}{(2z - \lambda_\ell)^2}.$$

Proof. This is immediate using the expression for optimal taxes from Lemma 3 and the expression for quantities after the intervention from Lemma 4. □

Using Lemma 5, we can rewrite the binding budget constraint $\sum_\ell \underline{q}_\ell \underline{\tau}_\ell = 0$ as

$$\sum_\ell (\underline{q}_\ell^0)^2 \frac{1}{1 - \lambda_\ell} \frac{z(z - \lambda_\ell)}{(2z - \lambda_\ell)^2} = 0. \quad (26)$$

As long as $\underline{q}_\ell^0 \neq 0$, it follows that either z must be equal to some λ_ℓ (with $\underline{q}_{\ell'}^0 = 0$ for all $\ell' \neq \ell$), or that $z - \lambda_\ell$ is positive for some ℓ and negative for some other ℓ . In either case, it follows that $z \in [\lambda_1, \lambda_n]$.

Each term of the sum induces an asymptote at $z = \lambda_\ell/2$. Let ℓ' be the smallest index so that $\underline{q}_{\ell'}^0 \neq 0$. (Recall $\lambda_n \leq \lambda_{n-1} \leq \dots \leq \lambda_1$.) The left-hand side of the above equation asymptotes to $-\infty$ as z decreases to $\lambda_{\ell'}/2$, and equals to $\sum_\ell (\underline{q}_\ell^0)^2 \frac{1}{4(1-\lambda_\ell)} > 0$ as z increases to $+\infty$. Thus there is a solution $z \geq \lambda_{\ell'}/2$ of the rewritten budget constraint, Equation 26. Moreover, any solution $z > 0$ of the equation along with the optimal tax formulas above gives (by the KKT theorem) a solution to our optimization problem. But the optimization problem is strictly convex, so that its solution must be unique, and thus there is a unique value of z that solves Equation 26.

If $\underline{q}_\ell \neq 0$ for all ℓ , which holds generically, we can conclude that

$$\max \left(\lambda_n, \frac{\lambda_1}{2} \right) \leq z \leq \lambda_1.$$

Indeed, all the inequalities discussed above must be strict.

Finally, since $2z - \lambda_\ell > 0$ for each ℓ (by the above argument) we have from Lemma 2 that $\underline{q}_{\ell'}$, the post-tax quantity, must have the same sign as $\underline{q}_{\ell'}^0$, the pre-tax quantity.

This concludes the proof of Proposition 2.