

ROBUST MARKET INTERVENTIONS

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

ABSTRACT. A large differentiated oligopoly yields inefficient market equilibria. An authority with imprecise information about the primitives of the market aims to design tax/subsidy interventions that increase efficiency *robustly*—i.e., with high probability. We identify a condition on demand that guarantees the existence of such interventions, and we show how to construct them using noisy estimates of demand complementarities and substitutabilities across products. The analysis works by deriving a novel description of the incidence of market interventions in terms of spectral statistics of a Slutsky matrix. Our notion of recoverable structure ensures that parts of the spectrum that are useful for the design of interventions are statistically recoverable from noisy demand estimates.

1. INTRODUCTION

The growing importance of firms with market power has led to a renewed interest in the study of imperfectly competitive markets (see, e.g., [De Loecker, Eeckhout, and Unger \(2020\)](#); [Azar and Vives \(2021\)](#); [Pellegrino \(2021\)](#); [Ederer and Pellegrino \(2021\)](#)). Markets with oligopolistic structures typically yield inefficient equilibrium outcomes, which can be improved through interventions. Intervention on one firm, however, creates spillovers on other firms due to their interrelated demands. An authority often has limited information on the structure of market demands. This leads us to develop a theory to determine when and how an authority with imprecise demand information can design interventions that increase welfare. Our approach combines classical welfare pass-through theory with statistical results on recoverable latent structures of large matrices to identify conditions on demand structure that ensure the robust achievement of welfare improvements even when many aspects of the demand system cannot be accurately estimated.

We study a multi-market oligopoly model—with general complementarities and substitutabilities across products—in which firms simultaneously set prices.

Date Printed. November 5, 2024.

Galeotti is at the London Business School, Golub is at Northwestern University, Goyal is at University of Cambridge and NYUAD, Talamàs is at IESE Business School, and Tamuz is at Caltech. We thank Yann Calvo-López and Yu-Chi Hsieh for excellent research assistance. 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). We thank Bruno Pellegrino for valuable exchanges and kindly sharing data with us, and Volker Nocke for a helpful discussion on a related project at the 2023 AEA meetings. We thank Pol Antràs, Modibo Camara, Yeon-Koo Che, Kevin Chen, Jeff Ely, Matthew Gentzkow, Tim Gowers, Matthew O. Jackson, Jose Luis Moraga, Stephen Morris, Martin Pietz, Ran Spiegler, Xavier Vives, and Kuang Xu for helpful comments.

The authority may be a marketplace operator (such as Alibaba or Amazon) or a government agency that regulates the market. This authority has access to a noisy signal about the market demand parameters—for instance, estimates produced from data on consumer behavior. We focus on intervention rules that prescribe a firm-specific tax/subsidy intervention for each possible signal realization. Our interest lies in the conditions under which intervention rules can increase total surplus *robustly*—with high probability, conditional on any true market state.

What makes the problem challenging is that, in markets with numerous and changing goods, such as those hosted on large online platforms, the demand system is very high-dimensional, and realistic signals will leave substantial uncertainty about many aspects of the market’s structure; see [Section 7.1](#) for a detailed discussion. Our main result is that if demand satisfies a property that we call “recoverable structure,” then there are feasible intervention rules that robustly increase the total surplus. Moreover, within a natural class of interventions—those which do not reduce consumer surplus¹—our feasible interventions achieve the largest gain in surplus that is possible for a given level of subsidy expenditure. Hence, these interventions are as good as those that could be designed by an authority with perfect information.

Furthermore, our results provide tight conditions for robust intervention in the following sense. First, the property of recoverable structure cannot be dispensed with; there exist reasonable demand systems without recoverable structure for which the authority cannot robustly increase total surplus. Second, we show that there are markets with recoverable structure in which *any* intervention rule that robustly increases total surplus is equivalent, in terms of how it allocates surplus, to the interventions identified by our main result.

We now explain what it means for demand to have a recoverable structure and we discuss how this property is used in the construction of robust interventions.

The demand structure of the oligopoly is reflected in the Slutsky matrix D . A given cell D_{ij} in this matrix gives us the derivative of product i ’s demand with respect to the price of product j . Thus the matrix specifies the complementarity and substitutability relationships across products. The authority’s signal consists of noisy estimates of the entries of this matrix, with standard assumptions on the errors. The recoverable structure property requires that D has some eigenvalues that (in absolute value) grow sufficiently fast as the market gets large, and that the space spanned by their associated eigenvectors correlates with the market demand.

Recoverable structure has both economic and statistical implications that, combined, form the basis of our main result.

The recoverable structure property captures the presence of shared large-scale patterns in the demand of firms. For instance, suppose products can be embedded in a finite-dimensional space, and the location of a product in this space reflects its relation with other products, so that the cell D_{ij} depends on product attributes. If there is a latent set of characteristics relevant for predicting D_{ij} across many products and relevant for predicting market demands

¹This constraint is natural for a platform that could lose customers to other marketplaces or a government that can lose political support.

(i.e., regressions based on these characteristics have nonvanishing explanatory power) then the matrix D would have a recoverable structure.

The statistical aspects of recoverable structure are central to the design of interventions when the Slutsky matrix is observed imperfectly. In particular, the Davis–Kahan Theorem implies that in large markets with a recoverable structure, the noisy observation of D can be used to precisely estimate the space of eigenvectors associated with the largest eigenvalues of D (Davis and Kahan, 1970).

To use these estimates, we turn to the economic implications of recoverable structure. A key concept in this analysis is pass-through—how changes in firms’ costs affect equilibrium prices, quantities and welfare. We provide a new spectral description of the pass-through of an intervention. That is, we diagonalize the Slutsky matrix to obtain a specific orthonormal basis in which we can express any intervention as a linear combination of orthogonal effects. These effects correspond to the projection of the intervention onto each eigenvector of the Slutsky matrix D . In large markets with a recoverable structure, an intervention that projects only onto eigenvectors of D with large eigenvalues has some very useful properties. First, and most importantly, the orthogonal decomposition means that the welfare effects of such interventions depend only on large eigenvalues of the Slutsky matrix and the subspace they jointly span—the parameters the authority can robustly identify. Second, there are interventions of this form that increase total surplus by more than they cost the authority, providing a net social gain with very low uncertainty. Third, such interventions have minimal effects on prices, which implies that consumer surplus is essentially unchanged, and so the surplus increase achieved by the intervention is incident entirely on producers. Fourth, such interventions actually achieve the maximum possible total surplus per dollar spent subject to the constraint that consumer surplus does not decrease.

Combining the economic and statistical implications of recoverable structure allows us to establish our main result. In markets possessing such structure, the authority can recover precise information about large-eigenvalue components of the spectral decomposition. Using the economic implications of this decomposition that we have just discussed, we then show that interventions based on this information have highly predictable welfare implications, and indeed can be feasibly chosen to achieve the largest possible surplus gain per unit of expenditure subject to not significantly harming consumers.

1.1. Related literature. Our paper contributes to the 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), Azar and Vives (2021), Nocke and Schutz (2018) and Nocke and Whinston (2022). A recent strand of research in macroeconomics and industrial organization uses network models of differentiated oligopoly—similar to the one we consider here—to provide empirical estimates of efficiency losses due to market power (e.g., Pellegrino (2021) and Ederer and Pellegrino (2021)).²

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

Given these estimates of inefficiencies, a natural theoretical question is: What feasible interventions can improve welfare? Our main contribution is to analyze interventions from the perspective of an authority uncertain about the demand structure. Our analysis combines new spectral pass-through formulas with results building on the statistical theory of large matrices, and we identify conditions on the demand structure that ensure the robust achievement of welfare improvements even when many aspects of the demand structure cannot be accurately estimated.³ This approach has significant implications for the sorts of empirical models that are needed to design interventions in large markets with many goods. We elaborate on these issues in [Section 7.2](#).

Methods in high-dimensional statistics are currently attracting considerable interest in econometric settings where covariates are high-dimensional (see, e.g., [Athey, Bayati, Doudchenko, Imbens, and Khosravi \(2021\)](#) and [Chernozhukov, Hansen, Liao, and Zhu \(2023\)](#)), and there is work applying related statistical models to informational or behavioral spillovers in social networks ([Golub and Jackson, 2012](#); [Dasaratha, 2020](#); [Cai, 2022](#); [Parise and Ozdaglar, 2023](#); [Chandrasekhar, Goldsmith-Pinkham, McCormick, Thau, and Wei, 2024](#)). However, we know little about the question of when noisy data can be effectively used in order to implement desirable interventions in the presence of strategic spillovers, particularly when it comes to market settings. We show that, in a large oligopoly market, the concepts developed in the literature on large network recovery can be useful for the design of socially desirable interventions.

Finally, our paper contributes to the theory of network interventions. Early contributions in this field include [Borgatti \(2006\)](#), [Ballester, Calvó-Armengol, and Zenou \(2006\)](#), and [Goyal \(1996\)](#).⁴ In recent years, spectral methods have been applied to study optimal intervention problems under the assumption that the authority has perfect information about the network of spillovers ([Galeotti, Golub, and Goyal, 2020](#); [Gaitonde, Kleinberg, and Tardos, 2021](#); [Liu and Tsyvinski, 2024](#)).⁵ By contrast, in the present paper, we study a setting in which the authority observes the structure of strategic spillovers with significant noise.⁶ The methods we develop for robust interventions can be applied to other network games more generally and we briefly discuss this in [Section 7.3](#).

³Our focus on pass-through builds on work emphasizing the value of pass-through as a conceptual tool, e.g., [Marshall \(1890\)](#), [Pigou \(1920\)](#), [Dixit \(1979\)](#) and, more recently, [Weyl and Fabinger \(2013\)](#), [Miklos-Thal and Shaffer \(2021\)](#) and [Norris \(2024\)](#).

⁴The literature on this subject is very large. Intervention design has been studied in models of information diffusion, advertising, finance, security, and pricing, among other topics—see e.g., [Banerjee, Chandrasekhar, Duflo, and Jackson \(2013\)](#), [Bloch and Querou \(2013\)](#), [Candogan, Bimpikis, and Ozdaglar \(2012\)](#), [Elliott et al. \(2014\)](#), [Belhaj and Deroian \(2017\)](#), [Demange \(2017\)](#), [Dziubinski and Goyal \(2017\)](#), [Galeotti and Goyal \(2009\)](#), [Galeotti and Rogers \(2013\)](#), and [Leduc, Jackson, and Johari \(2017\)](#).

⁵Some recent work uses spectral analysis to derive conditions for core-selecting reallocation auctions ([Rostek and Yoder, 2023](#)), and robust implementation ([Ollár and Penta, 2023](#)).

⁶We share with several prior papers the idea that decision-makers act under partial information about the network. For instance, [Galeotti, Goyal, Jackson, Vega-Redondo, and Yarif \(2010\)](#) study large network games where players have incomplete information about the network structure, described by a random graph; [Akbarpour, Malladi, and Saberi \(2020\)](#) considers seeding in a large random graph; the diffusion process there lacks any form of complementarity. Our questions and methods of analysis are very different from these papers.

2. MODEL

In this section we describe the differentiated oligopoly game, the interventions available to the authority, and our notion of robust interventions.

2.1. The differentiated oligopoly game.

2.1.1. *Demand side.* There is a set $\{1, \dots, n\}$ of distinct products. The demand for these products arises from the consumption choices of many optimizing households. Each household $h \in \{1, \dots, H\}$ takes prices as given and has a *choice utility* that is quasilinear in a numeraire m ,

$$U^h(\tilde{\mathbf{q}}^h, m) = V^h(\tilde{\mathbf{q}}^h) + m,$$

where V^h is a twice-differentiable and strictly concave function of the consumption profile $\tilde{\mathbf{q}}^h \in \mathbb{R}^n$ and m is a numeraire (“money”), in which all prices are denominated. Given a price profile $\tilde{\mathbf{p}}$, the household’s problem is to choose a bundle $\tilde{\mathbf{q}}^h$ to maximize $U^h(\tilde{\mathbf{q}}^h, m) - \tilde{\mathbf{p}} \cdot \tilde{\mathbf{q}}^h$. Letting $\mathbf{q}^h(\tilde{\mathbf{p}})$ be the (unique, by concavity of V^h) solution to household h ’s problem, total market demand is⁷

$$\mathbf{q}(\tilde{\mathbf{p}}) = \sum_{h=1}^H \mathbf{q}^h(\tilde{\mathbf{p}}).$$

Note that we use tilde notation for an arbitrary price or quantity, and then drop the tilde for these variables to indicate some optimal or equilibrium solution.

2.1.2. *Supply side.* There is a firm associated with each product: Firm i produces good i . Firms play a simultaneous pricing game; each firm chooses $\tilde{p}_i \geq 0$. For any realized profile of prices $\tilde{\mathbf{p}}$, firm i ’s profit is

$$q_i(\tilde{\mathbf{p}})(\tilde{p}_i - c_i), \tag{1}$$

where c_i is the (constant) marginal cost of production.

We fix a vector \mathbf{c}^0 of marginal costs and a pure-strategy Nash equilibrium \mathbf{p}^0 , and we refer to these as the *status quo* marginal costs and equilibrium, respectively.⁸ To facilitate unambiguous local comparative statics, we make the following assumption.

Assumption 1 (Local equilibrium uniqueness). There exist $\nu > 0$ and $\rho > 0$ such that, for all \mathbf{c} with $\|\mathbf{c} - \mathbf{c}^0\| < \nu$, there is a unique pure-strategy Nash equilibrium $\mathbf{p}(\mathbf{c})$ in a ρ -neighborhood of \mathbf{p}^0 .

From now on, we confine attention to cost perturbations within the set discussed in [Assumption 1](#), and when we refer to an *equilibrium* at any cost profile, we mean the locally unique one entailed by the above assumption.

⁷Because the households’ utilities are quasilinear in money, one can derive the same aggregate demand from a representative consumer, and this description is sufficient for studying producer and aggregate consumer surplus. However, some of our results will provide more refined results about the effects of consumers, showing that no household is significantly hurt.

⁸Existence of a pure strategy equilibrium is guaranteed if profits are quasi-concave in prices (see Theorem 1.2 in [Fudenberg and Tirole \(1991\)](#)). A sufficient condition for this is that the functions $1/q_i(\tilde{\mathbf{p}})$ are convex in \tilde{p}_i ([Vives \(1999\)](#) page 149). This holds, for instance, under the much stronger condition that demand is linear in prices; in this case the equilibrium is also unique. The general sufficient condition for local uniqueness is nonsingularity of the Jacobian of best responses at equilibrium, which will hold generically in our setting ([McLennan, 2018](#)).

2.2. Interventions and outcomes. An *authority*—an institution overseeing a marketplace—can intervene on the market. For concreteness, we focus on a canonical set of interventions: per-unit subsidies and taxes. For a consumption profile $\tilde{\mathbf{q}}$, a *per-unit subsidy intervention*

$$\boldsymbol{\sigma} = (\sigma_1, \dots, \sigma_n)$$

consists of a transfer $\sigma_i \tilde{q}_i$ from the authority to firm i ; a positive σ_i corresponds to a subsidy to firm i , while a negative σ_i corresponds to a tax. The authority's net spending associated with such an intervention is denoted by $S = \boldsymbol{\sigma} \cdot \tilde{\mathbf{q}}$. Recalling [Assumption 1](#), the set of feasible interventions is⁹

$$\Sigma = \{\boldsymbol{\sigma} \in \mathbb{R}^n : \|\boldsymbol{\sigma}\| < \nu\}.$$

The authority cares about the surplus that different market participants obtain in equilibrium. We focus on three canonical metrics: consumer surplus C , producer surplus P , and total surplus W (accounting for the intervention expenditure S). Given an intervention $\boldsymbol{\sigma}$, and quantity profiles $\{\mathbf{q}^h\}_{h=1,\dots,n}$ and $\mathbf{q} := \sum_h \mathbf{q}^h$, these are:

$$C = \sum_h C^h \quad \text{where} \quad C^h = V^h(\mathbf{q}^h) - \mathbf{q}^h \cdot \mathbf{p},$$

$$P = (\mathbf{p} - \mathbf{c}) \cdot \mathbf{q} \quad \text{and} \quad W = C + P - S. \quad (2)$$

We evaluate the effect of an intervention $\boldsymbol{\sigma}$ on an outcome variable Y by its first derivative. Formally, the first-order effect on any given outcome variable Y of changing subsidies in the direction $\boldsymbol{\sigma}$ is defined by¹⁰

$$\dot{Y}_{\boldsymbol{\sigma}} = \sum_i \frac{dY}{d\sigma_i} \sigma_i. \quad (3)$$

2.3. A locally linear oligopoly: Market states and signals. The authority receives a noisy signal about the structure of the market and, using this information, wants to implement interventions that robustly increase standard measures of surplus.

Before giving a general definition of this problem and of robust interventions, we introduce a special case—one with demand functions that are linear in a neighborhood of the status quo equilibrium, as captured by the following assumption.

Assumption 2 (Local linearity of demand). There exists $\rho > 0$ such that, in a neighborhood of radius ρ around the initial equilibrium \mathbf{p}^0 , the demand function $q_i(\tilde{\mathbf{p}})$ is linear in \tilde{p}_i for every firm i .

This setting is highly tractable and, as we will see, allows us to formulate an interesting robust intervention problem. We discuss how to extend the analysis for the case of non-linear demand in [Section 7.4](#). We now describe what the authority would ideally like to know, the signals it actually has, and then define the notion of a robust intervention.

⁹We restrict attention to deterministic interventions but this is immaterial to our results.

¹⁰We often omit the subscript $\boldsymbol{\sigma}$ when there is no ambiguity about the relevant intervention.

The firms' first-order conditions imply that equilibrium prices \mathbf{p} around the status quo \mathbf{p}^0 satisfy

$$q_i(\mathbf{p}) = -\frac{\partial q_i}{\partial p_i}(\mathbf{p})(p_i - c_i).$$

Local linearity implies that $\frac{\partial q_i}{\partial p_i}(\mathbf{p})$ is constant when prices change locally around \mathbf{p}^0 . Hence, we can replace the partial derivative in the above equation with $\frac{\partial q_i}{\partial p_i}(\mathbf{p}^0)$. By strict concavity of the consumer utility functions, this is a negative number; from now on, we maintain a normalization (by choosing suitable units in which to express the quantity produced by firm i) that $\frac{\partial q_i}{\partial p_i}(\mathbf{p}^0) = -1$ (see [Appendix A.1](#)). After this normalization, equilibrium behavior is summarized by the following system of equations:

$$\mathbf{q}(\mathbf{p}) = \mathbf{p} - \mathbf{c}. \quad (4)$$

Implicitly differentiating this (linear) system, we obtain that the effect of a small intervention $\boldsymbol{\sigma}$ on prices is determined by the following system of equations:

$$[\mathbf{I} - \mathbf{D}]\dot{\mathbf{p}} = -\boldsymbol{\sigma}, \quad (5)$$

where $\boldsymbol{\sigma} = \mathbf{c} - \mathbf{c}^0$ is the intervention (i.e., tax or subsidy offered by the authority), $\dot{\mathbf{p}}$ is the derivative of \mathbf{p} in the direction of $\boldsymbol{\sigma}$ (see [eq. \(3\)](#)), and $\mathbf{D} = \mathbf{D}(\mathbf{p}^0)$ is the Slutsky matrix (in the normalized units):

$$D_{ij} = \frac{\partial q_i}{\partial p_j}(\mathbf{p}^0).$$

For $i \neq j$, if $D_{ij} > 0$ (resp. $D_{ij} < 0$) then, around the equilibrium, products i and j are substitutes (resp. complements).¹¹ Note also that quantity changes following the intervention $\boldsymbol{\sigma}$ are pinned down by

$$\dot{\mathbf{q}} = \mathbf{D}\dot{\mathbf{p}}. \quad (6)$$

The local linearity of demand implies that comparative statics of prices and quantities are fully determined by the Slutsky matrix \mathbf{D} . This implies the following fact, which we will establish and discuss in our analysis below.

Fact 1. The matrix $\mathbf{D} = \mathbf{D}(\mathbf{p}^0)$ and the status quo quantities $\mathbf{q}^0 = \mathbf{q}(\mathbf{p}^0)$ together determine the first-order effects of an intervention $\boldsymbol{\sigma}$ on consumer surplus C , producer surplus P , and total surplus W , i.e., the derivatives \dot{C}_σ , \dot{P}_σ , and \dot{W}_σ .

We note that \mathbf{D} satisfies the following property ([Nocke and Schutz, 2017](#)).

Property NSD. The normalized Slutsky matrix \mathbf{D} is negative semidefinite, symmetric, and has diagonal entries $D_{ii} = -1$.

This property holds because the demand function can be taken to arise from a representative household (with a twice-differentiable utility function for goods equal to the sum of the consumers' utilities, V^h).

The following assumption is without loss of generality—it holds by suitably adjusting the units of the numeraire m , and will be useful later.

Assumption 3. The quantity vector's Euclidean norm $\|\mathbf{q}^0\|$ is at most 1.

¹¹In general, the matrix of demand derivatives need not be the same as the Slutsky matrix, but under quasilinearity the wealth effect is zero and the two coincide ([Nocke and Schutz, 2017](#)).

Fact 1 implies that if the authority knows the pair $(\mathbf{D}, \mathbf{q}^0)$ characterizing the market conditions at the status quo, then it knows the local surplus outcomes of interventions. We, therefore, call the tuple $(\mathbf{D}, \mathbf{q}^0)$, where \mathbf{D} is a negative semidefinite matrix with diagonal entries -1 and \mathbf{q}^0 is a vector of norm at most 1, the *market state*.

2.3.1. *The authority's signal.* The authority receives a signal about the market state $\boldsymbol{\theta} = (\mathbf{D}, \mathbf{q})$. This signal consists of random variables

$$\widehat{\mathbf{D}} = \mathbf{D} + \mathbf{E} \text{ and } \widehat{\mathbf{q}}^0 = \mathbf{q}^0 + \boldsymbol{\varepsilon}.$$

We make the following assumptions about the authority's signal. Let $\|\mathbf{E}\|$ denote the spectral norm of a matrix \mathbf{E} , which for a symmetric matrix is equal to its largest eigenvalue. Fix a sequence $b(n)$ and a positive constant \bar{V} .

Assumption 4. We assume that:

- (1) $\mathbb{E}[\|\mathbf{E}\|] \leq b(n)$;
- (2) all the ε_i are independent and $\mathbb{E}[\|\boldsymbol{\varepsilon}\|^2] \leq \bar{V}$.

The first part of the assumption bounds the matrix norm of the errors in estimating the normalized Slutsky matrix. In [Appendix C](#), we provide a simple procedure for sampling market data independently across product pairs (i, j) under which this assumption holds with $b(n) = \gamma n^{1/2}$ (for a positive constant $\gamma > 0$). We can think of this example as generating E_{ij} that are essentially independent, with variance of constant order (i.e., neither growing or decaying with n). This shows that the assumption can hold even when there is no entry of \mathbf{D} that can be recovered accurately.

If $b(n) = \gamma n^\beta$ with $\beta \in (1/2, 1)$, there are error structures consistent with [Assumption 4](#) where some terms in $\widehat{\mathbf{D}}$ have large covariances. For example, spatially correlated errors, with sufficiently "distant" parts of the matrix being at most slightly correlated, would satisfy this assumption.¹² An example of a situation that would cause part (1) of [Assumption 4](#) to be violated is the entries of \mathbf{E} all having correlation bounded away from zero (e.g., coming from a common shock to measured complementarities).

The second part of the assumption requires independence of errors across different products' quantities. The purpose of the assumption is to apply a law of large numbers for estimating outcomes such as the average quantity, as well as various linear combinations of quantities that are important for intervention outcomes. Independence is stronger than we need for our main result stated in [Theorem 1](#), and is made to facilitate exposition. When we use this assumption in the proof our main result, [Theorem 1](#), we rely on a substantially weaker condition (which is more technical to state) that [Assumption 4\(2\)](#) implies (see [Appendix A.2](#)). Concerning the assumption on the norm of $\boldsymbol{\varepsilon}$, recall that we assume that $\|\mathbf{q}^0\| \leq 1$; the assumption on $\boldsymbol{\varepsilon}$ scales the error to be of the same order of magnitude as the quantity vector.

We maintain Assumptions 1–4 unless specifically noted otherwise.

¹²The idea is the same as that of ergodicity-type conditions in time-series settings.

2.4. Robust interventions. In the linear oligopoly environment, [Fact 1](#) establishes that to determine the first-order surplus effects of any intervention σ , the only additional data needed is the tuple $(\mathbf{D}, \mathbf{q}^0)$, which we call the market state and denote by θ . We take a prior-free approach to modeling the authority's uncertainty over θ . That is, we assume the authority knows only that θ belongs to some set Θ and does not have a probability distribution over Θ . However, the authority knows the distribution of the signal $\mathbf{t} = (\hat{\mathbf{D}}, \hat{\mathbf{q}}^0)$ given $\theta \in \Theta$. We denote this distribution by μ_θ . We denote by \mathcal{T} the set of possible values that the signal \mathbf{t} can take.

Given the authority's signal \mathbf{t} , the authority designs an *intervention rule*¹³

$$\mathbf{R} : \mathcal{T} \rightarrow \Sigma,$$

prescribing an intervention $\sigma \in \Sigma$ for every possible signal \mathbf{t} .

A *market outcome* is a tuple (θ, σ) consisting of a market state and an intervention. We call this pair the outcome because it determines production, consumption and transfers. A *property* is a subset of the set $\Theta \times \Sigma$ of all possible outcomes. The authority wants to design an intervention rule that robustly achieves a desired property. We are interested in understanding which properties can be achieved with high probability in all market states that the authority considers possible:

Definition 1. An intervention rule \mathbf{R} achieves a property \mathbf{X} ϵ -robustly if the following holds: For every $\theta \in \Theta$, we have

$$\mu_\theta(\{\mathbf{t} : (\theta, \mathbf{R}(\mathbf{t})) \in \mathbf{X}\}) \geq 1 - \epsilon.$$

The only randomness in the definition is in the draw of the signal.

Notice that the definitions we have made do not depend on any details of the linear oligopoly environment. In any environment where (θ, σ) fully determines all outcomes of interest to the authority and there is a known distribution over signals given states, the definition can be applied.

3. RECOVERABLE STRUCTURE

We define a new concept, which we call recoverable structure, that plays a key role in the analysis of robust interventions that we will present in [Section 4](#). Recoverable structure imposes conditions on the pattern of complements and substitutes (which we will call interactions) among products, as summarized by the Slutsky matrix. It is a condition ensuring that some latent pattern of product interactions can be inferred robustly through the noisy observation and, at the same time, this latent structure is correlated with market quantities. We now introduce it formally and provide an economic interpretation.

A vector in \mathbb{R}^n describes a bundle of products. Given \mathbf{D} , we are interested in the subspace of bundles spanned by eigenvectors of \mathbf{D} with large eigenvalues. Formally, let $\mathcal{L}(\mathbf{D}, M) \subseteq \mathbb{R}^n$ be the subspace of the bundle space that is spanned by the eigenvectors of \mathbf{D} with eigenvalues at least M in absolute value.

Definition 2. The set of market states Θ has (M, δ) -recoverable structure if for every $(\mathbf{D}, \mathbf{q}^0) \in \Theta$ the projection of \mathbf{q}^0 onto $\mathcal{L}(\mathbf{D}, M)$ has norm at least δ .

¹³This should be measurable in a suitable sense, which is clear in our application.

To understand the definition, note that a Slutsky matrix, by virtue of being symmetric, can be orthogonally diagonalized: it can be written as a linear combination of orthogonal rank-one matrices:

$$\mathbf{D} = - \sum_{\mathbf{u}^\ell \in \mathcal{L}(\mathbf{D}, M)} |\lambda_\ell| \underbrace{\mathbf{u}^\ell (\mathbf{u}^\ell)^\top}_{\text{rank-1 matrix}} - \sum_{\mathbf{u}^\ell \notin \mathcal{L}(\mathbf{D}, M)} |\lambda_\ell| \underbrace{\mathbf{u}^\ell (\mathbf{u}^\ell)^\top}_{\text{rank-1 matrix}}$$

Here, $\lambda_1, \lambda_2, \dots, \lambda_n$ are the eigenvalues of \mathbf{D} (which are nonpositive numbers because \mathbf{D} is negative semidefinite), ordered from greatest to least in absolute value, and $(\mathbf{u}^1, \dots, \mathbf{u}^n)$ is a corresponding basis of orthonormal eigenvectors. All the summands are orthogonal to each other, and $|\lambda_\ell|$ is the norm of the contribution of the corresponding summand.

Having market states with (M, δ) -recoverable structure means that (i) \mathbf{D} has eigenvectors with eigenvalues larger than M in absolute value, so the first summation in the above formula is nonzero, and furthermore (ii) the vectors \mathbf{u}^ℓ in that sum can jointly account for a significant portion of the status quo quantities. As we will formally explain later, if we set M to be larger than the (suitably measured) size of the noise \mathbf{E} in the observation of \mathbf{D} , then condition (i) will ensure that the low-dimensional subspace $\mathcal{L}(\mathbf{D}, M)$ can be estimated precisely despite \mathbf{D} being observed with noise; it is intuitive that this is valuable for intervening in the market. However, if quantities were nearly orthogonal to this subspace, then it turns out that the recoverable structure would only make a small difference to surplus outcomes; the requirement (ii) entailed in the definition rules this out by imposing that the quantities in the market have nontrivial projection onto the recoverable subspace. In this sense the definition states that $(\mathbf{D}, \mathbf{q}^0)$ has a structure which can be recovered and is relevant to the market.

We provide an example to illustrate the notion of recoverable structure.

Example 1 (An illustration of recoverable structure). There are $n = 300$ products. The Slutsky matrix is the combination of two matrices:

$$\mathbf{D} = (1 - \gamma)\mathbf{D}_{\text{block}} + \gamma(-\mathbf{Z}\mathbf{Z}'). \quad (7)$$

The first matrix, $\mathbf{D}_{\text{block}}$, is a block matrix that divides the 300 products into 3 equally-sized blocks. The entries of $\mathbf{D}_{\text{block}}$ are based on a 3×3 matrix \mathbf{C} that governs the pattern of interactions within blocks (C_{ii}) and across them (C_{ij} with $i \neq j$). In particular,¹⁴

$$\mathbf{C} = \begin{pmatrix} -1 & 0.15 & 0.7 \\ 0.15 & -1 & 0.6 \\ 0.7 & 0.6 & -1 \end{pmatrix} \quad \text{and} \quad \mathbf{D}_{\text{block}} = \mathbf{C} \otimes \mathbf{J}_{n/3}.$$

Here \mathbf{J} is the matrix of ones. Products within each block are complements, and products across blocks are substitutes. For example, products in block 3 are highly substitutable with products in other blocks, whereas products in blocks 1 and 2 are only mildly substitutable with each other.

We can think of each block as a product category. In one category we have, for example, kitchen products, in another category we have digital entertainment products, and in the third category we have products related to sport activities. Within each of these categories, products are complements, but across the categories these products are substitutes. To take another example, we can think of

¹⁴The symbol \otimes denotes the Kronecker product.

three non-compatible operating systems that can be used to accomplish similar tasks; products within the same computer operating system are complements, and products across operating systems are substitutes.

The second term in eq. (7), a scaling of $-\mathbf{Z}\mathbf{Z}'$, is a negative definite matrix that adds heterogeneity to the very regular pattern of interactions in $\mathbf{D}_{\text{block}}$. In particular, we take \mathbf{Z} to be an $n \times 10$ matrix with rows drawn uniformly from the unit sphere.¹⁵

Finally, the observed Slutsky matrix is

$$\hat{\mathbf{D}} = \mathbf{D} + \mathbf{E}$$

where each entry of the error matrix \mathbf{E} —except for its diagonal entries, which are kept at zero—is drawn from $U[-1, 1]$; this noise structure results in the upper bound $b(n) = n^{1/2}$ on $\|\mathbf{E}\|$.

We are interested in illustrating some basic statistical implications of (M, δ) -recoverable structure that is strong enough relative to the noise $b(n)$ —that is, with M much larger than $b(n)$. In this example, the largest eigenvalue of $\mathbf{D}_{\text{block}}$ is much larger than $n^{1/2}$, and so when γ is low, then the \mathbf{D} of eq. (7) will also have some eigenvalues which are larger than $n^{1/2}$. In contrast, when γ is high, then the main contributor to \mathbf{D} is $-\mathbf{Z}\mathbf{Z}'$ and, in this case, the largest eigenvalues of \mathbf{D} turn out to be no larger than $n^{1/2}$. We now show that these two cases, low vs. high γ , have very different implications for inference.

First we consider a low value of $\gamma = 0.3$. The two largest eigenvalues of \mathbf{D} , in absolute value, are approximately 130 and 80, and these are considerably larger than $b(n) = \sqrt{300} \approx 17$, whereas the third largest eigenvalue is roughly 1.2; hence, $\mathcal{L}(\mathbf{D}, M)$ with $M > b(n)$ equal to, e.g., $n^{2/3}$, is the subspace spanned by eigenvectors \mathbf{u}^1 and \mathbf{u}^2 . As we will explain in Section 4.2, the Davis-Kahan Theorem tells us that we can use the noisy observation $\hat{\mathbf{D}}$ to accurately estimate the eigenvectors in $\mathcal{L}(\mathbf{D}, M)$. Here, for simplicity, we illustrate graphically how the first eigenvector of $\hat{\mathbf{D}}$ is a good description of the first eigenvector of \mathbf{D} .

Figure 1 depicts the Slutsky matrix when $\gamma = 0.3$. Panel A illustrates the true matrix \mathbf{D} and panel B illustrates the rank-1 matrix $\mathbf{u}^1(\mathbf{u}^1)^\top$ corresponding to the eigenvector with the highest eigenvalue. We can see that the rank-1 matrix of Panel B, despite being defined based on how \mathbf{D} acts on a very low-dimensional subspace, captures a lot of information about the pattern of demand interactions within and across blocks of \mathbf{D} . Panel C illustrates a realization of the observed matrix $\hat{\mathbf{D}}$. We now illustrate how the patterns in \mathbf{D} can be recovered from this observation. Let $\hat{\mathbf{u}}^1$ be the eigenvector of $\hat{\mathbf{D}}$ associated with its largest eigenvalue. Panel D illustrates the associated rank-1 matrix $\hat{\mathbf{u}}^1(\hat{\mathbf{u}}^1)^\top$. The close resemblance between the matrices in Panel D and Panel B is an informal illustration of how $\hat{\mathbf{u}}^1$ is a good approximation of the true \mathbf{u}^1 . That is, despite the fact that we observe \mathbf{D} with substantial noise, we can recover structure latent in \mathbf{D} , summarized by $\mathbf{u}^1(\mathbf{u}^1)^\top$, by simply calculating the largest-eigenvalue eigenvector of the noisy matrix $\hat{\mathbf{D}}$.

In Section 5 we extend this example to illustrate our main result Theorem 1 stated in Section 4, and show how we use the condition that recoverable structure imposes on the quantity vector.

¹⁵Note that entry (i, i) of $\mathbf{Z}\mathbf{Z}'$ is the norm of the i th row of \mathbf{Z} , which, by construction, is 1.

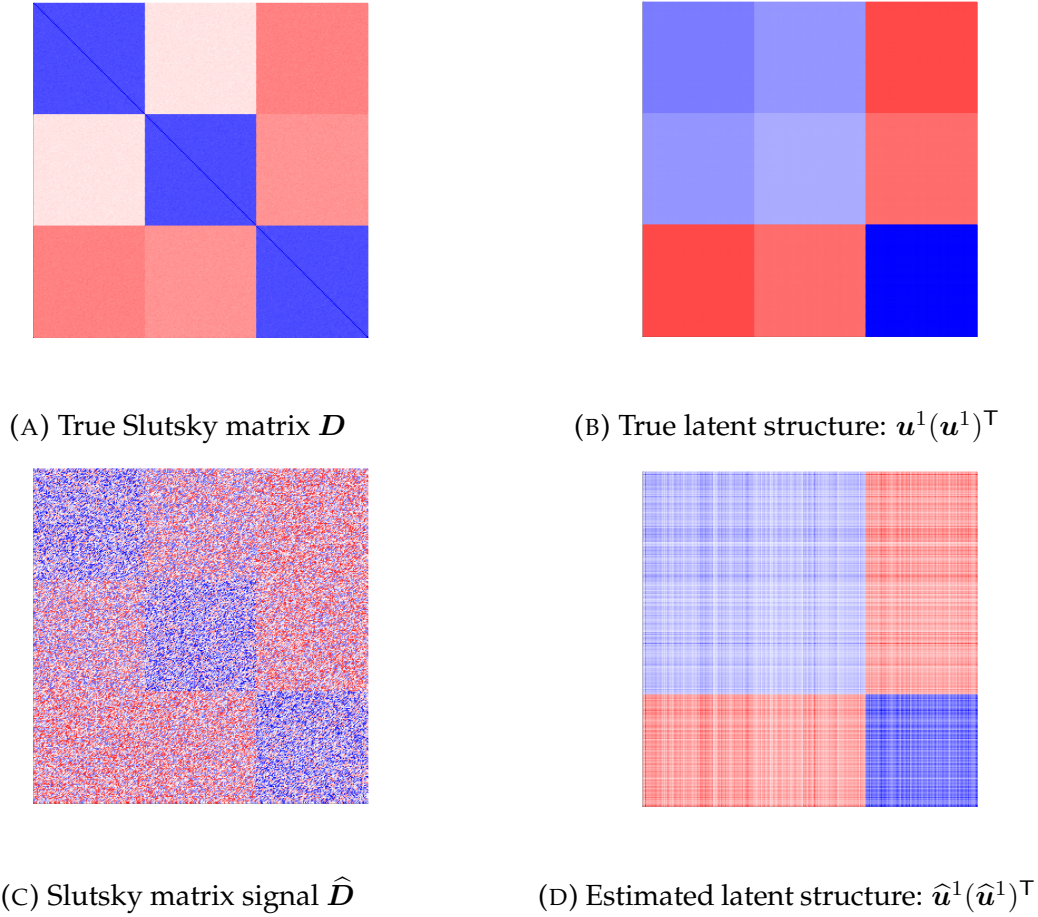


FIGURE 1. Illustration of true vs. estimated parameters when there is recoverable structure, $\gamma = 0.3$. Blue regions illustrate complementarities (a negative sign of the corresponding matrix entry) and the red regions signify substitutabilities (positive signs). The darker the color, the higher the corresponding entry in absolute value.

Figure 2 illustrates the Slutsky matrix when $\gamma = 0.9$. In this case the largest eigenvalue is comparable to \sqrt{n} and therefore is not much larger than the noise. Although $\gamma = 0.9$, and so D mainly consists of $-\mathbf{Z}\mathbf{Z}'$, the block matrix D_{block} is still visible in D (Panel A) and the true rank-1 matrix $\mathbf{u}^1(\mathbf{u}^1)^\top$ summarizes that pattern (Panel B). However, this structure cannot be recovered using noisy observation. This is illustrated by a realization of the estimated matrix \hat{D} (Panel C) and the associated rank-1 matrix $\hat{\mathbf{u}}^1(\hat{\mathbf{u}}^1)^\top$ (Panel D).

4. ANALYSIS OF ROBUST INTERVENTIONS

This section presents our main results, showing the existence of intervention rules that robustly achieve surplus improvements. To have a sense of what can be achieved at all, with or without noise, it is worth noting some basic properties of surplus under complete information:

Proposition 1. The following hold true:

- (1) Under any intervention, surplus outcomes satisfy $\frac{1}{2}\dot{P} + \dot{C} = \dot{S}$.

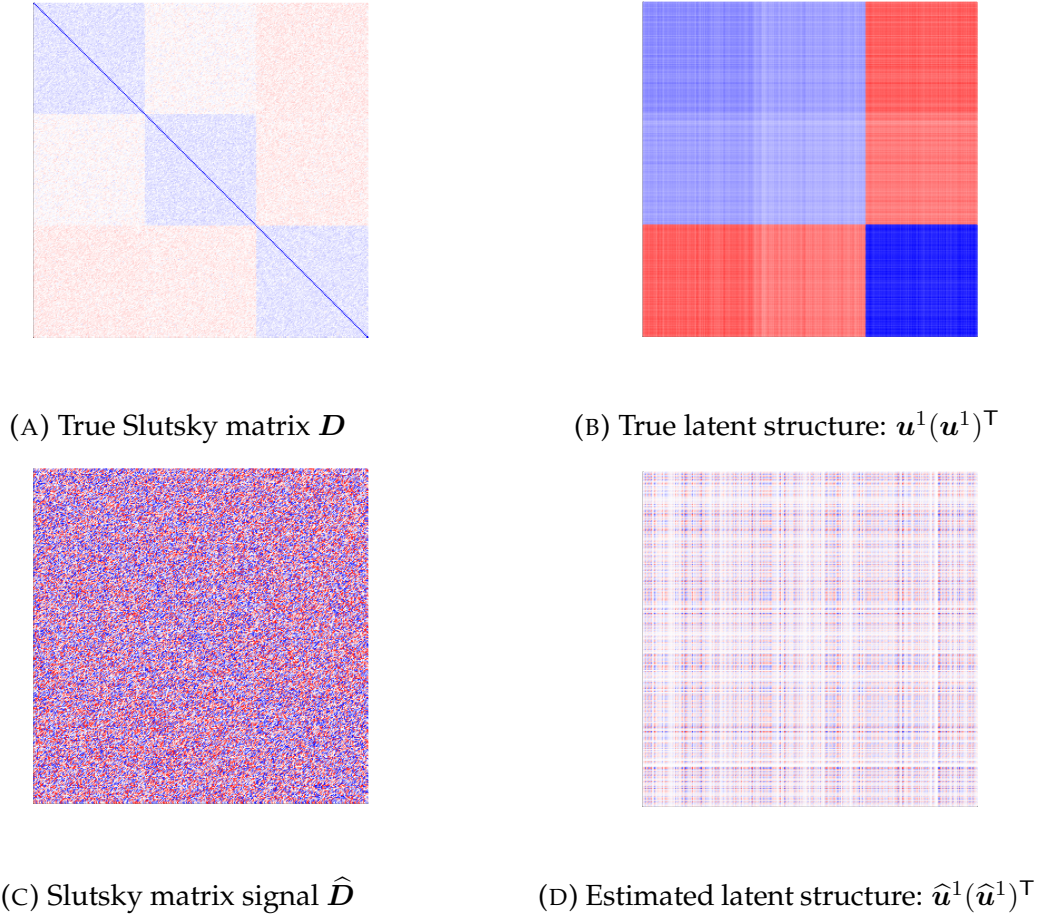


FIGURE 2. Illustration of true vs. estimated parameters when there is not recoverable structure, $\gamma = 0.9$. The meaning of the colors is as in Figure 2c

(2) For any intervention such that $\dot{C} \geq -\epsilon$, we have

$$\dot{P} + \dot{C} \leq (2 + \epsilon)\dot{S}.$$

Proof. Consumer surplus satisfies the standard Marshallian formula $\dot{C} = -\mathbf{q}^0 \cdot \dot{\mathbf{p}}$. Meanwhile, $\dot{P} = \mathbf{q}^0 \cdot (\dot{\mathbf{p}} + \boldsymbol{\sigma}) + (\mathbf{p} - \mathbf{c}) \cdot \dot{\mathbf{q}} = 2\mathbf{q}^0 \cdot \dot{\mathbf{q}}$, where we have used the equilibrium condition $\mathbf{p} - \mathbf{c} = \mathbf{q}$ from eq. (4). Finally, $\dot{S} = -\mathbf{q}^0 \cdot \dot{\mathbf{c}}$ by definition. These facts together imply part (1), and part (2) is an immediate corollary. \square

Part (1) implies that $\dot{P} + \dot{C} = 2\dot{S} - \dot{C}$, and so the market's surplus (the left-hand side) can potentially be increased without bound through the authority's expenditures, a reduction in consumer surplus, or a combination.¹⁶

Part (2) focuses on the set of interventions subject to a lower bound on the surplus of consumers, and derives an upper bound on surplus gain achievable in that case—even by an omniscient authority. The constraint is natural for authorities that want to limit consumer welfare loss—e.g., if the authority is a platform that can lose customers, or a government that can lose political support.

¹⁶Proposition 3 in Appendix B shows that the condition is not only necessary but also that, generically, all these outcomes can actually be achieved under complete information.

Our main result, [Theorem 1](#), examines what can be implemented robustly when the authority has partial information and the economy has a recoverable structure that is strong enough relative to the noise. It shows that there are interventions which robustly protect consumers from surplus loss and implement market surplus equal to the upper bound that an omniscient authority could achieve (the upper bound given by [Proposition 1\(2\)](#)).

Recall that we have fixed a sequence $b(n)$ under which [Assumption 4](#) holds—an upper bound on the noise in observations of the demand system.

Theorem 1. Let $M(n)$ be an increasing sequence with $M(n)/b(n) \rightarrow_n \infty$, and fix $\delta > 0$. Assume the set of market states has $(M(n), \delta)$ -recoverable structure. For every $\epsilon > 0$ and for every target expenditure $e > 0$, the following properties can be simultaneously achieved ϵ -robustly for sufficiently large n :

- (i) The sum of marginal consumer and producer surplus gains $\dot{P} + \dot{C}$ is at least twice the marginal expenditure, up to a small multiplicative error: $\dot{P} + \dot{C} \geq (2 - \epsilon)\dot{S}$.
- (ii) The marginal effect on consumer surplus is not significantly negative: $\dot{C} \geq -\epsilon$. Moreover, no individual consumer's surplus decreases significantly—i.e., $\dot{C}^h \geq -\epsilon$ for each household h .
- (iii) The marginal expenditure is arbitrarily close to the target expenditure e , i.e., $|\dot{S} - e| < \epsilon$.

The condition in [Theorem 1](#) stipulates that the market has $(M(n), \delta)$ -recoverable structure for some $M(n)$ that is asymptotically much larger than $b(n)$. This lower bound on $M(n)$ ensures that the recoverable structure in \mathbf{D} is substantially larger than the norm of the error matrix \mathbf{E} , which is $O(b(n))$ under [Assumption 4](#).

Part (i) states that, when this condition is satisfied, the authority can robustly achieve approximately two dollars of surplus gain per dollar spent.¹⁷

Point (ii) states that it is possible to do this in a way that leaves all households' welfare essentially unchanged; indeed, under the policy that we construct, the surplus gains are fully captured by the producers. Note that by [Proposition 1](#), the welfare gain described in point (i) is essentially the maximum total surplus change that an omniscient authority could implement with the same amount of spending without reducing consumer surplus.

Finally, point (iii) says the authority can precisely target the realized expenditure (and thus, total surplus impact) of the policy.

Our notion of ϵ -robustness says that the surplus properties above are achieved ex post with high probability. A natural question is whether they are also achieved in expectation (since, in principle, realizations with very negative surplus could occur with low probability). In our setting, it turns out that all the analysis would be unaffected if we added good ex ante expected performance to the definition of robustness.¹⁸

¹⁷Recalling the definition $\dot{W} = \dot{P} + \dot{C} - \dot{S}$, this implies that every dollar spent yields approximately one unit increase in \dot{W} , net total surplus.

¹⁸This is because the surplus pass-throughs are supported on $[0, 1]$ and all quantities in the proofs are bounded. Thus, convergence in probability is equivalent to convergence in L^1 , and so our proofs extend to show close approximations to the omniscient benchmark in terms of ex ante surplus.

There are two main steps in the proof of [Theorem 1](#). The first step is developed in [Section 4.1](#) and uses the spectral decomposition (i.e., orthogonal diagonalization) of the Slutsky matrix to represent the pass-through of an intervention to prices, quantities and surpluses as a linear combination of orthogonal effects. Each of these orthogonal effects is associated with the projection of the intervention onto an eigenvector of D . This approach highlights the potential to control the effects of an intervention by designing it in such a way that it projects exclusively onto subspaces spanned by eigenvectors of D that can be accurately identified.

In the second step, developed in [Section 4.2](#), we apply a statistical method that allows us to accurately identify a subspace spanned by some eigenvectors of the Slutsky matrix from noisy observations. We then show that interventions that project exclusively onto these recoverable subspaces are robust and have the desirable welfare properties stated in [Theorem 1](#).

4.1. Spectral decomposition and pass-throughs.

4.1.1. *Spectral decomposition of price and quantity effects.* A key preliminary step is a spectral decomposition of the pass-through of an intervention to prices, quantities, and surpluses. In this section we present this decomposition from the perspective of an analyst who has perfect knowledge about the market state (D, q^0) .

Denoting by U the matrix whose ℓ^{th} -column is the ℓ^{th} eigenvector u^ℓ of D , and by Λ the matrix whose non-diagonal elements are zero and whose ℓ^{th} diagonal element is λ_ℓ , we have:

$$D = U\Lambda U^\top.$$

An intervention σ that subsidizes (or taxes) a single product will in general affect not only the prices and quantities of that product, but also those of other products, since in equilibrium these are all tied through strategic interactions. If we think of the eigenvectors u^ℓ as representing bundles, then these bundles have the important property that an intervention $\sigma \propto u^\ell$ in the direction of such a bundle, will only affect the price $u^\ell \cdot p$ and quantity $u^\ell \cdot q$ of that bundle, leaving the prices and quantities of the bundles corresponding to the other eigenvectors unchanged. Generally, we can decompose $\sigma = \sum_\ell (u^\ell \cdot \sigma) u^\ell$ into a combination of n orthogonal interventions, each in the direction of an eigenvector. We can use this decomposition to solve the oligopoly game and get simple expressions for the pass-through of the intervention, in terms of the eigenvalues of the matrix.

Lemma 1. The pass-throughs from any intervention σ to prices and quantities of each eigenvector are as follows:

$$u^\ell \cdot \dot{p} = -\frac{1}{1 + |\lambda_\ell|} u^\ell \cdot \sigma \quad \text{and} \quad u^\ell \cdot \dot{q} = \lambda_\ell (u^\ell \cdot \dot{p}) = \frac{|\lambda_\ell|}{1 + |\lambda_\ell|} u^\ell \cdot \sigma.$$

Proof. From equation (5) we get $(I - U\Lambda U^\top) \dot{p} = \dot{c}$. Multiplying both sides by U^\top we get $U^\top \dot{p} = (I - \Lambda)^{-1} U^\top \dot{c}$ and, using equation (6), $U^\top \dot{q} = \Lambda(I - \Lambda)^{-1} U^\top \dot{c}$. \square

Thus, we can study the price and quantity pass-throughs of each of these n interventions separably across eigenvectors. In particular, each unit of subsidy

in direction \mathbf{u}^ℓ *exclusively* passes through to the price and quantities of bundle \mathbf{u}^ℓ , and it does so with coefficients $-(1 + |\lambda_\ell|)^{-1}$ and $|\lambda_\ell|(1 + |\lambda_\ell|)^{-1}$, respectively.

Note that the magnitudes of the price and quantity pass-throughs in the different \mathbf{u}^ℓ are ordered according to their corresponding eigenvalues: The larger is $|\lambda_\ell|$, the less a given subsidy $\mathbf{u}^\ell \cdot \boldsymbol{\sigma}$ reduces prices, but the more it increases quantities. This asymmetry is the result of two opposing forces: On the one hand, the strategic interactions among firms imply that the equilibrium price $\mathbf{u}^\ell \cdot \mathbf{p}^0$ is less sensitive to the subsidy $\mathbf{u}^\ell \cdot \boldsymbol{\sigma}$ the larger is $|\lambda_\ell|$. On the other, the demand $\mathbf{u}^\ell \cdot \mathbf{q}^0$ is more sensitive to the price $\mathbf{u}^\ell \cdot \mathbf{p}^0$ the larger¹⁹ is $|\lambda_\ell|$; this is just a fact about the market's demand function, rather than equilibrium pricing. Lemma 1 shows that the second effect dominates the first in the sense that the larger is $|\lambda_\ell|$, the more sensitive is the equilibrium quantity $\mathbf{u}^\ell \cdot \mathbf{q}^0$ to the subsidy $\mathbf{u}^\ell \cdot \boldsymbol{\sigma}$.

4.1.2. *Surplus metrics: Spectral decomposition and robust interventions.* We now combine the spectral decomposition with the surplus formulas to deduce the following second lemma.

Lemma 2. The pass-throughs to consumer, producer and total surpluses are:

$$\begin{aligned}\dot{C} &= - \sum_{\ell=1}^n (\mathbf{u}^\ell \cdot \mathbf{q}^0) (\mathbf{u}^\ell \cdot \dot{\mathbf{p}}) \\ \dot{P} &= 2 \sum_{\ell=1}^n (\mathbf{u}^\ell \cdot \mathbf{q}^0) (\mathbf{u}^\ell \cdot \dot{\mathbf{q}}) \\ \dot{W} &= \sum_{\ell=1}^n (\mathbf{u}^\ell \cdot \mathbf{q}^0) (\mathbf{u}^\ell \cdot \dot{\mathbf{q}}).\end{aligned}$$

Proof. The effect of the intervention on consumer surplus is $\dot{C} = -\mathbf{q}^0 \cdot \dot{\mathbf{p}}$. Multiplying this equation by $\mathbf{U}\mathbf{U}^\top$ gives $\dot{C} = -\mathbf{U}^\top \mathbf{q}^0 \cdot \mathbf{U}^\top \dot{\mathbf{p}}$. Similarly, the effect of the intervention on producer surplus is $\dot{P} = \mathbf{q}^0 \cdot (\dot{\mathbf{p}} + \boldsymbol{\sigma}) + (\mathbf{p} - \mathbf{c}) \cdot \dot{\mathbf{q}} = 2\mathbf{q}^0 \cdot \dot{\mathbf{q}}$. Multiplying the equation by $\mathbf{U}\mathbf{U}^\top$ we obtain the expression of \dot{P} in the Lemma. The expression for \dot{W} is obtained by aggregating \dot{P} , \dot{C} and the expenditure of the intervention. \square

Lemma 2 shows that the effect of an intervention on consumer, producer or total surplus is a weighted sum of pass-throughs to each of the eigenvectors \mathbf{u}^ℓ —with the weight being the corresponding bundle's quantity.

We now use Lemma 2 to illustrate how—by setting $\boldsymbol{\sigma}$ equal to \mathbf{u}^1 —the authority may be able to obtain the highest total surplus per-dollar spent possible subject to the constraint that the change in consumer surplus is not negative (recall Proposition 1). Because \mathbf{u}^1 is orthogonal to all the other \mathbf{u}^ℓ with $\ell \neq 1$, such intervention only changes the price and quantity of the bundle \mathbf{u}^1 :

$$\mathbf{u}^1 \cdot \dot{\mathbf{p}} = -\frac{1}{1 + |\lambda_1|} \quad \text{and} \quad \mathbf{u}^1 \cdot \dot{\mathbf{q}} = \frac{|\lambda_1|}{1 + |\lambda_1|},$$

¹⁹Indeed, it follows from (6) that $\mathbf{U}^\top \mathbf{q}^0 = \boldsymbol{\Lambda} \mathbf{U}^\top \mathbf{p}^0$, so the slope of the demand $\mathbf{u}^\ell \cdot \mathbf{q}^0$ with respect to own price $\mathbf{u}^\ell \cdot \mathbf{p}^0$ is equal to λ_ℓ .




FIGURE 3. An illustration of the Davis–Kahan theorem: How eigenvectors of \hat{D} (perturbed matrix) project onto eigenvectors of D (true matrix) with similar eigenvalues (off by at most $\|E\|$). This relationship ensures that the subspace generated by the “top” (i.e., large-eigenvalue) eigenvectors of \hat{D} is a good approximation of the subspace generated by the top eigenvectors of D . For our economic problem this implies that interventions based on top eigenvectors of \hat{D} yield high surplus pass-through, despite noise E .

leading to an overall change in consumer and producer surplus equal to:

$$\dot{C} = \frac{1}{1 + |\lambda_1|} \dot{S} \quad \text{and} \quad \dot{P} = 2 \frac{|\lambda_1|}{1 + |\lambda_1|} \dot{S}.$$

Theorem 1 focuses on the case of $M(n) \gg b(n)$, in which case $|\lambda_1| \rightarrow \infty$ as n grows large. Consequently, the intervention has a vanishing effect (of magnitude $1/(1 + |\lambda_1|)$) on equilibrium prices and consumer surplus. However, this intervention changes the quantity of bundle u^1 by almost one, and the equilibrium quantity of each product i changes proportionally to how much i is represented in u^1 , i.e., u_i^1 . Furthermore, recoverable structure ensures that the projection of u^1 on the status quo quantity vector is positive. This ensures that the intervention $\sigma = u^1$ has a non-zero level of expenditure, i.e., $\dot{S} = u^1 \cdot q^0 > 0$. Overall, the intervention changes quantities so that the increase in total surplus is the maximum possible increase per dollar spent (as n grows large), and these gains are captured entirely by the producers.

The next subsection explains that the property of recoverable structure allows us to estimate accurately objects like u^1 from noisy estimates.

4.2. Recovering and using the subspace of dominant eigenvectors.

4.2.1. *The Davis–Kahan theorem.* Recall that if $M(n) \gg b(n)$, then $(M(n), \delta)$ -recoverable structure requires that the normalized Slutsky matrix D has eigenvalues that, in absolute value, are much larger than $b(n)$; henceforth, for brevity, we say such eigenvalues are “large.”

The key tool in our statistical exercise that leverages this assumption is the Davis–Kahan theorem. Under the hypothesis that some eigenvalues of D are large, this theorem guarantees that, despite the noise in E , the large eigenvalues and associated eigenvectors of the observed matrix are a very good approximation of the true large eigenvalues. In other words, the noise in E cannot cause

the large eigenvalues of D to be “mixed up” with the eigenvalues far away in the spectrum; see Figure 3 for an illustration. A bit more formally, this theorem has the following two central implications in our setting:

- (i) $\widehat{D} = D + E$ has some eigenvalues which are themselves large;
- (ii) the eigenvectors of \widehat{D} associated with such eigenvalues are highly correlated with eigenvectors of D with similarly large eigenvalues, i.e., the eigenvectors of \widehat{D} associated with such eigenvalues can be written (up to a small error) as linear combinations of eigenvectors of D with similarly large eigenvalues.

4.2.2. *Using recovery of eigenvectors to design a subsidy policy: A simple illustration.* This powerful result forms the core of our strategy to recover and use structure about the oligopoly demand in a way that is robust to noise. To facilitate the illustration, we make a stronger assumption on D . The assumption is that the largest eigenvalue of D is sufficiently well-separated from all other eigenvalues by a “gap” much larger than $b(n)$. In that case, the Davis–Kahan Theorem has an even stronger implication: we can use \widehat{D} to recover a normalized eigenvector \widehat{u}^1 that is almost perfectly positively correlated to the corresponding eigenvector u^1 . This stronger “gap” condition holds in Example 1 when $\gamma = 0.3$ since the gap between the two largest eigenvalues is roughly 50 and $b(n)$ is roughly 17; see also Figure 1 as an illustration of the similarity between the true u^1 (in panel B) and the estimated \widehat{u}^1 (panel D).

To use this, note that as a consequence of Lemma 1 and Lemma 2, we have

$$\dot{W} = \sum_{\ell=1}^n (u^\ell \cdot q^0)(u^\ell \cdot \sigma) \frac{|\lambda_\ell|}{1 + |\lambda_\ell|}, \quad (8)$$

while expenditure is

$$\dot{S} = \sum_{\ell=1}^n (u^\ell \cdot q^0)(u^\ell \cdot \sigma). \quad (9)$$

Let us use the recovered \widehat{u}^1 to design an intervention with $\sigma \propto \widehat{u}^1$. The Davis–Kahan Theorem tells us that we can essentially take \widehat{u}^1 to be its true counterpart u^1 with a very small error, so from now on we will talk as if we know u^1 . Let us flip the sign of σ if necessary so that $(u^1 \cdot q^0)(u^1 \cdot \sigma)$ is positive. Then we can see from the equations above that \dot{W} will in fact be very similar to \dot{S} , since $\lambda_1 \gg b(n)$ is large and so $\frac{|\lambda_1|}{1 + |\lambda_1|} \approx 1$. Moreover, if we know $u^1 \cdot q^0$ and this differs from zero (which is a requirement of recoverable structure), we can scale the intervention to be of the size that we desire, and achieve $\dot{S} = 1$.

The argument just described contains some wishful thinking, however. When we arranged the sign of σ so that $(u^1 \cdot q^0)(u^1 \cdot \sigma)$ is positive, we did not consider that we have only a noisy observation \widehat{q}^0 of q^0 . So part of the challenge of the proof is to manage the observation error that makes \widehat{q}^0 different from q^0 , and to show that we can obtain a correct estimate of the sign with probability tending to 1 as $n \rightarrow \infty$. If we don’t manage to do this correctly, our intervention actually decreases efficiency with positive probability. This explains the need for the second part of Assumption 4 on ε in the signal.

However, assuming that noise is bounded is not enough: if $u^1 \cdot q^0$ is very small, then there may be no hope for consistently recovering the true magnitude

or sign of $\mathbf{u}^1 \cdot \mathbf{q}^0$ from the signal $\widehat{\mathbf{u}}^1 \cdot (\mathbf{q}^0 + \boldsymbol{\varepsilon})$ even if the noise is well-behaved: the asymptotically small noise may still overwhelm a similarly decaying underlying mean $\mathbf{u}^1 \cdot \mathbf{q}^0$. Such an unrecoverability would make it impossible to orient and scale our intervention appropriately. The definition of recoverable structure ensures that this problem does not occur by requiring that the projection of \mathbf{q}^0 onto eigenvectors with large eigenevalues is bounded away from zero.

The special case of the result that this discussion makes plausible is this: If the largest eigenvalue of \mathbf{D} is well-separated from others and if $\mathbf{q}^0 \cdot \mathbf{u}^1$ is not vanishingly small, then a subsidy profile that is proportional to \mathbf{u}^1 can, if it is suitably scaled, achieve all the properties of [Theorem 1](#).

4.2.3. The more general result. The proof of the main result improves on this sketch in two ways. First, it does not impose that we need to recover exactly \mathbf{u}^1 , but rather uses a much larger eigenspace of $\widehat{\mathbf{D}}$. Indeed, the general intervention projects $\widehat{\mathbf{q}}^0$ onto $\mathcal{L}(\mathbf{D}, M(n))$, the eigenspace of all eigenvectors of \mathbf{D} larger than $M(n)$. This makes it much easier for the analogue of $\mathbf{u}^1 \cdot \mathbf{q}^0$ not to be too small, since the projection of \mathbf{q}^0 onto a larger eigenspace will have a larger norm. Second, the general proof dispenses with assuming that any eigenvalues are well-separated. Instead, it deals with any possible spectrum of \mathbf{D} subject to our maintained assumptions. This is a considerable complication, because it is no longer generally possible to recover any true eigenvector \mathbf{u}^ℓ with any accuracy. We instead work directly with a recovered eigenspace $\mathcal{L}(\mathbf{D}, M(n))$ that generalizes the span of $\widehat{\mathbf{u}}^1$ and show that while we know little about the individual true eigenvectors of \mathbf{D} giving rise to this space, we can use the fact that all of them have large eigenvalues to generalize our argument for showing that (8) and (9) can be made very close and nonzero with a feasible intervention. This is where the arguments are more subtle and require going beyond standard applications of the Davis–Kahan theorem.

5. ILLUSTRATION: INTERVENTIONS BASED ON NOISY DEMAND ESTIMATES

We present simulations that illustrate the central ideas behind [Theorem 1](#). In particular, within the context of [Example 1](#), we illustrate the effects of the intervention rule that taxes and subsidizes firms in the direction of the eigenvector associated with the largest eigenvalue of the observed Slutsky matrix; we do so by evaluating the effect of this intervention in many randomly generated market states.

5.1. Generating the true and observed market state. We construct the true and observed Slutsky matrices \mathbf{D} and $\widehat{\mathbf{D}}$ as we described in [Example 1](#). In particular, $n = 300$, $\widehat{\mathbf{D}} = \mathbf{D} + \mathbf{E}$ and

$$\mathbf{D} = (1 - \gamma)\mathbf{D}_{\text{block}} + \gamma(-\mathbf{Z}\mathbf{Z}').$$

Recall that $\mathbf{D}_{\text{block}}$ has three blocks of equal size, products in the same block are complements and products across block are substitutes. The parameter $\gamma \in (0, 1)$ controls the size of the large eigenvalues of \mathbf{D} , with low values of γ creating large eigenvalues by weighting the block matrix \mathbf{D} whose eigenvalues grow linearly (and so much faster than $b(n) \sim n^{1/2}$).

The true initial quantity vector \mathbf{q}^0 has some regular block structure but also some idiosyncratic heterogeneity. It is constructed as follows:

$$q_i^0 = (\mathbf{q}_{\text{block}})_i X_i$$

Here, the quantity vector $\mathbf{q}_{\text{block}}$ provides a base quantity for each product that depends on its associated block (0.1 for products in the first two blocks, and 3 for the products in the third block). The random variable X_i is drawn independently of all others, and its logarithm is normal with variance 0.1 and mean 1. We use a multiplicative perturbation to avoid negative quantities.

The observed quantities are given by

$$\hat{q}_i^0 = q_i^0 Y_i$$

where Y_i is an independent error with the same distribution as X_i . This can be rewritten in terms of our additive error model, with $\varepsilon_i = q_i^0(Y_i - 1)$. Here again, the multiplicative error model avoids negative quantities.

5.2. Interventions. We focus on the following intervention rule: Recover the eigenvector associated to the largest eigenvalue, in absolute terms, of the estimated Slutsky matrix $\hat{\mathbf{D}}$. Intervene to subsidize firms in proportion to this eigenvector. This intervention aligns with the intervention rule behind our Theorem 1.²⁰

In order to meaningfully compare the effects of such intervention for different realizations of the market state, we scale the size of all interventions that we consider by requiring that they have the same expenditure based on the observed quantity vector.²¹

5.3. Evaluation of interventions. We consider different values of $\gamma \in [0, 1]$. For each of these values we generate 3000 market states according to the above description and we compute the changes in consumer and producer surplus under the true market state. Figure 4 summarizes this exercise: for each value of γ considered, it reports the median (blue dot) of the change in consumer surplus (panel A) and of the change in producer surplus (panel B) and the respective 5th and 95th percentiles associated with the 3000 market state realizations.

Figure 4 shows a sharp transition in the performance of the intervention. For γ less than roughly 0.7 the true market state has very large eigenvalues and so the authority can use the estimate of \mathbf{D} to precisely identify the underlying main eigenvector (see Figure 1 for $\gamma = 0.3$.) Note also that products in category 3 are the ones who are most substitutable with other products in categories 1 and 2 and so they are highly represented in the first eigenvector (i.e., they have a high eigenvector centrality). This implies that the estimated first eigenvector is sufficiently correlated to the status quo market quantity. Hence, for low γ , the true market state has recoverable structure. This allows the authority to implement interventions that robustly have the pass-through properties characteristic of high-eigenvalue eigenvectors—negligible impact on prices and

²⁰In this example, for low values of γ the largest eigenvalue of \mathbf{D} is sufficiently well-separated from all other eigenvalues and, consequently, the authority can use the first eigenvector of $\hat{\mathbf{D}}$ as a good approximation of \mathbf{u}^1 (see the illustrative discussion in Section 4.2.2).

²¹More precisely, we first project the observed quantity vector onto the recovered eigenvector, and use that to predict the expenditure size.




FIGURE 4. The effects of first eigenvector intervention on changes in consumer surplus (panel a) and changes in producer surplus (panel b) as a function of γ in a market described in Section 5.1. For each value of γ , we plot the median (in blue), 5th and 95th percentiles associated with 3000 market state realizations.

hence on consumers, along with an effect on producer surplus equal to twice the authority's spending γ , which is normalized to one unit.

However, as γ grows larger than 0.7, the property of recoverable structure fails (see Figure 2 for $\gamma = 0.9$) with the consequence that an intervention that taxes and subsidizes firms based on the estimated first eigenvector is very unpredictable and risky. The unpredictability is shown by the fast widening of the error bars as γ increases beyond 0.7. The riskiness is shown by the fact that for over a third of the outcomes, the realized change in producer surplus, and hence in total surplus, is negative.

In this example, the demand for products in block 3 is significantly higher than the demand for the other products. Hence, to a first approximation, consumer surplus increases when the price of products in block 3 decreases, while producer and total surplus increase when the quantity of these products increases.

When the authority can estimate the first eigenvector accurately, the first eigenvector intervention turns out to subsidize products in block 3 and tax all the other products. Subsidizing products in block 3 leads to a decrease in their price and hence an increase in their demand. Taxing products in blocks 1 and 2 leads to an increase of the price of these products, and hence an increase in the demand of products in block 3. Combining the effects, the intervention leads to a relatively high increase in the demand of good 3, while keeping prices roughly constant. As a result, both producer and total surplus increase dramatically without sizeable changes in consumer surplus.

This management of the spillovers can be achieved only by statistically identifying some relevant latent market structure from the noisy demand measurements. While in this example that structure takes the simple form of product categories, in general it might be less easy to describe, and yet equally useful for the design of robust interventions.

6. TIGHTNESS OF THE MAIN RESULT

We have established that if demand has recoverable structure, the authority can robustly achieve the maximum possible total surplus per dollar spent subject to the constraint that consumers are not harmed. The associated intervention rule boosts production with minimal price changes, resulting in all efficiency gains being captured by firms.

This raises two natural questions. First, can we find interventions that robustly increase total surplus when the demand does not have recoverable structure? Second, when there is recoverable structure, is the structure of robust interventions that we have just described in any sense necessary? For example, could we have found interventions that robustly increase consumer surplus, rather than leaving it unchanged?

In this section, we work with the case we have focused on in other illustrations, where \mathbf{E} has i.i.d. entries with a standard deviation that does not depend on n , giving $b(n) \sim n^{1/2}$; this choice is immaterial and the results apply to a wide range of alternative noise structures.

Proposition 2. The following hold:

1. There are environments such that (i) for any $\epsilon > 0$, there are no intervention rules that ϵ -robustly increase total surplus ($\dot{W} > 0$) for any n ; (ii) there are interventions achieving $\dot{P} = 2\dot{S}$ and $\dot{C} = 0$ for any $\dot{S} \geq 0$.
2. There are environments satisfying $(M(n), \delta)$ -recoverable structure with $\frac{M(n)}{b(n)} \rightarrow \infty$ and $\delta > 0$ such that (i) for any $\epsilon > 0$, there are no intervention rules that ϵ -robustly achieve $\dot{C} > \epsilon$ for all n ; (ii) there are interventions achieving $\dot{C} = \dot{S}$ and $\dot{P} = 0$ for any $\dot{S} \geq 0$.

In each case, [Proposition 2](#) describes the limits on what can be achieved by an authority with noisy information. It also notes that these limits really are about information: part (ii) of each case states that an omniscient authority would not be subject to the same limitation.

In more detail, Part 1 of [Proposition 2](#) tells us that we may not be able to design interventions that robustly increase total surplus $\dot{W} = \dot{C} + \dot{P} - \dot{S}$. (Since, by [Theorem 1](#) we know this can be done under the recoverable structure assumption, our construction must lack recoverable structure.) Intuitively, in this case the information about the market state learned from the signal can be very imprecise and, therefore, there are market states in which any intervention will lead to undesirable outcomes. The proof constructs a set of market states such that, with an uninformative signal, for any intervention there is a market state with $\dot{W} < 0$. The basic idea is to use the total surplus decomposition:

$$\dot{W} = \sum_{\ell} (\mathbf{u}^{\ell} \cdot \mathbf{q}^0)(\mathbf{u}^{\ell} \cdot \boldsymbol{\sigma}) \frac{|\lambda_{\ell}|}{1 + |\lambda_{\ell}|} \quad (10)$$

and construct the example so that the authority cannot accurately predict the signs of the terms for any given $\boldsymbol{\sigma}$.

Part 2 of [Proposition 2](#) tells us that even if the market state has recoverable structure, it may be impossible to design interventions that robustly increase total surplus *and* allow consumers to capture some of the resulting efficiency gains. [Lemma 2](#) tells us that to achieve such an outcome, the intervention must

project onto some u^ℓ where λ_ℓ is not too large, since only those eigenvectors have nonvanishing pass-through to consumer surplus. However, the noisy observation of \widehat{D} and q^0 gives very noisy estimates of the constituents of (10) corresponding to these eigenvectors. So, by targeting them, there is a sizeable chance (at least in some environments) that the policy will have negative consequences for the consumers.

7. DISCUSSION AND CONCLUDING REMARKS

We have developed a theory of robust interventions in large oligopolies. We identify a condition on demand under which an authority can robustly increase the total surplus per dollar spent as much as would be possible under perfect information subject to the constraint that consumers are not harmed. The methodological contribution is the development of spectral methods to analyze pass-through in oligopolies, and the application of these methods to gain leverage on statistical problems about oligopolies observed with noise.

We conclude with some observations about the scope of our analysis and connections to related research.

7.1. Marketplace big data in practice. In our model, the authority has “big data” about the Slutsky matrix D that may not allow precise estimation of any pairwise demand interactions or hedonic model parameters. This seems natural for markets with a large and changing collection of goods, such as those hosted on large online platforms. Such platforms collect immense amounts of data—about consumer browsing behavior, timing of purchases, consideration sets, etc.—and apply machine learning techniques to these data to form informative but imperfect estimates of interactions among various goods (Athey, 2018; Wager and Xu, 2021; Cai and Daskalakis, 2022; Bajari, Burdick, Imbens, Masoero, McQueen, Richardson, and Rosen, 2023). This is modeled by our notion of a market signal.²² Our paper offers an approach for calculating suitable statistics that suffice for effective interventions despite the noise in this signal.

Our approach contrasts with a standard one in empirical industrial organization, where markets are defined tightly so that each contains only a relatively small number of similar goods with strong demand interactions, and then a small number of hedonic parameters and demand elasticities are precisely estimated. That approach would correspond in our notation to a very precise signal (i.e., an error matrix E with small norm) and a small number of goods.

In terms of interventions, online marketplaces are an appealing setting for our results because the authorities in charge of them can finely target policies that function as taxes and subsidies, including, e.g., commission rates, discount coupons, free advertising, etc.—and regularly experiment with such perturbations. It is worth noting, however, that the policies our analysis recommends need not be specific to individual products. This is because when we take a large matrix (in our case, the Slutsky matrix) that reflects relationships among

²²The distributional properties of the matrix E capturing the errors in these estimates would depend on the application, as would the conditions for E to have small enough norm for recoverable structure to be present. It would be interesting to investigate these issues in specific applications.

units (in our case, products) and look at the eigenvectors with large eigenvalues, the coordinates of those eigenvectors typically yield low-dimensional embeddings capturing substantively natural categories (Chen, Chi, Fan, and Ma, 2021). For instance, in our Example 1, the top two eigenvectors are sufficient to recover the blocks to which the goods belong. Relatedly, *spectral clustering* analyses based on the top few eigenvectors sort items into natural “similarity” classes, where similarity is defined by having similar relationships to other classes (Spielman and Teng, 1996). Once again, the spectral statistics used in these techniques tend to pick up interpretable “broad” features of the products, rather than idiosyncracies specific to individual products. As a result, the policies our interventions recommend—which project all variation onto these eigenvectors—will often be close to a policy that depends mostly on category—e.g., a subsidy on the sale of smartphones along with a tax on certain types of accessories. Though the policies will not be perfectly regular (note the irregularities of Panel D of Figure 1) the above observations lead us to conjecture that an authority constrained to design policies that discriminate only at a coarse product level could, under natural assumptions, achieve a substantial amount of the gains of our policies. We leave these interesting considerations to future work.

7.2. Relationship with hedonic utility models. Recent work by Pellegrino (2021) and Ederer and Pellegrino (2021) uses an oligopoly model to empirically quantify the evolution of market power. The relationship of our model to their work sheds light on the types of empirical models that can capture recoverable structure. In Pellegrino (2021), the model of demand is hedonic, in the spirit of Lancaster (1966): the household’s utility is additively separable in the contributions of various *characteristics*, and a product provides a bundle of these characteristics. The Slutsky matrix D derived from this demand model can be expressed as a transformation of the cosine similarity matrix of products’ characteristics, which was estimated for a large set of consumer goods by Hoberg and Phillips (2016) using text data.²³ We have calculated that in the Slutsky matrix derived this way, the eigenvalues are all small and the recoverable structure condition fails.²⁴ It is useful to reflect on why this is the case.

In the model of Pellegrino (2021), a simple calculation shows that it is not possible for the Slutsky matrix to have large eigenvalues.²⁵ For an economic intuition, note that in the Lancaster (1966) type of model, the “direct” relationship between any pair of goods is substitution. With substitution, if some demand

²³Pellegrino (2021) and Ederer and Pellegrino (2021) consider quantity competition, but the Slutsky matrix does not depend on this choice.

²⁴There are more than $n > 3000$ products in the data and for the case of i.i.d. noise we would require that the largest eigenvalue is considerably larger than $b(n) = \sqrt{n} \approx 54$; this fails as largest eigenvalue of D in absolute value is about 2.

²⁵The (un-normalized) Slutsky matrix in Pellegrino (2021) is $-B^{-1}$, where $B = I + \alpha(\Sigma - I)$. The matrix Σ is positive semidefinite because it can be written as $V^T V$, where the columns of V are the characteristic vectors of various products. Thus all eigenvalues of B are real numbers bounded below by $1 - \alpha$, and all eigenvalues of $-B^{-1}$ are at most $1/(1 - \alpha)$ in magnitude. Pellegrino uses the value $\alpha = 0.12$, which prevents any eigenvalue from exceeding 1.13. We do not work with exactly the same Slutsky matrix because of the normalization in Appendix A.1. Its eigenvalues are a bit different, but they can still be bounded by a constant. Numerically we see that the normalization makes little difference.

is diverted from one good due an increase in its price, the sum of effects on all substitute goods is bounded, since, loosely speaking, the demand gained by these other goods must come out of the demand lost by the more expensive one. So the sum of positive entries in D corresponding to this effect is bounded. This in turn places bounds on any complementarities in the Slutsky matrix.²⁶ In short, in a hedonic model where the basic force is substitution, overall spillovers are bounded, and the fact that D has no large eigenvalues is the mathematical manifestation of this.

Quite different behavior is seen in models where utility arises directly out of the consumption of complementary goods, and such complementarities play a central role in our examples of recoverable structure. A leading practical example comes from the use of computers: a consumer's utility from a computer depends on the hardware, the operating system, and the applications. Two firms selling distinct components—a hardware device and an operating system, for instance—are supplying complementary goods; on the other hand, two firms selling the same component (say, operating systems) are supplying substitute goods (Matutes and Regibeau, 1988, 1992). The illustrative example 1 we develop in Section 3 shows how Slutsky matrices with large eigenvalues arise naturally in such settings.²⁷ But, as we have seen, it is impossible to produce the same patterns in models of the Lancaster (1966) type, because they cannot have large eigenvalues; one would need to add terms corresponding to the idea that some characteristics provide more value when enjoyed together. There is a simple economic intuition for why such complementarities are more likely to produce recoverable structure: when the price of one good goes down, then all its complements can become more demanded by comparable nonvanishing amounts if they are not substitutes for one another. This leaves the kinds of clusters of nonvanishing entries in D which are the hallmark of recoverable structure.

In summary, direct complementarities of this sort seem practically important, as a natural source of the recoverable structure that plays a central role in our results. We hope these observations will motivate further empirical research on the structure of large-scale oligopoly models with complementary goods.

7.3. Games on networks. One can view our exercise as a special case of an intervention, under noisy information, in a game among a large number of agents. In our case, the game is a standard oligopoly pricing model. Note that, under the assumption of linear demand, the pricing game can be seen as a network game with linear best replies where the Slutsky matrix is the network. Our analysis shows that if the oligopoly exhibits a recoverable structure, then there are robust interventions that maximise changes in producer surplus. The methods we have developed can be extended to other settings. For example, in a public goods setting, interventions aim to realign private marginal returns with social marginal returns. The literature has developed

²⁶Note that complementarity—demand of one good is decreasing in the price of the other—can arise in Pellegrino (2021) model. This happens through indirect effects: the substitute of my substitute can be my complement. However, since the “direct” substitution effect is bounded in magnitude, so are the indirect consequences.

²⁷The complementarities there happen to be within-category, but that is not important for our point here.

tools to understand how to do this when the authority has precise information on the spillovers causing the underprovision of public goods.²⁸ However, we know little about the design of interventions under noisy information about such externalities. Similarly, in contracting for teams under moral hazard, network methods have recently been developed for locally perturbing contracts to achieve better outcomes for a principal (Dasaratha et al., 2024). But it is a considerable challenge to extend these results to the realistic case where the strategic interactions among members of an organization are only imperfectly known. General games will lack some of the structure we have taken advantage of. This includes the properties of the spillovers structure coming from a symmetric, positive semidefinite Slutsky matrix. So there are challenges to overcome in extending our results. We hope this paper stimulates research in these exciting new directions.

7.4. Nonlinear demand. We have assumed that demand is exactly linear in a neighborhood around the status quo equilibrium point. This assumption implies that the pass-through of an intervention to prices and quantities and, therefore, to welfare, depends only on the Slutsky matrix D . We use this simplification to develop new concepts useful for robust market interventions. These concepts can be extended to nonlinear demand settings. We briefly explain how.

When demand is not locally linear, the pass-through of cost shocks depends not only on the Slutsky matrix (which is the Jacobian of demand) but also on the Hessian of the demand function, the matrix whose (i, j) entry is $\partial^2 q_i(\mathbf{p})/\partial p_i \partial p_j$ (see, e.g., Miklos-Thal and Shaffer (2021)). Thus, our calculations would change, and there is no guarantee that our linear tax/subsidy interventions based only on the Slutsky matrix at the status quo, would perform as they do in the linear model.

However, our main result can be extended once we allow the authority to use nonlinear interventions, i.e., to commit a vector of functions specifying a payment to each producer i as a function of all prices and quantities realized after the intervention. With this broader set of instruments, the authority can use nonlinear rebates based on post-intervention quantities to reduce the problem to the one we have studied. The key idea is to effectively linearize the demand the firms face around the status quo—i.e., use transfers to make up the difference between realized demand and a linear demand function. Once demand has been “linearized” in this way, the problem that firms face becomes equivalent to the one we have studied and we can use the results developed to design per-unit tax/subsidy interventions with desirable welfare properties. If we assume that the curvature of the demand of each product is locally bounded by a known constant, the payments needed to linearize demand can be bounded by a small fraction of the first-order gains of an intervention, so our welfare guarantees go through. Such assumptions on curvature also allow us specify concrete sizes of interventions that achieve a given level of welfare gain, rather than just characterizing the behavior of derivatives.

²⁸See, for instance, Bramoullé et al. (2014) and Elliott and Golub (2019).

REFERENCES

- AKBARPOUR, M., S. MALLADI, AND A. SABERI (2020): "Just a few seeds more: value of network information for diffusion," *Available at SSRN 3062830*.
- ATHEY, S. (2018): "The impact of machine learning on economics," *The economics of artificial intelligence: An agenda*, 507–547.
- ATHEY, S., M. BAYATI, N. DOUDCHENKO, G. IMBENS, AND K. KHOSRAVI (2021): "Matrix completion methods for causal panel data models," *Journal of the American Statistical Association*, 116, 1716–1730.
- AZAR, J. AND X. VIVES (2021): "General equilibrium oligopoly and ownership structure," *Econometrica*, 89, 999–1048.
- BAI, Z. D. AND Y. Q. YIN (1988): "Necessary and sufficient conditions for almost sure convergence of the largest eigenvalue of a Wigner matrix," *Annals of Probability*, 16, 1729–1741.
- BAJARI, P., B. BURDICK, G. W. IMBENS, L. MASOERO, J. MCQUEEN, T. S. RICHARDSON, AND I. M. ROSEN (2023): "Experimental design in marketplaces," *Statistical Science*, 38, 458–476.
- BALLESTER, C., A. CALVÓ-ARMENGOL, AND Y. ZENOU (2006): "Who's who in Networks. Wanted: the Key Player," *Econometrica*, 74, 1403–1417.
- BANERJEE, A., A. G. CHANDRASEKHAR, E. DUFLO, AND M. O. JACKSON (2013): "The diffusion of microfinance," *Science*, 341, 1236–1239.
- BELHAJ, M. AND F. DEROIAN (2017): "Contracting on Networks," Mimeo.
- BLOCH, F. AND N. QUEROU (2013): "Pricing in Social Networks," *Games and Economic Behavior*, 80, 263–281.
- BORGATTI, S. (2006): "Identifying sets of key players in a social network," *Computational and Mathematical Organization Theory*, 12, 21–34.
- BRAMOULLÉ, Y., R. KRANTON, AND M. D'AMOURS (2014): "Strategic interaction and networks," *The American Economic Review*, 104, 898–930.
- CAI, Y. (2022): "Linear Regression with Centrality Measures," .
- CAI, Y. AND C. DASKALAKIS (2022): "Recommender systems meet mechanism design," in *Proceedings of the 23rd ACM Conference on Economics and Computation*, 897–914.
- CANDOGAN, O., K. BIMPIKIS, AND A. OZDAGLAR (2012): "Optimal Pricing in Networks with Externalities," *Operations Research*, 60, 883–905.
- CHANDRASEKHAR, A. G., P. GOLDSMITH-PINKHAM, T. MCCORMICK, S. A. THAU, AND J. WEI (2024): "Non-robustness of diffusion estimates on networks with measurement error," Tech. rep., National Bureau of Economic Research.
- CHEN, Y., Y. CHI, J. FAN, AND C. MA (2021): "Spectral methods for data science: A statistical perspective," *Foundations and Trends® in Machine Learning*, 14, 566–806.
- CHERNOZHUKOV, V., C. HANSEN, Y. LIAO, AND Y. ZHU (2023): "Inference for low-rank models," *The Annals of statistics*, 51, 1309–1330.
- DALLAPORTA, S. (2012): "Eigenvalue variance bounds for Wigner and covariance random matrices," *Random Matrices: Theory and Applications*, 1.
- DASARATHA, K. (2020): "Distributions of centrality on networks," *Games and Economic Behavior*, 122, 1–27.
- DASARATHA, K., B. GOLUB, AND A. SHAH (2024): "Incentive Design With Spillovers," SSRN Working Paper, available at SSRN: <https://ssrn.com/abstract/4688888>

[//ssrn.com/abstract=4853054](https://ssrn.com/abstract=4853054) or <http://dx.doi.org/10.2139/ssrn.4853054>.

- DAVIS, C. AND W. M. KAHAN (1970): "The rotation of eigenvectors by a perturbation. III," *SIAM Journal on Numerical Analysis*, 7, 1–46.
- DE LOECKER, J., J. EECKHOUT, AND G. UNGER (2020): "The rise of market power and the macroeconomic implications," *The Quarterly Journal of Economics*, 135, 561–644.
- DEMANGE, G. (2017): "Optimal Targeting Strategies in a Network under Complementarities," *Games and Economic Behaviour*, 105, 84–103.
- DIXIT, A. (1979): "Price changes and optimum taxation in a many-consumer economy," *Journal of Public Economics*, 11, 143–157.
- (1986): "Comparative statics for oligopoly," *International Economic Review*, 27, 107–122.
- DZIUBINSKI, M. AND S. GOYAL (2017): "How Do You Defend a Network?" *Theoretical Economics*, 12, 331–376.
- EDERER, F. AND B. PELLEGRINO (2021): "A Tale of Two Networks: Common Ownership and Product Market Rivalry," Working paper, University of Maryland.
- ELLIOTT, M. AND A. GALEOTTI (2019): "The role of networks in antitrust investigations," *Oxford Review of Economic Policy*, 35, 614–637.
- ELLIOTT, M. AND B. GOLUB (2019): "A Network Approach to Public Goods," *Journal of Political Economy*, 127, 730–776.
- ELLIOTT, M., B. GOLUB, AND M. O. JACKSON (2014): "Financial Networks and Contagion," *American Economic Review*, 104, 3115–53.
- FUDENBERG, D. AND J. TIROLE (1991): *Game Theory*, MIT press.
- GAITONDE, J., J. KLEINBERG, AND É. TARDOS (2021): "Polarization in geometric opinion dynamics," in *Proceedings of the 22nd ACM Conference on Economics and Computation*, 499–519.
- GALEOTTI, A., B. GOLUB, AND S. GOYAL (2020): "Targeting interventions in networks," *Econometrica*, 88, 2445–2471.
- GALEOTTI, A. AND S. GOYAL (2009): "Influencing the Influencers: a theory of strategic diffusion," *The Rand Journal of Economics*, 40, 509–532.
- GALEOTTI, A., S. GOYAL, M. JACKSON, F. VEGA-REDONDO, AND L. YARIV (2010): "Network Games," *Review of Economic Studies*, 77, 218–244.
- GALEOTTI, A. AND B. W. ROGERS (2013): "Strategic Immunization and Group Structure," *American Economic Journal: Microeconomics*, 5, 1–32.
- GOLUB, B. AND M. O. JACKSON (2012): "How homophily affects the speed of learning and best-response dynamics," *The Quarterly Journal of Economics*, 127, 1287–1338.
- GOYAL, S. (1996): "Interaction structure and social change," *Journal of Institutional and Theoretical Economics (JITE)/Zeitschrift für die gesamte Staatswissenschaft*, 472–494.
- HOBERG, G. AND G. PHILLIPS (2016): "Text-based network industries and endogenous product differentiation," *Journal of political economy*, 124, 1423–1465.
- LANCASTER, K. J. (1966): "A new approach to consumer theory," *Journal of political economy*, 74, 132–157.
- LEDUC, M. V., M. O. JACKSON, AND R. JOHARI (2017): "Pricing and referrals in diffusion on networks," *Games and Economic Behavior*, 104, 568–594.

- LIU, E. AND A. TSYVINSKI (2024): “A Dynamic Model of Input–Output Networks,” *Review of Economic Studies*, rdae012.
- MARSHALL, A. (1890): *Principles of Economics*, New York: Macmillan.
- MATUTES, C. AND P. REGIBEAU (1988): ““ Mix and match”: product compatibility without network externalities,” *The RAND Journal of Economics*, 221–234.
- (1992): “Compatibility and bundling of complementary goods in a duopoly,” *The Journal of Industrial Economics*, 37–54.
- MCLENNAN, A. (2018): *Advanced fixed point theory for economics*, vol. 25, Springer.
- MIKLOS-THAL, J. AND G. SHAFFER (2021): “Pass-through as an economic tool: on exogenous competition, social incidence, and price discrimination,” *Journal of Political Economy*, 129, 323–335.
- NOCKE, V. AND N. SCHUTZ (2017): “Quasi-linear integrability,” *Journal of Economic Theory*, 169, 603–628.
- (2018): “Multiproduct Firm Oligopoly: An Aggregative Games Approach,” *Econometrica*, 86, 523–557.
- NOCKE, V. AND M. WHINSTON (2022): “Concentration Thresholds for Horizontal Mergers,” *American Economic Review*, forthcoming.
- NORRIS, J. J. (2024): “ Bounding High Dimensional Comparative Statics,” *Mimeo*.
- OLLÁR, M. AND A. PENTA (2023): “A network solution to robust implementation: The case of identical but unknown distributions,” *Review of Economic Studies*, 90, 2517–2554.
- PARISE, F. AND A. OZDAGLAR (2023): “Graphon games: A statistical framework for network games and interventions,” *Econometrica*, 91, 191–225.
- PELLEGRINO, B. (2021): “Product Differentiation and Oligopoly: a Network Approach,” Working paper, University of Maryland.
- PIGOU, A. (1920): *The Economics of Welfare*, London: Macmillan.
- ROSTEK, M. AND N. YODER (2023): “Reallocative auctions and core selection,” *Mimeo*.
- SPIELMAN, D. A. AND S.-H. TENG (1996): “Spectral partitioning works: Planar graphs and finite element meshes,” in *Proceedings of 37th conference on foundations of computer science*, IEEE, 96–105.
- VIVES, X. (1999): *Oligopoly pricing: old ideas and new tools*, London, England: MIT press.
- WAGER, S. AND K. XU (2021): “Experimenting in equilibrium,” *Management Science*, 67, 6694–6715.
- WEYL, E. G. AND M. FABINGER (2013): “Pass-through as an economic tool: Principles of incidence under imperfect competition,” *Journal of Political Economy*, 121, 528–583.

APPENDIX A. OMITTED PROOFS AND DETAILS FOR MAIN RESULTS

A.1. Normalization of spillover matrix. For any differentiable 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 \mathbf{q} .

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

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

and by the chain rule we have that

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

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

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

Now set $\gamma_i = 1/\sqrt{|D_{ii}^q|}$. 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 .

A.2. Proof of Theorem 1. We prove the theorem under Assumptions 1–3 but we replace [Assumption 4](#) with the following weaker assumption:

Assumption 5. We assume that:

- (1) $\mathbb{E}[\|\mathbf{E}\|] \leq b(n)$;
- (2) for any sequence of linear subspaces $V(n)$ of \mathbb{R}^n with dimension $d(n)$, where $d(n)/n \rightarrow_n 0$, the norm $\|P_{V(n)}\varepsilon\|$ tends to 0 in probability.

Note that parts (1) of [Assumption 4](#) and [Assumption 5](#) are identical. However, part (2) of [Assumption 4](#), which stated that all ε_i are independent, implies part (2) of [Assumption 5](#).

Some notation: For any matrix \mathbf{M} , we define $\Lambda(\mathbf{M}, \underline{\lambda})$ as the set of eigenvalues of \mathbf{M} with absolute value greater than or equal to $\underline{\lambda}$ and $\mathcal{L}(\mathbf{M}, \underline{\lambda})$ as the space spanned by corresponding eigenvectors.

Recall that under [Assumption 5](#), a signal with $b(n)$ -bounded noise can be written $\hat{\mathbf{D}} = \mathbf{D}(\mathbf{p}^0) + \mathbf{E}$, where, by Markov's inequality,

$$\|\mathbf{E}\| \leq kb(n) \tag{11}$$

with high probability, where k is a large constant; we can absorb this constant into $b(n)$, which we will do from now on.

Recall also the conditions of the theorem imply that we have sequences $b(n)$ and $M(n)$ such that, for large enough n , \mathbf{D} has eigenvalues exceeding $M(n)$ where²⁹ $b(n) \ll M(n)$ and $\Lambda(\mathbf{D}, M(n))$ is nonempty. We also choose two other sequences $\underline{M}(n)$ and $\widehat{M}(n)$, such that $\underline{M}(n) \ll \widehat{M}(n) \ll a(n)$ and the differences between these successive sequences also dominate $b(n)$.

We use the following notation:

$$\begin{aligned} \underline{\Lambda}(n) &:= \Lambda(\mathbf{D}; \underline{M}(n)), & \widehat{\Lambda}(n) &:= \Lambda(\widehat{\mathbf{D}}; \widehat{M}(n)), & \Lambda(n) &:= \Lambda(\mathbf{D}; M(n)) \\ \underline{\mathcal{L}}(n) &:= \mathcal{L}(\mathbf{D}, \underline{M}(n)), & \widehat{\mathcal{L}}(n) &:= \mathcal{L}(\widehat{\mathbf{D}}, \widehat{M}(n)), & \mathcal{L}(n) &:= \mathcal{L}(\mathbf{D}, M(n)). \end{aligned}$$

²⁹We use the notation $a(n) \ll b(n)$ to mean that $a(n)/b(n) \rightarrow_n 0$.

Let P_V be the projection operator onto subspace V and P_V^\perp its orthogonal complement. Let $(\lambda_1, \mathbf{u}^1), \dots, (\lambda_n, \mathbf{u}^n)$ be eigenpairs of \mathbf{D} , with $|\lambda_1| \geq |\lambda_2| \geq \dots \geq |\lambda_n|$.

We now define our intervention:

$$\boldsymbol{\sigma} = \frac{P_{\hat{L}(n)} \hat{\mathbf{q}}^0}{\|P_{\hat{L}(n)} \hat{\mathbf{q}}^0\|^2} \quad (12)$$

The expenditure of this intervention is

$$\dot{S} = \boldsymbol{\sigma} \cdot \mathbf{q}^0 = \frac{P_{\hat{L}(n)} \hat{\mathbf{q}}^0}{\|P_{\hat{L}(n)} \hat{\mathbf{q}}^0\|^2} \cdot \mathbf{q}^0 \quad (13)$$

Our first main lemma, which we will prove shortly, will assert that \dot{S} converges in probability to 1. The challenge in proving such a result is that the actual expenditure depends on true quantities, whereas the intervention is built based on *estimated* quantities projected onto an “estimated” eigenspace of $\hat{\mathbf{D}}$. We begin with a technical result that will be key to controlling the differences between actual and estimated objects. It relies on the Davis–Kahan Theorem, which is important for our analysis.

Lemma 3. The norm $\|P_{\underline{L}(n)}^\perp P_{\hat{L}(n)}\|$ converges to 0 in probability as $n \rightarrow \infty$.

Proof. The Davis–Kahan Theorem guarantees that, for any $\eta > 0$, there exists $N_1(\eta)$ such that for all $n > N_1(\eta)$, with probability greater than or equal to $1 - \eta$, two key properties hold. First, every eigenvalue in the set $\Lambda(n)$, which is nonempty by the recoverable structure assumption, has a corresponding eigenvalue within distance $O(b(n))$, and therefore in $\hat{\Lambda}(n)$. Second,

$$\|P_{\underline{L}(n)}^\perp P_{\hat{L}(n)}\| \leq \frac{2\|\mathbf{E}\|}{\text{gap}} \quad (14)$$

where gap is the minimum distance between some eigenvalue of $\hat{\mathbf{D}}$ in $\hat{\Lambda}(n)$ and some eigenvalue of \mathbf{D} not contained in $\underline{\Lambda}(n)$. This gap is at least $\widehat{M}(n) - \underline{M}(n) \gg b(n)$. Thus, increasing $N_1(\eta)$ if necessary, we conclude the statement of the lemma. \square

Lemma 4. As $n \rightarrow \infty$, the expenditure derivative \dot{S} converges in probability to 1.

Proof. Write $\mathbf{q}^0 = \hat{\mathbf{q}}^0 - \boldsymbol{\varepsilon}$ and calculate

$$\begin{aligned} \dot{S} = \boldsymbol{\sigma} \cdot \mathbf{q}^0 &= \frac{P_{\hat{L}(n)} \hat{\mathbf{q}}^0}{\|P_{\hat{L}(n)} \hat{\mathbf{q}}^0\|^2} \cdot (\hat{\mathbf{q}}^0 - \boldsymbol{\varepsilon}) \\ &= 1 - \frac{P_{\hat{L}(n)} \hat{\mathbf{q}}^0}{\|P_{\hat{L}(n)} \hat{\mathbf{q}}^0\|^2} \cdot P_{\hat{L}(n)} \boldsymbol{\varepsilon} \end{aligned}$$

By the Cauchy–Schwarz inequality

$$|P_{\hat{L}(n)} \hat{\mathbf{q}}^0 \cdot P_{\hat{L}(n)} \boldsymbol{\varepsilon}| \leq \|P_{\hat{L}(n)} \hat{\mathbf{q}}^0\| \cdot \|P_{\hat{L}(n)} \boldsymbol{\varepsilon}\|,$$

which, combined with the above, gives

$$\dot{S} \geq 1 - \frac{\|P_{\hat{L}(n)} \boldsymbol{\varepsilon}\|}{\|P_{\hat{L}(n)} \hat{\mathbf{q}}^0\|}.$$

The fact that $L(n)$ is a subspace of $\widehat{L}(n)$ and the assumption of recoverable structure together ensure that $\|P_{\widehat{L}(n)}\widehat{\mathbf{q}}^0\| \geq \delta$.

So to obtain the conclusion, it suffices to show that $\|P_{\widehat{L}(n)}\boldsymbol{\varepsilon}\|$ tends in probability to 0. This will be established via [Assumption 5](#). To apply it, we need to show that $\dim \widehat{L}(n)/n \rightarrow_n 0$. The reason this holds is that if there were more than $n/M(n)$ eigenvalues in $\wedge(\mathbf{D}, M(n))$, then the absolute value of sum of these eigenvalues would exceed n , but the trace (and hence sum of eigenvalues) of the negative semidefinite matrix \mathbf{D} is $-n$. Since $M(n) \rightarrow \infty$, the statement is established. \square

Now, to prove the theorem, we use [Lemma 2](#) to write:

$$\dot{W} = \underbrace{\sum_{\lambda_\ell \in \underline{\Delta}(n)} (\mathbf{u}^\ell \cdot \mathbf{q}^0)(\mathbf{u}^\ell \cdot \boldsymbol{\sigma}) \frac{|\lambda_\ell|}{1 + |\lambda_\ell|}}_{\dot{W}_M} + \underbrace{\sum_{\lambda_\ell \notin \underline{\Delta}(n)} (\mathbf{u}^\ell \cdot \mathbf{q}^0)(\mathbf{u}^\ell \cdot \boldsymbol{\sigma}) \frac{|\lambda_\ell|}{1 + |\lambda_\ell|}}_{\dot{W}_R}. \quad (15)$$

where we have divided the quantity into a main (M) part and the rest (R). A similar (but simpler) decomposition applies to expenditure

$$\dot{S} = \underbrace{\sum_{\lambda_\ell \in \underline{\Delta}(n)} (\mathbf{u}^\ell \cdot \mathbf{q}^0)(\mathbf{u}^\ell \cdot \boldsymbol{\sigma})}_{\dot{S}_M} + \underbrace{\sum_{\lambda_\ell \notin \underline{\Delta}(n)} (\mathbf{u}^\ell \cdot \mathbf{q}^0)(\mathbf{u}^\ell \cdot \boldsymbol{\sigma})}_{\dot{S}_R}. \quad (16)$$

The proof can be completed with two further lemmas.

Lemma 5. As $n \rightarrow \infty$, both \dot{W}_R and \dot{S}_R converge in probability to 0.

Proof. This follows from [Lemma 3](#) and the fact that by construction, $\boldsymbol{\sigma} \in \widehat{L}(n)$. \square

Lemma 6. As $n \rightarrow \infty$,

$$\dot{W}_M - \dot{S}_M \xrightarrow{p} 0.$$

Proof. Using the expressions above, write the difference

$$|\dot{W} - \dot{S}| = \sum_{\lambda_\ell \in \underline{\Delta}(n)} \left[\frac{|\lambda_\ell|}{1 + |\lambda_\ell|} - 1 \right] (\mathbf{u}^\ell \cdot \mathbf{q}^0)(\mathbf{u}^\ell \cdot \boldsymbol{\sigma})$$

Note that for all $\ell \in \underline{\Delta}(n)$, $|\lambda_\ell| \geq \underline{M}(n)$, so $1 - \frac{1}{\underline{M}(n)} \leq \frac{|\lambda_\ell|}{1 + |\lambda_\ell|} \leq 1$. By Hölder's inequality (with $p = \infty$ and $q = 1$), we have

$$|\dot{W} - \dot{S}| \leq \frac{1}{\underline{M}(n)} \sum_{\lambda_\ell \in \underline{\Delta}(n)} |(\mathbf{u}^\ell \cdot \mathbf{q}^0)(\mathbf{u}^\ell \cdot \boldsymbol{\sigma})|.$$

Then applying Cauchy-Schwarz term by term, we have

$$|\dot{W} - \dot{S}| \leq \frac{1}{\underline{M}(n)} \sum_{\lambda_\ell \in \underline{\Delta}(n)} \|P_{\mathbf{u}^\ell} \mathbf{q}^0\| \|P_{\mathbf{u}^\ell} \boldsymbol{\sigma}\|.$$

$$|\dot{W} - \dot{S}| \leq \frac{1}{\underline{M}(n)} \|P_{\underline{L}(n)} \mathbf{q}^0\| \|P_{\underline{L}(n)} \boldsymbol{\sigma}\|.$$

Now using our [Assumption 3](#)—that $\|\mathbf{q}^0\|$ is bounded in n —gives the result. \square

We can put everything together. [Lemma 4](#) gives that $\dot{S} \xrightarrow{p} 1$. Combining this with [\(16\)](#) and [Lemma 5](#) (which says that $\dot{S}_R \xrightarrow{p} 0$) we see that $\dot{S}_M \xrightarrow{p} 1$. Then using [Lemma 6](#), we find that $\dot{W}_M \xrightarrow{p} 1$. Another application of [Lemma 5](#) gives that $\dot{W}_R \xrightarrow{p} 0$, so that $\dot{W} \xrightarrow{p} 1$. The claim about the effect on \dot{C} follows immediately from [Proposition 1](#).

Finally, we consider the effect on individual consumer surpluses. For any consumer h ,

$$\dot{C}^h = -\mathbf{q}^h \cdot \dot{\mathbf{p}} = -\sum_{\ell=1}^n (\mathbf{u}^\ell \cdot \mathbf{q}^h)(\mathbf{u}^\ell \cdot \dot{\mathbf{p}}) \quad (17)$$

Using [Lemma 1](#):

$$\mathbf{u}^\ell \cdot \dot{\mathbf{p}} = \frac{1}{1 + |\lambda_\ell|} (\mathbf{u}^\ell \cdot \boldsymbol{\sigma}) \quad (18)$$

Then by an argument very similar to the proof of [Lemma 5](#), we conclude that $\dot{C}^h \xrightarrow{p} 0$.

A.3. Proof of [Proposition 2](#). We prove each part separately.

Part 1: We construct an environment (i.e, a set of market states Θ and, for each $\theta \in \Theta$ a distribution μ_θ of (\hat{D}, \hat{q}^0)) where no intervention robustly increases total surplus, but where an omniscient authority could achieve total surplus gains $\dot{P} + \dot{C} = 2\dot{S}$ with $\dot{C} \geq 0$.

Let

$$D^* = -\frac{n}{n-1} \mathbf{I} + \frac{1}{n-1} \mathbf{1}\mathbf{1}^\top$$

where \mathbf{I} is the $n \times n$ identity matrix and $\mathbf{1}$ is the n -dimensional vector of ones. Let $Q \subset \mathbb{R}^n$ be the $(n-1)$ -dimensional sphere of vectors orthogonal to $\mathbf{1}$ with norm 1. The set of market states is

$$\Theta = \left\{ (D, \mathbf{q}^0) : D = D^* \text{ and } \mathbf{q}^0 = \frac{1}{n} \left(\mathbf{1} + \frac{1}{2} \mathbf{r} \right), \text{ for some } \mathbf{r} \in Q. \right\}$$

In this proof it will be useful to define random variables on the extended reals $\mathbb{R} \cup \{-\infty, \infty\}$, which allows us to take $\varepsilon_i = \infty$ for all i .³⁰ Note that D is known and the signal about \mathbf{q}^0 is always ∞ in every entry, so that any intervention rule implements a single intervention—call it $\boldsymbol{\sigma}^*$.

Note that D is a normalized Slutsky matrix. The spectral decomposition is as follows:

- $\lambda_1 = 0$ with multiplicity 1; the corresponding eigenvector is $\mathbf{u}^1 = n^{-1/2} \mathbf{1}$;
- $\lambda_2 = -\frac{n}{n-1}$ with multiplicity $n-1$.

Now write

$$\boldsymbol{\sigma} = a_1 \mathbf{u}^1 + a_2 \mathbf{u}^2, \quad (19)$$

where \mathbf{u}^2 is a vector orthogonal to \mathbf{u}^1 and $a_2 \leq 0$ (we can achieve this sign of a_2 by choosing the sign of the vector \mathbf{u}^1 appropriately).

³⁰In [Assumption 4](#), we may take $\bar{V} = \infty$; this is immaterial since that assumption plays a role only in our positive results. Since D is deterministic, we can take E to be the zero matrix.

Complete $(\mathbf{u}^1, \mathbf{u}^2)$ to an orthonormal basis of eigenvectors of \mathbf{D} , $(\mathbf{u}^1, \mathbf{u}^2, \mathbf{u}^3, \dots, \mathbf{u}^n)$. Consider the market state $(\mathbf{D}, \mathbf{q}^0) \in \Theta$ given by $\mathbf{D} = \mathbf{D}^*$ and

$$\mathbf{q}^0 = \frac{1}{n} \left(\mathbf{1} + \frac{1}{2} \mathbf{u}^2 \right).$$

For any intervention σ , Lemma 2 states that the change in total surplus is given by:

$$\dot{W} = - \sum_{\ell=1}^n \frac{|\lambda_\ell|}{1 + |\lambda_\ell|} (\mathbf{u}^\ell \cdot \mathbf{q}^0) (\mathbf{u}^\ell \cdot \sigma) \quad (20)$$

It follows by our choice of \mathbf{q}^0 that $\mathbf{u}^\ell \cdot \sigma$ is equal to a_1 and a_2 for $\ell = 1$ and $\ell = 2$ respectively, and to 0 for $\ell > 2$. Also, $\mathbf{u}^\ell \cdot \mathbf{q}^0$ is equal to $n^{-1/2}$ for $\ell = 1$ and to $n^{-1}/2$ for $\ell = 2$. Since $\lambda_1 = 0$, only the term corresponding to $\ell = 2$ is nonzero, and so

$$\dot{W} = - \frac{|\lambda_2|}{1 + |\lambda_2|} \frac{1}{2n} a_2,$$

which is non-positive by our choice of $a_2 \leq 0$. Therefore, in this market state σ achieves $\dot{W} > 0$ with probability 0. We have thus shown that no intervention can achieve $\dot{W} > 0$ ϵ -robustly, for any $\epsilon > 0$.

On the other hand, by Proposition 3 in Appendix B, any surplus outcome $(\dot{P}, \dot{C}, \dot{S})$ satisfying $\frac{1}{2}\dot{P} + \dot{C} = \dot{S}$ can be achieved.

Part 2: Note the setting here is different from Part 1: in that case we were constructing an environment without aggregate structure, while here we will give an example of an environment with aggregate structure where no intervention can robustly increase consumer surplus.

The set of market states is as follows. Since we have to show that there is no intervention that has the claimed property for all n , we will work with $n = 2^m$ firms, where m is a positive integer. Let $\mathbf{r}^1, \mathbf{r}^2, \dots, \mathbf{r}^n \in \{-1, +1\}^n$ be vectors that are mutually orthogonal, with \mathbf{r}^1 being the vector of all ones. (This is possible by Sylvester's construction of Hadamard matrices.) Let $\mathbf{u}^\ell = \frac{1}{\sqrt{n}} \mathbf{r}^\ell$. Let

$$\mathbf{D}^* = -\frac{n}{2} (\mathbf{u}^1)^\top \mathbf{u}^1 - \frac{1}{2} \mathbf{I}.$$

Note this matrix has eigenvalues $\lambda_1 = -\frac{n+1}{2}$ with multiplicity 1 and $-1/2$ with multiplicity $n-1$. The set of possible market states is $(\mathbf{D}^*, \mathbf{q}^0)$, where \mathbf{q}^0 is given by any vector of the form

$$\mathbf{q}^0 = \frac{1}{2} \left(\mathbf{u}^1 + f(n) \sum_{\ell=2}^n s_\ell \mathbf{u}^\ell \right), \quad (21)$$

where each $s_\ell \in \{-1, +1\}$, and $f(n)$ is a real-valued function we will specify. Each $(\mathbf{D}^*, \mathbf{q}^0) \in \Theta$ satisfies our maintained assumptions. Moreover, Θ has $(M(n), \delta)$ -recoverable structure for any $M(n) \leq n/2$ and $\delta = .49$.

In any environment, suppose there is an intervention rule that achieves the property $\dot{C} \geq \epsilon > 0$ robustly over a set of market states Θ . Then, by definition of robustly achieving the property, the intervention rule must achieve it with high probability when θ is drawn from any distribution over Θ . We will use this fact now, specifying the distribution over Θ given by drawing the s_ℓ in (21)

with equal probability from ± 1 , and derive a contradiction to our assumption at the start of this paragraph.

Recall

$$\dot{C} = - \sum_{\ell=1}^n \frac{1}{1 + |\lambda_\ell|} (\mathbf{u}^\ell \cdot \mathbf{q}^0) (\mathbf{u}^\ell \cdot \boldsymbol{\sigma}).$$

The $\ell = 1$ converges in probability to 0 as $n \rightarrow \infty$, since $|\lambda_1| \rightarrow \infty$. Note we may take each \mathbf{u}^ℓ to be a scaling of $\mathbf{r}^{(\ell)}$.

To handle the rest of the summation, we introduce the error structure in this environment: let the ε_i be independent and normal, each with variance $1/n$, which satisfies [Assumption 4](#). We note a useful fact. Fixing any $\delta > 0$ and n , if $f(n)$ is small enough, then (under [Assumption 4](#)), for large enough n , the conditional distribution over $\mathbf{q}^0 - \frac{1}{2}\mathbf{u}^1$ given the authority's signal (in this example just $\hat{\mathbf{q}}^0$) is within δ of the prior in total variation norm, except for signals having probability at most δ .³¹

Now, let

$$\dot{C}_{\geq 2} = - \sum_{\ell=2}^n \frac{1}{1 + |\lambda_\ell|} (\mathbf{u}^\ell \cdot \mathbf{q}^0) (\mathbf{u}^\ell \cdot \boldsymbol{\sigma}) = -\frac{2}{3} \sum_{\ell=2}^n (\mathbf{u}^\ell \cdot \mathbf{q}^0) (\mathbf{u}^\ell \cdot \boldsymbol{\sigma}).$$

We will study the distribution of this random variable for any fixed $\boldsymbol{\sigma}$, under the prior. We claim that $\dot{C}_{\geq 2}$ has the same distribution as $-\dot{C}_{\geq 2}$. This is because in the above expression only \mathbf{q}^0 is random, and $\mathbf{u}^\ell \cdot \mathbf{q}^0$ are independently equal to plus or minus the same constant by definition. Thus, for any $\epsilon > 0$,

$$P(\dot{C}_{\geq 2} \leq -\epsilon \mid \hat{\mathbf{q}}^0) - P(\dot{C}_{\geq 2} \geq \epsilon \mid \hat{\mathbf{q}}^0) \rightarrow_n 0.$$

Now, taking expectations over $\hat{\mathbf{q}}^0$, and using the useful fact, we find that

$$P(\dot{C}_{\geq 2} \leq -\epsilon) - P(\dot{C}_{\geq 2} \geq \epsilon) \rightarrow_n 0.$$

Thus we have a contradiction to robustly achieving $\dot{C} \geq \epsilon$.

On the other hand, again by [Proposition 3](#) in [Appendix B](#), any surplus outcome $(\dot{P}, \dot{C}, \dot{S})$ satisfying $\frac{1}{2}\dot{P} + \dot{C} = \dot{S}$ can be achieved.

Remark 1. For simplicity this proof has worked with a case where all uncertainty comes from \mathbf{q}^0 , and the projections of \mathbf{q}^0 onto various low eigenvectors are small. A stronger negative result can be obtained by elaborating the construction a bit—making the λ_ℓ different from one another, so that the eigenspaces orthogonal to $\mathbf{1}$ are not deterministic. Then, by showing that the eigenvectors are often impossible to recover with any precision, we can obtain a similar conclusion even if all the projections $\mathbf{u}^\ell \cdot \mathbf{q}^0$ are bounded below.

³¹Let's work in the basis $(\mathbf{u}^1, \mathbf{u}^2, \dots, \mathbf{u}^n)$. For each ℓ , the authority observes a signal $z_\ell = f(n)s_\ell \mathbf{u}^\ell + \zeta_\ell$, where ζ_ℓ are i.i.d. normal with variance $1/n$ (we have used that the rotation into the new basis does not affect the distribution of errors). Now fix n and condition on error realizations satisfying $\|\varepsilon\| \leq K(n)$, where $K(n)$ will be specified later. For a fixed $K(n)$, if $f(n)$ is chosen small enough, the density of the signal z_ℓ depends arbitrarily little on s_ℓ , and thus the posterior distribution about s_ℓ depends arbitrarily little on the signal.

APPENDIX B. COMPLETE INFORMATION

Recall from [Proposition 1](#) that

$$\frac{1}{2}\dot{P} + \dot{C} = \dot{S} \quad (22)$$

always holds. We now show the following result.

Proposition 3. If all products are independent (i.e., $\mathbf{D} = -\mathbf{I}$), then an intervention that spends \dot{S} dollars implements $\dot{C} = \frac{\dot{S}}{2}$ and $\dot{P} = \dot{S}$. Otherwise, for generic \mathbf{q}^0 , any $(\dot{C}, \dot{P}, \dot{S})$ that satisfies (22) can be implemented by an intervention.

Proof. First, consider the case in which products are independent. In this case, the effect of a subsidy to firm i and a subsidy to firm j to surpluses are separable. So it is sufficient to prove the result for an intervention which only subsidizes one firm: take $\boldsymbol{\sigma}$ with $\sigma_j = 0$ for all $j \neq 1$ and $\sigma_1 > 0$, so that $\dot{S} = \sigma_1 q_1^0$. Note the basis diagonalizing \mathbf{D} is the standard basis, and using the formulas in [Lemma 2](#) with $\lambda_\ell = -1$ for all ℓ , we conclude $\dot{C} = \sigma q_1^0/2$ and $\dot{P} = \sigma q_1^0$, as claimed.

Second, if products are not all independent ($\mathbf{D} \neq -\mathbf{I}$), then there exists $\ell \neq \ell'$ with $\lambda_\ell \neq \lambda_{\ell'}$. For generic \mathbf{q} , we also have that $(\mathbf{u}^\ell \cdot \mathbf{q})(\mathbf{u}^{\ell'} \cdot \mathbf{q}) \neq 0$. Without loss, take $\ell = 1$ and $\ell' = 2$. We now construct a set of interventions under which we can obtain any outcome $(\dot{C}, \dot{P}, \dot{S})$ that satisfies (22).

For any real number $\beta \neq 0$ and any \dot{S} , consider the intervention

$$\boldsymbol{\sigma} = \beta ((\mathbf{u}^1 \cdot \mathbf{q})\mathbf{u}^1 + \alpha(\mathbf{u}^2 \cdot \mathbf{q})\mathbf{u}^2),$$

where α is chosen so that $\boldsymbol{\sigma} \cdot \mathbf{q} = \dot{S}$. Note that $\boldsymbol{\sigma} \cdot \mathbf{q} = \beta((\mathbf{u}^1 \cdot \mathbf{q})^2 + \alpha(\mathbf{u}^2 \cdot \mathbf{q})^2)$ so this implies that

$$\alpha = \frac{\dot{S}/\beta - (\mathbf{u}^1 \cdot \mathbf{q})^2}{(\mathbf{u}^2 \cdot \mathbf{q})^2}$$

Hence, as we vary β , we keep $\boldsymbol{\sigma} \cdot \mathbf{q}$ constant and equal to \dot{S} . Next, note that

$$\begin{aligned} \dot{C} &= \sum_{\ell} \frac{1}{1 + |\lambda_{\ell}|} (\mathbf{u}^{\ell} \cdot \mathbf{q})(\mathbf{u}^{\ell} \cdot \boldsymbol{\sigma}) \\ &= \frac{\beta}{1 + |\lambda_1|} (\mathbf{u}^1 \cdot \mathbf{q})^2 + \frac{\alpha\beta}{1 + |\lambda_2|} (\mathbf{u}^2 \cdot \mathbf{q})^2 \\ &= \frac{\beta}{1 + |\lambda_1|} (\mathbf{u}^1 \cdot \mathbf{q})^2 + \frac{\dot{S} - \beta(\mathbf{u}^1 \cdot \mathbf{q})^2}{(\mathbf{u}^2 \cdot \mathbf{q})^2} \frac{1}{1 + |\lambda_2|} (\mathbf{u}^2 \cdot \mathbf{q})^2 \\ &= \frac{\beta}{1 + |\lambda_1|} (\mathbf{u}^1 \cdot \mathbf{q})^2 + \frac{\dot{S} - \beta(\mathbf{u}^1 \cdot \mathbf{q})^2}{1 + |\lambda_2|} \\ &= \beta(\mathbf{u}^1 \cdot \mathbf{q})^2 \left(\frac{1}{1 + |\lambda_1|} - \frac{1}{1 + |\lambda_2|} \right) + \frac{1}{1 + |\lambda_2|} \dot{S} \end{aligned}$$

Since this is a non-constant linear function in β , for a given \dot{S} , we can achieve any given \dot{C} by choosing β appropriately. \square

APPENDIX C. A NOISE STRUCTURE UNDER WHICH ASSUMPTION 4 HOLDS

For a concrete interpretation of Assumption 4, we will present a sampling procedure and an associated estimator for the normalized demand matrix \mathbf{D} . As in Appendix A.1, we will not assume any normalization to begin with (since the authority does not have the luxury of the market being normalized).

Since all our notation about the Slutsky matrix (D_{ij} , its estimators \widehat{D}_{ij} , etc.) was for the normalized version of this matrix, for this proof we use underlines to refer the un-normalized matrix, whose entries are $\underline{D}_{ij} = \partial q_i / \partial p_j$. We will also use underlines in our notation for estimators of the un-normalized entries.

For simplicity we assume that all households share a single representative utility function for goods and that the number of households exceeds n^2 , where n is the number of firms. Moreover, we assume that, for all i and j , the entries \underline{D}_{ij} are uniformly bounded, i.e., $|\underline{D}_{ij}| \leq \underline{D}_{\max} < \infty$ for some positive constant \underline{D}_{\max} . Also assume the true diagonal entries \underline{D}_{ii} are bounded away from 0: $0 < \underline{D}_{\min} \leq |\underline{D}_{ii}|$ for all i . The constants \underline{D}_{\min} and \underline{D}_{\max} do not depend on i, j , or n .

For each product pair (i, j) , the authority samples a distinct household—call it $h(i, j)$ —with the representative preferences; it performs a demand experiment to obtain an unbiased estimate $\widehat{\underline{D}}_{ij}^{h(i,j)}$ of \underline{D}_{ij} and an unbiased estimate $\widehat{\underline{D}}_{ii}^{h(i,j)}$ of \underline{D}_{ii} . Note there is a distinct estimate of \underline{D}_{ii} for each household $h(i, j)$.

We drop the superscript $h(i, j)$ on $\widehat{\underline{D}}_{ij}^{h(i,j)}$ and write

$$\widehat{\underline{D}}_{ij} = \underline{D}_{ij} + F_{ij}.$$

We assume that the F_{ij} are mean-zero and independent across (i, j) , that their support is uniformly bounded, and that the matrix \mathbf{F} is symmetric.³² Let \bar{V} be an upper bound on $\text{Var}[F_{ij}]$ that does not depend on i, j , or n .

Turning to estimates of \underline{D}_{ii} , we write

$$\widehat{\underline{D}}_{ii}^{h(i,j)} = \underline{D}_{ii} + G_{ij}^{h(i,j)}.$$

We assume that the $G_{ij}^{h(i,j)}$ are mean-zero and independent across (i, j) , and that $G_{ij}^{h(i,j)}$ is independent of $F_{i'j'}$ unless $(i', j') = (i, j)$. We also assume that the variances of the $G_{ij}^{h(i,j)}$ are bounded by some number \bar{V} that does not depend on i, j , or n . Now, letting

$$\widehat{\underline{D}}_{ii} = \frac{1}{n} \sum_j \widehat{\underline{D}}_{ij}^{h(i,j)},$$

we may write

$$\widehat{\underline{D}}_{ii} = \underline{D}_{ii} + G_{ii}, \text{ where } G_{ii} = \frac{1}{n} \sum_j G_{ij}^{h(i,j)}.$$

Next, as in Appendix A.1, let $\widehat{\Gamma}$ be the diagonal matrix whose (i, i) diagonal entry is $\widehat{\Gamma}_{ii} = 1/\sqrt{\widehat{\underline{D}}_{ii}}$, and construct

$$\widehat{\mathbf{D}} = \widehat{\Gamma} \widehat{\underline{\mathbf{D}}} \widehat{\Gamma}.$$

³²Since \mathbf{D} is a symmetric matrix, we view the parameters as being the numbers on or above the diagonal.

Next, turning to true quantities (as opposed to hatted estimators) recall the analogous definitions $\Gamma_{ii} = 1/\sqrt{D_{ii}}$ and also recall that

$$\mathbf{D} = \Gamma \underline{\mathbf{D}} \Gamma$$

is the true normalized Slutsky matrix. Write $\mathbf{E} = \widehat{\mathbf{D}} - \mathbf{D}$. Notice that $\text{Var}[G_{ii}] \leq \bar{V}/(n-1)$, a fact we will use throughout. (The reason for the very precise estimate is that we are using $n-1$ different, independent households' signals to estimate the same number \underline{D}_{ii} .)

Lemma 7. Under the sampling procedure described, $\mathbb{E}[\|\mathbf{E}\|] < b(n)$, where $b(n) = \gamma n^{1/2}$ for some constant $\gamma > 0$.

Proof. Let $\Phi = \widehat{\Gamma} - \Gamma$. We start by rewriting the error matrix \mathbf{E} :

$$\mathbf{E} = \widehat{\mathbf{D}} - \mathbf{D} = \widehat{\Gamma} \widehat{\underline{\mathbf{D}}} \widehat{\Gamma} - \Gamma \underline{\mathbf{D}} \Gamma.$$

Substituting $\widehat{\Gamma} = \Gamma + \Phi$ and $\widehat{\underline{\mathbf{D}}} = \underline{\mathbf{D}} + \mathbf{F}$, we have:

$$\mathbf{E} = (\Gamma + \Phi)(\underline{\mathbf{D}} + \mathbf{F})(\Gamma + \Phi) - \Gamma \underline{\mathbf{D}} \Gamma.$$

Expanding the product and using the symmetry of all the matrices involved, we obtain:

$$\begin{aligned} \mathbf{E} &= \Gamma \underline{\mathbf{D}} \Gamma + \Gamma \underline{\mathbf{D}} \Phi + \Gamma \mathbf{F} \Gamma + \Gamma \mathbf{F} \Phi \\ &\quad + \Phi \underline{\mathbf{D}} \Gamma + \Phi \underline{\mathbf{D}} \Phi + \Phi \mathbf{F} \Gamma + \Phi \mathbf{F} \Phi - \Gamma \underline{\mathbf{D}} \Gamma \\ &= \Gamma \mathbf{F} \Gamma + 2\Gamma \underline{\mathbf{D}} \Phi + 2\Gamma \mathbf{F} \Phi + \Phi \underline{\mathbf{D}} \Phi + \Phi \mathbf{F} \Phi. \end{aligned}$$

Our goal is to bound the expected spectral norm of \mathbf{E} :

$$\mathbb{E}[\|\mathbf{E}\|] = O(\sqrt{n}).$$

Using the triangle inequality for the spectral norm $\|\cdot\|$, we have:

$$\|\mathbf{E}\| \leq \|\Gamma \mathbf{F} \Gamma\| + 2\|\Gamma \underline{\mathbf{D}} \Phi\| + 2\|\Gamma \mathbf{F} \Phi\| + \|\Phi \underline{\mathbf{D}} \Phi\| + \|\Phi \mathbf{F} \Phi\|.$$

We will bound the expected value of each term on the right-hand side. In this argument, we adopt the standard practice that the meaning of the constant C can change from line to line, but this symbol always stands for a deterministic constant that does not depend on i, j , or n .

First term: $\mathbb{E}[\|\Gamma \mathbf{F} \Gamma\|]$. Since Γ is a diagonal matrix with entries $\Gamma_{ii} = 1/\sqrt{D_{ii}}$, and D_{ii} is bounded away from zero and infinity, there exist constants $\Gamma_{\min}, \Gamma_{\max} > 0$ such that:

$$\Gamma_{\min} \leq \Gamma_{ii} \leq \Gamma_{\max}.$$

Thus, $\|\Gamma\| \leq \Gamma_{\max}$. The matrix \mathbf{F} has independent, mean-zero entries F_{ij} with variances $\text{Var}(F_{ij}) \leq \bar{V}$.

Consider the matrix $\mathbf{H} = \Gamma \mathbf{F} \Gamma$ with entries $H_{ij} = \Gamma_{ii} F_{ij} \Gamma_{jj}$. The variances of H_{ij} satisfy:

$$\text{Var}(H_{ij}) = \Gamma_{ii}^2 \Gamma_{jj}^2 \text{Var}(F_{ij}) \leq \Gamma_{\max}^4 \bar{V}.$$

By standard results on the expected spectral norm of symmetric random matrices with independent entries (Bai and Yin, 1988), we have that $\|\mathbf{H}\| = O_p(\sqrt{n})$. Under our assumption that \mathbf{F} has uniformly bounded support, this can be extended (Dallaporta, 2012, Theorem 2) to the statement that $\mathbb{E}[\|\mathbf{H}\|] = O(\sqrt{n})$.

Second term: $\mathbb{E}[\|\Gamma \underline{D} \Phi\|]$. Since $\Phi = \widehat{\Gamma} - \Gamma$ and $\widehat{\Gamma}_{ii} = 1/\sqrt{\widehat{D}_{ii}}$, we use a Taylor expansion around D_{ii} to approximate³³ Φ_{ii} :

$$\Phi_{ii} = \widehat{\Gamma}_{ii} - \Gamma_{ii} = -\frac{1}{2}\Gamma_{ii}^3 G_{ii} + O(G_{ii}^2).$$

The variance of Φ_{ii} satisfies:

$$\text{Var}(\Phi_{ii}) \leq \left(\frac{1}{2}\Gamma_{\max}^3\right)^2 \text{Var}(G_{ii}) \leq \frac{C}{n} \quad (23)$$

for some constant $C > 0$. The entries of $\Gamma \underline{D} \Phi$ are:

$$(\Gamma \underline{D} \Phi)_{ij} = \Gamma_{ii} \underline{D}_{ij} \Phi_{jj}.$$

Therefore, the squared Frobenius norm is:

$$\|\Gamma \underline{D} \Phi\|_F^2 = \sum_{i,j} (\Gamma_{ii} \underline{D}_{ij} \Phi_{jj})^2.$$

Taking expectations:

$$\mathbb{E}[\|\Gamma \underline{D} \Phi\|_F^2] \leq \Gamma_{\max}^2 (\underline{D}_{\max})^2 \sum_{i,j} \mathbb{E}[\Phi_{jj}^2] \leq Cn^2 \cdot \frac{1}{n} = Cn,$$

where in the penultimate step we have used eq. (23). Thus, using the fact that Frobenius norm bounds the spectral norm and Jensen's inequality, we have:

$$\mathbb{E}[\|\Gamma \underline{D} \Phi\|] \leq \mathbb{E}[\|\Gamma \underline{D} \Phi\|_F] \leq \sqrt{\mathbb{E}[\|\Gamma \underline{D} \Phi\|_F^2]} \leq C\sqrt{n}.$$

Third term: $\mathbb{E}[\|\Gamma \mathbf{F} \Phi\|]$. The entries of $\Gamma \mathbf{F} \Phi$ are:

$$(\Gamma \mathbf{F} \Phi)_{ij} = \Gamma_{ii} F_{ij} \Phi_{jj}.$$

The squared Frobenius norm is:

$$\|\Gamma \mathbf{F} \Phi\|_F^2 = \sum_{i,j} (\Gamma_{ii} F_{ij} \Phi_{jj})^2.$$

Taking expectations and using the Cauchy–Schwarz inequality:

$$\mathbb{E}[\|\Gamma \mathbf{F} \Phi\|_F^2] \leq \Gamma_{\max}^2 \bar{V} \sum_{i,j} \mathbb{E}[\Phi_{jj}^2] \leq Cn^2 \cdot \frac{1}{n} = Cn.$$

Therefore:

$$\mathbb{E}[\|\Gamma \mathbf{F} \Phi\|] \leq C\sqrt{n}.$$

Fourth term: $\mathbb{E}[\|\Phi \underline{D} \Phi\|]$ The entries of $\Phi \underline{D} \Phi$ are:

$$(\Phi \underline{D} \Phi)_{ij} = \Phi_{ii} \underline{D}_{ij} \Phi_{jj}.$$

The squared Frobenius norm is:

$$\|\Phi \underline{D} \Phi\|_F^2 = \sum_{i,j} (\Phi_{ii} \underline{D}_{ij} \Phi_{jj})^2.$$

³³Let $f(x) = 1/\sqrt{x}$. The derivative is $f'(x) = -\frac{1}{2}x^{-3/2}$. Expanding $f(D_{ii} + G_{ii})$ around $x = D_{ii}$ using Taylor's theorem, we obtain $f(D_{ii} + G_{ii}) = f(D_{ii}) + f'(D_{ii})G_{ii} + O(G_{ii}^2)$. Therefore, $\Phi_{ii} = \widehat{\Gamma}_{ii} - \Gamma_{ii} = f(D_{ii} + G_{ii}) - f(D_{ii}) = f'(D_{ii})G_{ii} + O(G_{ii}^2) = -\frac{1}{2}\Gamma_{ii}^3 G_{ii} + O(G_{ii}^2)$, where $\Gamma_{ii} = f(D_{ii})$.

Taking expectations:

$$\mathbb{E}[\|\Phi \underline{D} \Phi\|_F^2] \leq (\underline{D}_{\max})^2 \sum_{i,j} \mathbb{E}[\Phi_{ii}^2] \mathbb{E}[\Phi_{jj}^2] \leq Cn^2 \left(\frac{C}{n}\right)^2 = C.$$

Thus:

$$\mathbb{E}[\|\Phi \underline{D} \Phi\|] \leq C.$$

Fifth term: $\mathbb{E}[\|\Phi \mathbf{F} \Phi\|]$ The entries of $\Phi \mathbf{F} \Phi$ are:

$$(\Phi \mathbf{F} \Phi)_{ij} = \Phi_{ii} F_{ij} \Phi_{jj}.$$

The squared Frobenius norm is:

$$\|\Phi \mathbf{F} \Phi\|_F^2 = \sum_{i,j} (\Phi_{ii} F_{ij} \Phi_{jj})^2.$$

Taking expectations and using the Cauchy–Schwarz inequality,

$$\mathbb{E}[\|\Phi \mathbf{F} \Phi\|_F^2] \leq \bar{V} \sum_{i,j} \mathbb{E}[\Phi_{ii}^2] \mathbb{E}[\Phi_{jj}^2] \leq Cn^2 \left(\frac{C}{n}\right)^2 = \frac{C}{n}.$$

Therefore:

$$\mathbb{E}[\|\Phi \mathbf{F} \Phi\|] \leq \frac{C}{\sqrt{n}}.$$

Combining all the bounds:

$$\mathbb{E}[\|\mathbf{E}\|] \leq C_1 \sqrt{n} + 2C_2 \sqrt{n} + 2C_3 \sqrt{n} + C_4 + \frac{C_5}{\sqrt{n}} = O(\sqrt{n}),$$

where C_1, C_2, C_3, C_4, C_5 are constants independent of n , as desired. \square