ROBUST MARKET INTERVENTIONS

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

ABSTRACT. When can interventions in markets be designed to increase surplus *robustly*—i.e., with high probability—accounting for uncertainty due to imprecise information about economic primitives? In a setting with many strategic firms, each possessing some market power, we present conditions for such interventions to exist. The key condition, *significant structure*, requires large-scale complementarities among families of products. The analysis works by decomposing the incidence of interventions in terms of principal components of a Slutsky matrix. Under significant structure, a noisy signal of this matrix reveals enough about these principal components to design robust interventions. Our results demonstrate the usefulness of spectral methods for analyzing imperfectly observed strategic interactions with many agents.

Date Printed. August 22, 2026.

Galeotti is at the London Business School, Golub is at Northwestern University, Goyal is at the University of Cambridge, Talamàs is at IESE Business School, and Tamuz is at Caltech.

We thank Yann Calvó 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 Keynes Fund for Applied Economics (Goyal), 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 for sharing data with us, and Volker Nocke for a helpful discussion on a related project at the 2023 AEA meetings. We thank Ed Glaeser for an important early nudge to pursue this research direction, and we thank Pol Antràs, Itai Ashlagi, Modibo Camara, Yeon-Koo Che, Kevin Chen, Rebecca Diamond, Jeff Ely, Joey Feffer, Alex Frankel, Matthew Gentzkow, Tim Gowers, Matthew O. Jackson, Zi Yang Kang, Ernest Liu, Margaret Meyer, Jose Luis Moraga, Stephen Morris, Martin Peitz, Ilya Segal, Andy Skrzypacz, Ran Spiegler, Ludwig Straub, Xavier Vives, Kuang Xu, and Frank Yang for helpful comments.

We thank the editor and four excellent referees for very helpful comments and guidance on previous versions that greatly improved the manuscript.

The paper subsumes the working paper Galeotti et al. (2021).

1. INTRODUCTION

Market power has recently attracted renewed attention and is thought to have significant and growing welfare implications (see, e.g., Syverson, 2019). While many applied studies of competition confine the study of market power to tightly defined product markets, it is becoming clear that important welfare-relevant spillovers operate *across* such markets (Baqaee and Farhi, 2020; Azar and Vives, 2021; Pellegrino, 2021; Ederer and Pellegrino, 2021). Our theoretical understanding of such spillovers in environments with many firms interacting via general demand systems remains limited. This paper is about the welfare theory of such interactions.

Consider firms that set prices in a setting where demands exhibit arbitrary complementarities and substitutabilities. For instance, one firm's product—e.g., a Samsung smartphone—may be a substitute for some products—e.g., Apple smartphones—and a complement to others—e.g., compatible accessories such as earbuds, watches, and smart home appliances—which may, in turn, be substitutes or complements to one another.

We are interested in the nature of inefficiencies in such an environment and interventions to respond to them. We consider these issues from the perspective of an authority that recognizes the possibility of inefficiency due to market power and can intervene through taxes and subsidies on firms' sales. For concreteness, we can think of this authority as the operator of a large marketplace, such as Amazon, that mediates retail sales, or as a governmental institution. The authority's objective is to increase the equilibrium economic surplus generated by the marketplace. (Once this is achieved, the authority can reclaim and redistribute this additional surplus via transfers.) What principles should guide the design of such interventions? When can such policies be implemented under realistic uncertainty about market parameters?

What makes the problem challenging is that, once we broaden our perspective beyond one traditionally defined market (e.g., a set of imperfectly substitutable products such as smartphones) and consider spillovers to a variety of other complements and substitutes, there is a kind of curse of dimensionality. A marketplace with numerous and changing goods is described by a demand system that is high-dimensional and a priori unstructured. As any firm's cost changes, the number of potential spillover effects to consider is equal to the number of products; thus, the number of interactions scales quadratically in this number. Realistic signals leave substantial uncertainty about many aspects of the environment, so that the authority will have nothing close to a precise estimate of the entire demand system. This raises the question of whether there exist policies that can reliably improve surplus despite this uncertainty.

Our main result is that if the economic environment and the statistical signal about it jointly satisfy a property that we call *significant structure*, then there are feasible intervention rules that robustly (i.e., with high probability) increase equilibrium total surplus. They do so despite the fact that the authority faces large errors in observing every detail of the system. Moreover, within a natural class of interventions—those whose direct (pre-transfer) incidence on consumer surplus is nonnegative—our feasible interventions achieve the largest gain in surplus that is possible for a given level of subsidy expenditure. Hence, within

this class, these interventions are as good as those that could be designed by an authority with perfect information.

The key condition in the paper is significant structure. We now explain what it means and how it is used.

The demand structure is encoded in a matrix $\boldsymbol{D}$ of demand derivatives (which in our setting is equal to the Slutsky matrix). A given cell D_{ij} in this matrix is the derivative of product i's demand with respect to product j's price. Thus, the matrix specifies the complementarity and substitutability relationships across products. Mathematically, the significant structure property requires that a normalized version of $\boldsymbol{D}$ possesses a low-dimensional subspace of eigenvectors with large eigenvalues. In other words, there exist one or more principal components of the demand system that account for a significant amount of demand behavior.

We show that significant structure entails an economic property—substantial large-scale complementarities in the economy. These large-scale complementarities give rise to cross-market double marginalization problems, where many goods exert externalities on one another. In turn, these externalities create the potential for small subsidies to have large spillover effects, raising the consumption of many underproduced goods and thereby substantially improving welfare.

Significant structure is also central to realizing this potential. In particular, it permits uncovering the large-scale structure of demand relationships when the Slutsky matrix is observed imperfectly. The authority's signal consists of noisy estimates of the entries of this matrix. In principle, this noise can create large uncertainty in the operation of a given intervention. The key step in dealing with this is an application of a theorem of Davis and Kahan (1970). We use this theorem to show that when the economy satisfies significant structure, the noisy observation of $\boldsymbol{D}$ suffices to precisely predict the effects of *some* well-chosen interventions—specifically, those operating in the space of eigenvectors associated with the largest eigenvalues of $\boldsymbol{D}$.

Combining the economic and statistical implications of significant structure allows us to establish our main result. In markets possessing such structure, the authority can recover precise information about large-eigenvalue components of $\boldsymbol{D}$, and this information suffices to construct interventions with highly predictable surplus implications.

The robust interventions constructed in this way keep consumer surplus essentially unchanged while generating total-surplus gains. Though the direct incidence of these gains is on producers, they can be shared with consumers through additional transfer instruments. Proposition 3 shows that, in some environments with significant structure, the authority *cannot* robustly guarantee any fixed positive fraction of expenditure accruing to consumer surplus. Thus, there are some intrinsic constraints on the direct incidence of robust interventions.

At a technical level, to perform our analysis we develop 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 the implications of any intervention as a linear combination of orthogonal effects. These effects correspond to the projection of the intervention onto each

eigenvector of a (normalized) Slutsky matrix $\boldsymbol{D}$. By characterizing the pass-throughs of subsidies to prices, quantities, and welfare separately across these principal components, we are able to prove that targeting the high-eigenvalue principal components yields precisely predictable results achieving our claimed welfare properties.

1.1. **Structure of the paper.** We introduce the framework in Section 2 and, as a benchmark, Section 3 characterizes the outcome an authority can implement with complete information (Proposition 1).

Section 4 introduces an important running example of a demand system, consistent with a representative consumer with a quadratic utility function. There are K product categories; products within a category are substitutes, while across categories they are complements. We illustrate some of the main ideas of the paper in the context of this example and present a special case of our main result, Proposition 2, in this setting: as the number of product categories grows large, the authority can design robust interventions that are welfare-improving. This example brings out the role of large-scale complementarities and a tractable decomposition of effects of taxes/subsidies in the design of robust intervention policies.

Section 5.1 introduces a new spectral description of the pass-through of an intervention. The spectral decomposition is of independent interest, yielding a useful basis in which price and welfare pass-throughs of cost shocks behave intuitively despite arbitrary spillovers across firms. Section 5.2 formally defines the property of significant structure and shows how to check it within the product-category model.

Section 6 presents the main result of the paper, Theorem 1.

Section 7 develops a Monte Carlo experiment to demonstrate how significant structure permits the recovery of welfare-relevant structure from noisy observations of the demand system, and to illustrate the limitations of this recovery.

Section 8 further explores the scope of our analysis. Section 8.1 clarifies the limits of our robust interventions and shows that it may be impossible to robustly improve welfare when significant structure fails (Proposition 3). Section 8.2 discusses foundations for the property of significant structure and develops an empirical diagnostic for it, which we illustrate using the Monte Carlo experiment. Section 8.3 discusses how the linear–quadratic utility model and the hedonic model of demand relate to our theory. Section 8.4 discusses the implications of our assumption of local linear demand and how to generalize our insights once we dispense with that assumption.

1.2. **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 differentiated oligopoly network models—similar to the one we consider here—to provide empirical estimates of efficiency losses due to market power (e.g., Pellegrino, 2021; Ederer and Pellegrino, 2021).[1]

[1]See 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.[2] This approach has significant implications for understanding which kinds of empirical models are needed to design interventions in large markets with many goods. In particular, we note that the significant structure condition entails widespread complementarities in the economy, and this is distinctive relative to the focus on substitutabilities in the models of hedonic utility used in the literature. We elaborate on these issues in Section 8.

Methods in high-dimensional statistics are currently attracting considerable interest in econometric settings (Athey, 2018; Athey, Bayati, Doudchenko, Imbens, and Khosravi, 2021; Chernozhukov, Hansen, Liao, and Zhu, 2023), including work applying related statistical models to informational or behavioral spillovers in social networks and marketplaces (Golub and Jackson, 2012; Bajari, Burdick, Imbens, Masoero, McQueen, Richardson, and Rosen, 2023; Cai and Daskalakis, 2022; Dasaratha, 2020; Cai, 2022; Parise and Ozdaglar, 2023; Chandrasekhar, Goldsmith-Pinkham, McCormick, Thau, and Wei, 2024; Wager and Xu, 2021).

However, we know little about when noisy data can be effectively used in order to implement desirable interventions in the presence of strategic spillovers, particularly in market settings. We show that, in a large oligopoly market, spectral methods for high-dimensional statistics developed in the literature on large network recovery can be useful for designing socially desirable interventions.[3]

Our paper contributes to the theory of network interventions. Early contributions include Borgatti (2006), Ballester, Calvó-Armengol, and Zenou (2006), and Goyal (1996).[4] Spectral methods have recently been applied to optimal intervention problems when spillovers are known (Galeotti, Golub, and Goyal, 2020; Gaitonde, Kleinberg, and Tardos, 2021; Liu and Tsyvinski, 2024).[5] By contrast, in the present paper, the authority observes strategic spillovers with significant noise. The methods we develop for robust interventions can be applied

[2]Our 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).

[3]See the monograph Chen et al. (2021) for a detailed treatment of the mathematical background and some marquee applications.

[4]The 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), Belhaj and Deroian (2017), Demange (2017), Dziubinski and Goyal (2017), Galeotti and Goyal (2009), and Leduc, Jackson, and Johari (2017).

[5]Some recent work uses spectral analysis to derive conditions for core-selecting reallocative auctions (Rostek and Yoder, 2023), and robust implementation (Ollár and Penta, 2023). See also Aguiar and Serrano (2017) on spectral methods to study Slutsky matrices in a consumer theory setting.

to other network games and we briefly discuss this in Section 8.5. Our analysis of perturbations of taxes and subsidies is related to the classic "tax reform approach" in public finance (Feldstein, 1976; Tirole and Guesnerie, 1981); the study of uncertain spillovers distinguishes our work.

Our approach to robustness is conceptually related to, but methodologically distinct from, an extensive literature in economic theory. That literature focuses on understanding the design of mechanisms and contracts that achieve desired outcomes even when assumptions about the environment (e.g., agents' preferences, beliefs, and rationality) are relaxed; see Carroll (2019) for a survey. Our definition of robustness aligns with the spirit of this literature. However, in our context, the motivation for analyzing robust interventions arises from the high-dimensional nature of the market state, and we use methods that align with statistical work in this type of setting.

2. Framework

The foundation of our framework is a differentiated oligopoly game. Within this game, we introduce a statistical framework describing the signals available to the authority about the oligopoly, a class of interventions available to the authority, and a notion of rules that use these signals to achieve good outcomes robustly.

2.1. **Demand side.**

There is a set $N = \{1, \ldots, n\}$ of distinct products, with the number of products $n \geq 1$; a typical product is denoted by i. The demand for these products arises from the consumption choices of a finite number of optimizing households. Each household $h \in \{1, \ldots, H\}$ takes prices as given and has a *choice utility* that is quasilinear in a numeraire m,

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

where V^h is a twice-differentiable and strictly concave function of the consumption profile $\tilde{\boldsymbol{q}}^h \in \mathbb{R}^n$ and m is a numeraire ("money"), in which all prices are denominated. Given a price profile $\tilde{\boldsymbol{p}}$, the household's problem is to choose a bundle $\tilde{\boldsymbol{q}}^h$ to maximize $V^h(\tilde{\boldsymbol{q}}^h) - \tilde{\boldsymbol{p}} \cdot \tilde{\boldsymbol{q}}^h$. Letting $\boldsymbol{q}^h(\tilde{\boldsymbol{p}})$ be the solution to household h's problem (unique by strict concavity of V^h), total market demand is[6]

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

We use tilde notation for an arbitrary price or quantity, and then drop the tilde for these variables to indicate optimal or equilibrium solutions to be introduced later.

2.2. **Supply side.**

There is a firm associated with each product: Firm i produces product i. Firms play a simultaneous pricing game; each firm chooses $\tilde{p}_i \geq 0$. For any profile of prices $\tilde{\boldsymbol{p}}$, firm i's profit is

$$q_i(\tilde{\boldsymbol{p}})(\tilde{p}_i - c_i + \sigma_i), \tag{1}$$

where c_i is firm i's (constant) marginal cost of production and σ_i is a per-unit subsidy on firm i. In particular, when σ_i is positive, the authority subsidizes

[6]Because the households' utilities are quasilinear in money, one can derive the same aggregate demand from a single representative consumer.

firm i's production and transfers $\sigma_i \tilde{q}_i$ to firm i, whereas a negative σ_i corresponds to a tax.

We call $\boldsymbol{\sigma} = (\sigma_1, \ldots, \sigma_n)$ an intervention. We denote by $\boldsymbol{p}(\boldsymbol{\sigma})$ a pure strategy Nash equilibrium given $\boldsymbol{\sigma}$, and we refer to $\boldsymbol{p}^0 = \boldsymbol{p}(\mathbf{0})$ and $\boldsymbol{q}^0 = \boldsymbol{q}(\boldsymbol{p}^0)$ as the *status quo* vectors of equilibrium price and quantity. Throughout, $\|\cdot\|$ denotes the Euclidean norm for vectors and the induced operator norm for matrices. To facilitate unambiguous local comparative statics, we assume that $\boldsymbol{p}(\boldsymbol{\sigma})$ is unique and continuous in $\boldsymbol{\sigma}$ around the status quo $\boldsymbol{\sigma} = \mathbf{0}$, as the following assumption, maintained throughout the analysis, formalizes:

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

For a discussion of conditions for existence and uniqueness of Nash equilibrium in oligopoly models, see Fudenberg and Tirole (1991), Vives (1999), and Cumbul and Virág (2018).[7] A sufficient condition for local uniqueness is nonsingularity of the Jacobian of best responses at equilibrium; this holds generically in our setting.[8]

We confine attention to small interventions (i.e., $\|\boldsymbol{\sigma}\| < \nu$) and evaluate the effect of an intervention $\boldsymbol{\sigma}$ on any outcome variable Y (e.g., producer or consumer surplus) by its first derivative evaluated at the status quo intervention $\boldsymbol{\sigma} = \mathbf{0}$:

$$\dot{Y}_{\boldsymbol{\sigma}} = \sum_i \frac{dY}{d\sigma_i}(\mathbf{0})\sigma_i. \tag{2}$$

Throughout, we consider the effects of an intervention starting from the status quo equilibrium, which implies that the expenditure $S = \boldsymbol{q} \cdot \boldsymbol{\sigma}$ associated with intervention $\boldsymbol{\sigma}$ satisfies

$$\dot{S}_{\boldsymbol{\sigma}} = \boldsymbol{q}^0 \cdot \boldsymbol{\sigma}. \tag{3}$$

2.3. **Interventions and their effects.** We start by focusing on environments in which each good's demand is linear in its own price in a neighborhood of the status quo equilibrium, as captured by the following assumption.

Assumption 2 (Each good's demand is linear in its own price, locally)**.** There exists $\rho > 0$ such that $\frac{\partial q_i(\boldsymbol{p})}{\partial p_i} = \frac{\partial q_i(\boldsymbol{p}^0)}{\partial p_i}$ for all $\boldsymbol{p}$ satisfying $\|\boldsymbol{p} - \boldsymbol{p}^0\| < \rho$.

This assumption facilitates our analysis of interventions, yielding simple formulas for comparative statics. We discuss the content of this assumption and the extension of our methods to the case in which this assumption does not hold in Section 8.4.

The firms' first-order conditions imply that equilibrium prices $\boldsymbol{p}(\boldsymbol{\sigma})$ satisfy

$$q_i(\boldsymbol{p}(\boldsymbol{\sigma})) = -\frac{\partial q_i(\boldsymbol{p}(\boldsymbol{\sigma}))}{\partial p_i}(p_i(\boldsymbol{\sigma}) - c_i + \sigma_i) \quad \text{for each firm } i. \tag{4}$$

Assumption 2 implies that we can replace the $\boldsymbol{p}(\boldsymbol{\sigma})$-dependent partial derivative in the above equation with the constant $\frac{\partial q_i(\boldsymbol{p}^0)}{\partial p_i}$. Implicitly differentiating

[7] Assuming compact action sets, existence of a pure strategy equilibrium is guaranteed if profits are quasi-concave in prices.

[8] For a discussion of these issues, see, e.g., McLennan (2018).

the first-order-condition system with respect to $\boldsymbol{\sigma}$ gives the price equation below. The quantity equation follows from applying the chain rule to $\boldsymbol{q}(\boldsymbol{p}(\boldsymbol{\sigma}))$ at $\boldsymbol{\sigma} = \mathbf{0}$:

$$[-\text{diag}(\boldsymbol{D}) - \boldsymbol{D}]\dot{\boldsymbol{p}}_{\boldsymbol{\sigma}} = \text{diag}(\boldsymbol{D})\boldsymbol{\sigma} \quad \text{and} \quad \dot{\boldsymbol{q}}_{\boldsymbol{\sigma}} = \boldsymbol{D}\dot{\boldsymbol{p}}_{\boldsymbol{\sigma}}, \tag{5}$$

where $\boldsymbol{D} = \boldsymbol{D}(\boldsymbol{p}^0)$ is the Slutsky matrix:[9]

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

and $\text{diag}(\boldsymbol{D})$ is a diagonal matrix defined by $\text{diag}(\boldsymbol{D})_{ii} = D_{ii}$. Note that $D_{ii} < 0$ by strict concavity of the consumers' utility functions. For $i \neq j$, if $D_{ij} > 0$ (resp. $D_{ij} < 0$) then, around the equilibrium, products i and j are substitutes (resp. complements).

Assumption 2 implies that comparative statics of prices and quantities are fully determined by the Slutsky matrix $\boldsymbol{D}$. We note that $\boldsymbol{D}$ satisfies the following property (Nocke and Schutz, 2017).

Property NSD. The Slutsky matrix $\boldsymbol{D}$ is negative semidefinite.

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).

2.4. **Surpluses.** 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 $W = C + P - S$, where S is the authority's expenditure. Let C^h denote the consumer surplus accruing to consumer h. Given an intervention $\boldsymbol{\sigma}$, equilibrium price $\boldsymbol{p}(\boldsymbol{\sigma})$ and equilibrium quantity profiles $\{\boldsymbol{q}^h(\boldsymbol{p}(\boldsymbol{\sigma}))\}_{h=1,\dots,H}$, these are:

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

$$P = (\boldsymbol{p}(\boldsymbol{\sigma}) - \boldsymbol{c} + \boldsymbol{\sigma}) \cdot \boldsymbol{q}(\boldsymbol{p}(\boldsymbol{\sigma})) \tag{6}$$

2.5. **The statistical framework and robust interventions.** A *market state*, denoted by $\boldsymbol{\theta}$, is a pair $(\boldsymbol{D}, \boldsymbol{q}^0)$; the set of possible market states is $\boldsymbol{\Theta}$. The authority receives a signal, denoted by $\widehat{\boldsymbol{\theta}} \in \widehat{\boldsymbol{\Theta}}$, about the market state. In our setting, $\widehat{\boldsymbol{\theta}}$ consists of signals about the Slutsky matrix and quantities:

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

The joint distribution of $(\boldsymbol{E}, \boldsymbol{\varepsilon})$ may depend on $\boldsymbol{\theta}$; we denote the resulting signal distribution by $\varphi_{\boldsymbol{\theta}} \in \Delta(\widehat{\boldsymbol{\Theta}})$. We will assume for simplicity throughout that, conditional on the market state $\boldsymbol{\theta}$, the quantity errors ε_i have expectation zero, are i.i.d., and are independent of the matrix observation noise $\boldsymbol{E}$.

[9] In general, the matrix of derivatives of Marshallian demand need not be the same as the Slutsky matrix (which works with compensated demand). However, in this demand system, the wealth effect is zero due to the fact that the goods utility and money are additively separable. Thus, the two matrices coincide (Nocke and Schutz, 2017), and so we use the term "Slutsky matrix" throughout.

An *intervention rule* $\boldsymbol{R}$ maps each signal $\widehat{\boldsymbol{\theta}} = (\widehat{\boldsymbol{D}}, \widehat{\boldsymbol{q}}^0)$ to an intervention $\boldsymbol{\sigma}$.[10] A *market outcome* is a pair $(\boldsymbol{\theta}, \boldsymbol{\sigma})$ consisting of a market state and an intervention; we refer to this pair as an outcome because it determines production, consumption, and transfers. A *property* is a measurable subset $\mathscr{P}$ of all possible market outcomes. For example, a desirable property is increasing total surplus: $\mathscr{P} = \{(\boldsymbol{\theta}, \boldsymbol{\sigma}) : \dot{W}_{\boldsymbol{\sigma}} > 0\}$. We are interested in understanding which properties $\mathscr{P}$ can be achieved with high probability in all market states.

Definition 1. An intervention rule $\boldsymbol{R}$ achieves a property $\mathscr{P}$ *ϵ–robustly* if, for every $\boldsymbol{\theta} \in \boldsymbol{\Theta}$,

$$\varphi_{\boldsymbol{\theta}}\left(\widehat{\boldsymbol{\theta}} : (\boldsymbol{\theta}, \boldsymbol{R}(\widehat{\boldsymbol{\theta}})) \in \mathscr{P}\right) \geq 1 - \epsilon.$$

Note that we require this condition to hold in every state, so the only randomness in this definition is in the signal draw.[11] This is the frequentist approach following Wald's (1950) statistical decision theory. A recent line of research in economics works in this paradigm: welfare or regret guarantees are given uniformly over structural economic parameters $\boldsymbol{\theta}$ (over which the analyst does not have a prior); signals conditional on such parameters are analyzed probabilistically, with their distribution being implied by the analyst's statistical procedure.[12] For examples of recent work following this paradigm, see Manski (2004), Kitagawa and Tetenov (2018) and Athey and Wager (2021). It is worth remarking that Definition 1 defines success in terms of the probability of a good outcome. In our setting, we will be able to deduce from this statements about expected performance conditional on $\boldsymbol{\theta}$ (which hold uniformly over $\boldsymbol{\theta}$), as is typical in the literature we have just mentioned.

In our asymptotic results, we say that an intervention rule achieves a property *robustly* if, for any $\epsilon > 0$, the rule achieves the property ϵ-robustly for all sufficiently large n.

2.6. **Asymptotic analysis in large markets.** It will be convenient to have notation for expressing that some quantity asymptotically vanishes relative to others.

An *environment* is a sequence

$$\mathcal{E} := (\boldsymbol{\Theta}(n), \varphi(n))_n,$$

where $\boldsymbol{\Theta}(n)$ is a set of market states with n firms and $\varphi(n)$ describes the distribution of noise.[13] Given $\mathcal{E}$, for a sequence of real-valued random variables $X_n(\boldsymbol{\theta})$ and a constant x, we say that $X_n \xrightarrow{p} x$ *uniformly over* $\boldsymbol{\theta} \in \boldsymbol{\Theta}(n)$ if for every $\eta > 0$,

$$\sup_{\boldsymbol{\theta} \in \boldsymbol{\Theta}(n)} \mathbb{P}_{\varphi_{\boldsymbol{\theta}}(n)}\big(|X_n(\boldsymbol{\theta}) - x| > \eta\big) \to_n 0.$$

[10] We require this map to be measurable in a suitable sense, which is clear in our application.

[11] Equivalently, the condition is the requirement that the probability of achieving $\mathscr{P}$ is at least $1 - \epsilon$ under every prior over $\boldsymbol{\Theta}$.

[12] Our Appendix D describes a sampling procedure giving rise to signals about the state in our framework.

[13] Given that the distribution of noise can depend on the market state, $\varphi(n)$ is a map $\boldsymbol{\theta} \mapsto \varphi_{\boldsymbol{\theta}}(n)$ assigning to each state $\boldsymbol{\theta} \in \boldsymbol{\Theta}(n)$ a signal distribution.

We write $X_n = o_{\mathrm{p},\boldsymbol{\Theta}}(1)$ if $X_n \xrightarrow{p} 0$ uniformly over $\boldsymbol{\theta} \in \boldsymbol{\Theta}(n)$. More generally, we write $X_n = o_{\mathrm{p},\boldsymbol{\Theta}}(f(n))$ if $X_n/f(n) \xrightarrow{p} 0$ uniformly over $\boldsymbol{\theta} \in \boldsymbol{\Theta}(n)$. Finally, we write $X_n = O_{\mathrm{p},\boldsymbol{\Theta}}(1)$ if for every $\epsilon > 0$ there exists $M < \infty$ such that

$$\sup_{\boldsymbol{\theta} \in \boldsymbol{\Theta}(n)} \mathbb{P}_{\varphi_{\boldsymbol{\theta}}(n)}(|X_n(\boldsymbol{\theta})| > M) < \epsilon$$

for all sufficiently large n. In this notation, the roman subscript p denotes convergence in probability and is not an additional parameter, while $\boldsymbol{\Theta}$ records that the convergence is uniform over the sequence of state spaces $(\boldsymbol{\Theta}(n))_n$ in the environment.

We will maintain the following technical assumption, which makes asymptotics well-behaved.

Assumption 3. The environment $\mathcal{E} := (\boldsymbol{\Theta}(n), \varphi(n))_n$ is such that, for all $(\boldsymbol{D}, \boldsymbol{q}^0) \in \boldsymbol{\Theta}(n)$, there exist universal constants $0 < d_{\min} \leq d_{\max} < \infty$ such that

$$-D_{ii} \in [d_{\min}, d_{\max}] \quad \text{for all } i.$$

Assumption 3 rules out the possibility that own-price demand effects diverge relative to one another as we progress along the sequence defining the environment.

2.7. **Normalization of the Slutsky matrix.** For a given market state $\boldsymbol{\theta} = (\boldsymbol{D}, \boldsymbol{q}^0)$, it will be convenient to describe the market outcomes with respect to a normalized Slutsky matrix, denoted by $\underline{\boldsymbol{D}}$. Specifically, the normalized Slutsky matrix is:

$$\underline{\boldsymbol{D}} = [\mathrm{diag}(-\boldsymbol{D})]^{-1/2} \boldsymbol{D} [\mathrm{diag}(-\boldsymbol{D})]^{-1/2}.$$

where $\mathrm{diag}(-\boldsymbol{D})$ is a diagonal matrix retaining only the diagonal elements of $-\boldsymbol{D}$ and setting all off-diagonal elements to zero. This means that

$$\underline{D}_{ii} = -1 \quad \text{and} \quad \underline{D}_{ij} = \frac{D_{ij}}{\sqrt{D_{ii} D_{jj}}} \quad \text{for all } i \neq j.$$

This normalization entails a change of units for quantities from $\boldsymbol{q}^0$ to

$$\underline{\boldsymbol{q}}^0 = [\mathrm{diag}(-\boldsymbol{D})]^{-1/2} \boldsymbol{q}^0.$$

Keeping units of money fixed, the corresponding transformations for prices, costs, and interventions are:

$$\underline{\boldsymbol{p}}^0 = [\mathrm{diag}(-\boldsymbol{D})]^{1/2} \boldsymbol{p}^0, \quad \underline{\boldsymbol{c}}^0 = [\mathrm{diag}(-\boldsymbol{D})]^{1/2} \boldsymbol{c}^0, \quad \underline{\boldsymbol{\sigma}} = [\mathrm{diag}(-\boldsymbol{D})]^{1/2} \boldsymbol{\sigma}.$$

This normalization amounts to choosing units for the goods so that each good's own-price effect is -1, making the diagonal entries of $\underline{\boldsymbol{D}}$ homogeneous.[14] The normalized coordinates are convenient because, in these, the equilibrium first-order conditions of equation (4) become:

$$\underline{q}_i(\underline{\boldsymbol{p}}(\underline{\boldsymbol{\sigma}})) = \underline{p}_i(\underline{\boldsymbol{\sigma}}) - \underline{c}_i^0 + \underline{\sigma}_i, \tag{7}$$

which we exploit throughout.

[14] Every underlined variable (e.g., $\underline{\boldsymbol{q}}^0, \underline{\boldsymbol{p}}^0$) is defined relative to the Slutsky matrix $\boldsymbol{D}$, which depends on the state $\boldsymbol{\theta}$. A fully explicit notation would subscript the normalization operation (underline) by $\boldsymbol{\theta}$, but we omit this and take, in any single expression, the state to be a fixed $\boldsymbol{\theta}$.

3. THE CASE OF COMPLETE INFORMATION: SURPLUS OUTCOMES

As a benchmark, we characterize the effects of interventions undertaken by an authority with complete information on the market state. Formally, n is an arbitrary fixed number and $\mathbf{\Theta}(n)$ is a singleton throughout this section.

Proposition 1.

1. For any intervention $\boldsymbol{\sigma}$, we have:
$$\dot{C}_{\boldsymbol{\sigma}} + \frac{1}{2}\dot{P}_{\boldsymbol{\sigma}} = \dot{S}_{\boldsymbol{\sigma}}. \tag{8}$$
Hence, if $\dot{C}_{\boldsymbol{\sigma}} \geq 0$, then $\dot{C}_{\boldsymbol{\sigma}} + \dot{P}_{\boldsymbol{\sigma}} \leq 2\dot{S}_{\boldsymbol{\sigma}}$.
2. Assume $\boldsymbol{D}$ has at least one nonzero off-diagonal entry. For generic $\boldsymbol{q}^0$, any surplus outcome $(\dot{C}, \dot{P}, \dot{S})$ that satisfies (8) can be implemented by some intervention $\boldsymbol{\sigma}$.

It is useful to relate Proposition 1 to the standard single-product monopolist facing linear demand with constant marginal cost. In this case, a small per-unit subsidy σ lowers the equilibrium price by $\sigma/2$ and so the consumer gains half the subsidy expenditure ($\dot{C} = \frac{1}{2}\dot{S}$). The producer, instead, gains $\dot{S}/2$ from the static incidence on the initial quantity, plus another $\dot{S}/2$ from the profit on the newly induced volume, totaling the full subsidy amount $\dot{S}$.

In the monopoly ($n = 1$) setting, these changes satisfy (8). Proposition 1 shows that this aggregate constraint carries over to an oligopoly with price competition.

Proof of part (1). We fix the state $(\boldsymbol{D}, \boldsymbol{q}^0)$ and present the surplus formula under the normalization introduced in Section 2.7. (Surpluses are invariant to the changes of units performed in the normalization, because the unit of money was not changed.) The effect of $\boldsymbol{\sigma}$ on consumer surplus satisfies the Marshallian formula $\dot{C}_{\boldsymbol{\sigma}} = -\underline{\boldsymbol{q}}^0 \cdot \underline{\dot{\boldsymbol{p}}}_{\boldsymbol{\sigma}}$. The effect on producer surplus is

$$\dot{P}_{\boldsymbol{\sigma}} = \underline{\boldsymbol{q}}^0 \cdot (\underline{\dot{\boldsymbol{p}}}_{\boldsymbol{\sigma}} + \underline{\boldsymbol{\sigma}}) + (\underline{\boldsymbol{p}}^0 - \underline{\boldsymbol{c}}) \cdot \underline{\dot{\boldsymbol{q}}}_{\boldsymbol{\sigma}} = 2\underline{\boldsymbol{q}}^0 \cdot \underline{\dot{\boldsymbol{q}}}_{\boldsymbol{\sigma}},$$

where the second equality follows because Equation (7) at $\underline{\boldsymbol{\sigma}} = \mathbf{0}$ yields $\underline{\boldsymbol{q}}^0 = \underline{\boldsymbol{p}}^0 - \underline{\boldsymbol{c}}^0$. Applying $\frac{d}{dt}\big|_{t=0}$ to both sides of Equation (7) along the path $t\underline{\boldsymbol{\sigma}}$ gives $\underline{\dot{\boldsymbol{q}}}_{\boldsymbol{\sigma}} = \underline{\dot{\boldsymbol{p}}}_{\boldsymbol{\sigma}} + \underline{\boldsymbol{\sigma}}$. Hence, any $(\dot{C}_{\boldsymbol{\sigma}}, \dot{P}_{\boldsymbol{\sigma}}, \dot{S}_{\boldsymbol{\sigma}})$ satisfy

$$\dot{C}_{\boldsymbol{\sigma}} + \frac{1}{2}\dot{P}_{\boldsymbol{\sigma}} = -\underline{\boldsymbol{q}}^0[\underline{\dot{\boldsymbol{p}}}_{\boldsymbol{\sigma}} - \underline{\dot{\boldsymbol{q}}}_{\boldsymbol{\sigma}}] = \underline{\boldsymbol{q}}^0 \cdot \underline{\boldsymbol{\sigma}} = \dot{S}_{\boldsymbol{\sigma}},$$

where the first equality follows by substituting the expressions for $\dot{C}_{\boldsymbol{\sigma}}$ and $\dot{P}_{\boldsymbol{\sigma}}$, the second equality follows by using $\underline{\dot{\boldsymbol{q}}}_{\boldsymbol{\sigma}} = \underline{\dot{\boldsymbol{p}}}_{\boldsymbol{\sigma}} + \underline{\boldsymbol{\sigma}}$, and the third equality by (3).

The proof of part (2) uses the techniques we develop in Section 5.1, and we relegate it to Appendix A.1. □

If products are independent or there is only one product (i.e., $\boldsymbol{D}$ is a diagonal matrix), then an intervention implements a specific point of (8): $\dot{C} = \dot{S}/2$ and $\dot{P} = \dot{S}$. Part (2) of Proposition 1 shows that, as long as there is demand interdependency across products (i.e., $\boldsymbol{D}$ is not a diagonal matrix), the authority can leverage the resulting demand spillovers to generate any point that satisfies (8).

This, however, requires precise information about $\boldsymbol{D}$. To see this, note that the total surplus pass-through of an intervention $\boldsymbol{\sigma}$ is[15]

$$\dot{W}_{\boldsymbol{\sigma}} = \underline{\boldsymbol{q}}^0 \cdot \underline{\boldsymbol{D}}\underline{\dot{\boldsymbol{p}}}_{\boldsymbol{\sigma}} = -\underline{\boldsymbol{q}}^0 \cdot \underline{\boldsymbol{D}}[\boldsymbol{I} - \underline{\boldsymbol{D}}]^{-1}\underline{\boldsymbol{\sigma}}. \tag{9}$$

Matrix inverses such as $[\boldsymbol{I} - \underline{\boldsymbol{D}}]^{-1}$ can be extremely sensitive to the entries of $\boldsymbol{D}$. Therefore, without precise knowledge of $\boldsymbol{D}$ and $\boldsymbol{q}^0$, the authority may not be able to implement a desired point on the Pareto frontier defined by (8). In particular, the authority may not even be sure that a given intervention will increase total surplus ($\dot{W}_{\boldsymbol{\sigma}} > 0$) rather than decrease it ($\dot{W}_{\boldsymbol{\sigma}} < 0$). Indeed, it seems hard to justify the detailed study of comparative statics such as (9), at least when n is large, without confronting the uncertainty about the ingredients of the formula an analyst or authority is likely to face.

These observations motivate the central question of this paper: *Which interventions have surplus effects that can be predicted with confidence by an authority facing substantial uncertainty about market primitives?* The rest of the paper addresses this question.

The main result of this paper (Theorem 1 in Section 6) provides a sufficient condition—called "significant structure"—for the existence of intervention rules that improve surplus robustly, and characterizes these interventions. The next section illustrates these ideas using a simple example.

4. Illustrative example: A block model of demand

To motivate our main definition (presented in Section 5.2) we consider a simple example within a canonical oligopoly model and give a version of our main result in this setting (Proposition 2).

Suppose there are $K > 1$ different product categories, each with a fixed number $m > 1$ of products, for a total of $n = mK$ products. Throughout this example, the asymptotic sequence fixes m and lets K grow, so $n = mK$ grows by adding product categories. Products in each category are (imperfect) substitutes, while products across categories are complements. For instance, one product category may consist of different tennis rackets, while another consists of different tennis shoes.

Formally, we take $\mathcal{E} = (\boldsymbol{\Theta}(n), \varphi(n))_n$ to be an environment defined by a block model of demand. For a given state $\boldsymbol{\theta} = (\boldsymbol{D}, \boldsymbol{q}^0) \in \boldsymbol{\Theta}(n)$ the demand relationships between products are specified by a K-by-K symmetric matrix $\boldsymbol{B}(\boldsymbol{\theta})$:

$$D_{ij} = \begin{cases} -1 & \text{if } i = j, \\ B_{k(i),k(i)} > 0 & \text{if } k(i) = k(j) \text{ and } i \neq j \\ B_{k(i),k(j)} < 0 & \text{if } k(i) \neq k(j), \end{cases}$$

where $k(i)$ is the product category of i.[16] We assume without loss of generality (Remark 3 in the appendix) that the quantities satisfy $\|\boldsymbol{q}^0\| = 1$. Figure 1 depicts a two-product-category demand model.

Condition 1. Assume that $\mathcal{E} = (\boldsymbol{\Theta}(n), \varphi(n))_n$ satisfies the following conditions:

[15] To see this, note that $\dot{W}_{\boldsymbol{\sigma}} = \dot{C}_{\boldsymbol{\sigma}} + \dot{P}_{\boldsymbol{\sigma}} - \dot{S}_{\boldsymbol{\sigma}}$ and, using (8) and the expressions for $\dot{P}$, we obtain $\dot{W}_{\boldsymbol{\sigma}} = \dot{P}_{\boldsymbol{\sigma}}/2 = \underline{\boldsymbol{q}}^0 \cdot (\underline{\boldsymbol{D}}\underline{\dot{\boldsymbol{p}}}_{\boldsymbol{\sigma}}) = -\underline{\boldsymbol{q}}^0 \cdot \underline{\boldsymbol{D}}[\boldsymbol{I} - \underline{\boldsymbol{D}}]^{-1}\underline{\boldsymbol{\sigma}}$.

[16] Note that, in this example, we have made the simplifying assumption that, for each market state $\boldsymbol{\theta} = (\boldsymbol{D}, \boldsymbol{q}^0)$, $D_{ii} = -1$ for all i (so that no normalization is required).

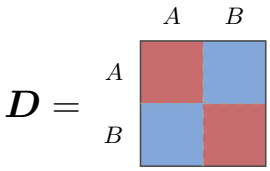


FIGURE 1. Two product-category demand: products within a category are substitutes (red blocks) and across categories products are complements (blue blocks)

(a) There exists $\kappa < 0$ and $\delta > 0$ such that, in every market state $\boldsymbol{\theta} \in \boldsymbol{\Theta}(n)$:
  (1) There is a minimum level of complementarity across product categories: $D_{ij} < \kappa$ whenever $k(i) \neq k(j)$;
  (2) The quantities satisfy $\sum_i q_i^0 > \delta\sqrt{n}$.
(b) The errors E_{ij} in observing $\boldsymbol{D}$ are independent, mean-zero, and bounded uniformly over $\boldsymbol{\Theta}(n)$ and n. The quantity $\boldsymbol{q}^0$ is observed without noise.

Section 8.3 shows that Condition 1(a) is consistent with a demand function derived by optimal consumption of a representative consumer with quadratic preferences (see, e.g., Singh and Vives, 1984), and that the pricing game outcome is well behaved when K grows large. Appendix D shows that Condition 1(b) is consistent with an authority estimating the market state by sampling households making two-good tradeoffs.

Part (a)(2) of Condition 1 is worth a brief remark: it is satisfied by the norm-1 quantity vector given by $q_i^0 = \frac{1}{\sqrt{n}}$ for all i (uniform quantities), but not satisfied by a vector that puts quantity 1 in some fixed i and 0 elsewhere. It amounts to a condition that quantity is not concentrated too tightly on a small set of indices.

The following proposition is the main result of our paper instantiated in this block model of demand (it follows from Theorem 1 in Section 6).

Proposition 2. Consider the block model of demand and assume Condition 1 holds. For every $\varepsilon > 0$ and for every target expenditure $s > 0$, there is an intervention rule that, if the number of categories K is large enough along this sequence, ϵ-robustly achieves the property that $(\dot{P}, \dot{C}, \dot{S})$ is within ε Euclidean distance of $(2s, 0, s)$.

Recall that Proposition 1 established that gross surplus $\dot{C} + \dot{P}$ cannot be increased by more than twice the level of expenditure without reducing consumer surplus. Proposition 2 provides a set of assumptions that are sufficient for the existence of robust intervention rules that achieve this bound.

These assumptions allow recovery of some partial information about the market state that is sufficient to design intervention rules that increase total surplus robustly without decreasing consumer surplus. In particular, our approach to constructing robust interventions takes the following route:

(1) Show that, for targeting interventions to achieve good surplus outcomes, it is sufficient to have information on certain eigenvectors of $\underline{\boldsymbol{D}}$—those associated with large eigenvalues.
(2) State conditions under which those eigenvectors can be recovered from the noisy signal $\widehat{\boldsymbol{D}} = \boldsymbol{D} + \boldsymbol{E}$.

The argument for (1) leverages a new spectral decomposition of how interventions pass through to surpluses that we present in Section 5.1. Section 5.2 then formalizes the conditions on market states and noise needed to recover eigenvectors with associated large eigenvalues.

5. Spectral price theory and significant structure

5.1. **A key tool: Spectral price theory.** Denoting by $\boldsymbol{U}$ the matrix whose ℓ^{th} column is the ℓ^{th} eigenvector $\boldsymbol{u}^\ell$ of $\underline{\boldsymbol{D}}$, and by $\boldsymbol{\Lambda}$ the matrix whose off-diagonal elements are zero and whose ℓ^{th} diagonal element is the eigenvalue λ_ℓ associated with $\boldsymbol{u}^\ell$, we have:

$$\underline{\boldsymbol{D}} = \boldsymbol{U}\boldsymbol{\Lambda}\boldsymbol{U}^{\mathsf{T}}.$$

We can think of the eigenvectors of $\underline{\boldsymbol{D}}$ as the principal components of the normalized Slutsky matrix. An intervention $\boldsymbol{\sigma}$ that subsidizes (or taxes) a single product affects the prices and quantities of that product and, through strategic interaction, also those of other products. However, if we think of the eigenvector $\boldsymbol{u}^\ell$ as a bundle of products, an intervention $\underline{\boldsymbol{\sigma}} \propto \boldsymbol{u}^\ell$ only passes through to the price $\boldsymbol{u}^\ell \cdot \underline{\boldsymbol{p}}$ and quantity $\boldsymbol{u}^\ell \cdot \underline{\boldsymbol{q}}$ of that bundle; the prices and quantities of all the other bundles $\boldsymbol{u}^{\ell'}$ are unchanged. Hence, we can decompose an intervention $\underline{\boldsymbol{\sigma}} = \sum_\ell (\boldsymbol{u}^\ell \cdot \underline{\boldsymbol{\sigma}})\boldsymbol{u}^\ell$ into a combination of n orthogonal interventions, each in the direction of an eigenvector. We can use this decomposition to obtain simple expressions for the pass-through of the intervention in terms of the eigenvalues of $\underline{\boldsymbol{D}}$ as summarized by the next Lemma.

Lemma 1. The pass-throughs from any intervention $\boldsymbol{\sigma}$ to prices and quantities of each eigenvector are as follows:

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

Proof. Using the normalization, the equation (5) becomes $[\boldsymbol{I} - \underline{\boldsymbol{D}}]\underline{\dot{\boldsymbol{p}}}_{\boldsymbol{\sigma}} = -\underline{\boldsymbol{\sigma}}$; substituting $\underline{\boldsymbol{D}} = \boldsymbol{U}\boldsymbol{\Lambda}\boldsymbol{U}^{\mathsf{T}}$ we get $\left(\boldsymbol{I} - \boldsymbol{U}\boldsymbol{\Lambda}\boldsymbol{U}^{\mathsf{T}}\right)\underline{\dot{\boldsymbol{p}}}_{\boldsymbol{\sigma}} = -\underline{\boldsymbol{\sigma}}$. Multiplying both sides by $\boldsymbol{U}^{\mathsf{T}}$, we get the vector equation $\boldsymbol{U}^{\mathsf{T}}\underline{\dot{\boldsymbol{p}}}_{\boldsymbol{\sigma}} = -(\boldsymbol{I} - \boldsymbol{\Lambda})^{-1}\boldsymbol{U}^{\mathsf{T}}\underline{\boldsymbol{\sigma}}$ and, using equation (5), we also get $\boldsymbol{U}^{\mathsf{T}}\underline{\dot{\boldsymbol{q}}}_{\boldsymbol{\sigma}} = -\boldsymbol{\Lambda}(\boldsymbol{I} - \boldsymbol{\Lambda})^{-1}\boldsymbol{U}^{\mathsf{T}}\underline{\boldsymbol{\sigma}}$. The scalar formulas in Lemma 1 follow from (i) taking the ℓ-th coordinate of these vector equations, using that $\boldsymbol{U}$ is orthonormal and $(\boldsymbol{I} - \boldsymbol{\Lambda})^{-1}$ is diagonal, and (ii) recalling from Property NSD that $\underline{\boldsymbol{D}}$ is negative semidefinite, so $\lambda_\ell \leq 0$ and hence $1 - \lambda_\ell = 1 + |\lambda_\ell|$. □

The magnitudes of the price and quantity pass-throughs in each direction $\boldsymbol{u}^\ell$ are ordered according to their corresponding eigenvalues: The larger $|\lambda_\ell|$ is, the less a subsidy $\boldsymbol{u}^\ell \cdot \underline{\boldsymbol{\sigma}}$ 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 $\boldsymbol{u}^\ell \cdot \underline{\boldsymbol{p}}$ is less sensitive to the subsidy $\boldsymbol{u}^\ell \cdot \underline{\boldsymbol{\sigma}}$ the larger is $|\lambda_\ell|$. On the other, the demand $\boldsymbol{u}^\ell \cdot \underline{\boldsymbol{q}}$ is more sensitive to the price $\boldsymbol{u}^\ell \cdot \underline{\boldsymbol{p}}$ the larger is $|\lambda_\ell|$; this is just a fact about the market's demand function, rather than equilibrium pricing.[17] 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 $\boldsymbol{u}^\ell \cdot \underline{\boldsymbol{q}}$ to the subsidy $\boldsymbol{u}^\ell \cdot \underline{\boldsymbol{\sigma}}$.

[17] Indeed, it follows from (5) that $\boldsymbol{U}^{\mathsf{T}}\underline{\dot{\boldsymbol{q}}} = \boldsymbol{\Lambda}\boldsymbol{U}^{\mathsf{T}}\underline{\dot{\boldsymbol{p}}}$, so the derivative of the demand $\boldsymbol{u}^\ell \cdot \underline{\boldsymbol{q}}$ with respect to $\boldsymbol{u}^\ell \cdot \underline{\boldsymbol{p}}$ is equal to λ_ℓ.

By combining the spectral decomposition of Lemma 1 with the surplus formulas derived in Section 2.4 we obtain that the effect of an intervention on consumer, producer, or total surplus is a weighted sum of pass-throughs to each of the eigenvectors $\boldsymbol{u}^\ell$—with the weight being the corresponding bundle's quantity $\boldsymbol{u}^\ell \cdot \underline{\boldsymbol{q}}^0$.

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

$$\dot{C}_{\boldsymbol{\sigma}} = -\sum_{\ell=1}^{n} (\boldsymbol{u}^\ell \cdot \underline{\boldsymbol{q}}^0)(\boldsymbol{u}^\ell \cdot \underline{\dot{\boldsymbol{p}}}_{\boldsymbol{\sigma}}) \quad \dot{P}_{\boldsymbol{\sigma}} = 2\sum_{\ell=1}^{n} (\boldsymbol{u}^\ell \cdot \underline{\boldsymbol{q}}^0)(\boldsymbol{u}^\ell \cdot \underline{\dot{\boldsymbol{q}}}_{\boldsymbol{\sigma}}) \quad \text{and} \quad \dot{W}_{\boldsymbol{\sigma}} = \dot{P}_{\boldsymbol{\sigma}}/2.$$

Proof. The effect of the intervention on consumer surplus is $\dot{C}_{\boldsymbol{\sigma}} = -\underline{\boldsymbol{q}}^0 \cdot \underline{\dot{\boldsymbol{p}}}_{\boldsymbol{\sigma}}$. Using the fact that the eigenvectors form an orthonormal basis ($\boldsymbol{U}\boldsymbol{U}^\mathsf{T} = \boldsymbol{I}$), we insert the identity matrix into the dot product to obtain $\dot{C}_{\boldsymbol{\sigma}} = -\boldsymbol{U}^\mathsf{T}\underline{\boldsymbol{q}}^0 \cdot \boldsymbol{U}^\mathsf{T}\underline{\dot{\boldsymbol{p}}}_{\boldsymbol{\sigma}}$. Similarly, the effect of the intervention on producer surplus is $\dot{P}_{\boldsymbol{\sigma}} = \underline{\boldsymbol{q}}^0 \cdot (\underline{\dot{\boldsymbol{p}}}_{\boldsymbol{\sigma}} + \underline{\boldsymbol{\sigma}}) + (\underline{\boldsymbol{p}}^0 - \underline{\boldsymbol{c}}) \cdot \underline{\dot{\boldsymbol{q}}}_{\boldsymbol{\sigma}} = 2\underline{\boldsymbol{q}}^0 \cdot \underline{\dot{\boldsymbol{q}}}_{\boldsymbol{\sigma}}$. Multiplying by $\boldsymbol{U}\boldsymbol{U}^\mathsf{T}$ yields the expression of $\dot{P}_{\boldsymbol{\sigma}}$ in the Lemma. The expression for $\dot{W}_{\boldsymbol{\sigma}}$ is obtained by aggregating $\dot{P}_{\boldsymbol{\sigma}}$, $\dot{C}_{\boldsymbol{\sigma}}$ and the intervention expenditure. □

We now explain why identifying large-eigenvalue eigenvectors of $\underline{\boldsymbol{D}}$ is helpful for the design of interventions that achieve a certain point on the full-information surplus frontier, i.e., Equation (8). Indeed, suppose the authority knows $\boldsymbol{u}^1$ and evaluates the effect of intervention $\underline{\boldsymbol{\sigma}} = \boldsymbol{u}^1$. This intervention costs $\dot{S}_{\boldsymbol{\sigma}} = \boldsymbol{u}^1 \cdot \underline{\boldsymbol{q}}^0$ and leads to an overall change in consumer and producer surplus equal to

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

Hence, if:

a. the normalized status quo quantity $\underline{\boldsymbol{q}}^0$ is not orthogonal to $\boldsymbol{u}^1$, which implies that $\dot{S}_{\boldsymbol{\sigma}} \neq 0$, and
b. the eigenvalue λ_1 associated with $\boldsymbol{u}^1$ is, in absolute value, sufficiently large,

then the intervention $\underline{\boldsymbol{\sigma}} = \boldsymbol{u}^1$ achieves $\dot{C}_{\boldsymbol{\sigma}}/\dot{S}_{\boldsymbol{\sigma}} \approx 0$ and $\dot{P}_{\boldsymbol{\sigma}}/\dot{S}_{\boldsymbol{\sigma}} \approx 2$. By Proposition 1, this rate of surplus increase per dollar spent is the highest possible subject to the constraint that the change in consumer surplus is non-negative. At this point, the sign of $\dot{S}$, $\dot{P}$, $\dot{W}$ is undetermined, because the sign of $\boldsymbol{u}^1 \cdot \underline{\boldsymbol{q}}^0$ is undetermined.[18] By choosing the sign of the intervention appropriately, we can achieve an intervention that increases total surplus. It is not important that we design the intervention proportional to $\boldsymbol{u}^1$; if an intervention is approximately contained in a subspace spanned by eigenvectors that satisfy conditions (a) and (b), then the same reasoning applies to it.

This discussion suggests that recovering information about large eigenvalues of $\underline{\boldsymbol{D}}$ and their eigenvectors is very helpful for designing an intervention. The next section develops a set of conditions on the underlying market state and observation noise, called *significant structure*, that ensures the authority can do this.

[18] Recall that $\boldsymbol{u}^1$ and $-\boldsymbol{u}^1$ are equally good eigenvectors of $\underline{\boldsymbol{D}}$, so by bad luck we may have chosen a $\boldsymbol{u}^1$ that makes spending and surplus change negative.

5.2. **Significant structure.** For matrices, $\|\cdot\|$ is the induced operator norm, which for symmetric matrices is the maximum absolute value of the eigenvalues. We also denote by $\mathcal{L}(\underline{\boldsymbol{D}}, b) \subseteq \mathbb{R}^n$ the space spanned by eigenvectors of $\underline{\boldsymbol{D}}$ with eigenvalues at least b in absolute value.

Definition 2. An environment $\mathcal{E} = (\boldsymbol{\Theta}(n), \varphi(n))_n$ has significant structure if there is a sequence $b(n) \to \infty$ satisfying the following conditions.

(1) The quantity error vector $\boldsymbol{\varepsilon}$ is comparable to the quantity vector, in the sense that $\|\boldsymbol{\varepsilon}\| = O_{\mathrm{p},\boldsymbol{\Theta}}(\|\boldsymbol{q}^0\|)$.
(2) The demand observation noise is negligible relative to $b(n)$ in the sense that $\|\boldsymbol{E}\| = o_{\mathrm{p},\boldsymbol{\Theta}}(b(n))$.
(3) Some eigenvalues of $\underline{\boldsymbol{D}}$ have absolute value at least $b(n)$ and normalized quantities have non-vanishing projection onto the eigenspaces spanned by corresponding eigenvectors: there exists $\delta > 0$ such that, uniformly over market states and n,

$$\left\| P_{\mathcal{L}(\underline{\boldsymbol{D}}, b(n))}\, \underline{\boldsymbol{q}}^0 \right\| \geq \delta \|\underline{\boldsymbol{q}}^0\|.$$

Note that in (1) and (2), the asymptotic notation requires that the statement hold uniformly over market states, under the error distribution $\varphi_{\boldsymbol{\theta}}(n)$ corresponding to that market state.

Condition (1) says that the quantity signal $\widehat{\boldsymbol{q}}^0 = \boldsymbol{q}^0 + \boldsymbol{\varepsilon}$ is not swamped by noise in overall magnitude. It permits the application of laws of large numbers to identify weighted moments of $\boldsymbol{q}^0$ from $\widehat{\boldsymbol{q}}^0$. Intuitively, if quantities are too small relative to the noise, the authority has no way to identify or target the sign of terms in the welfare decomposition, so it is not surprising that such a condition should be involved.

In the remaining conditions, a sequence $b(n)$ defines a diverging threshold for eigenvalues being considered "large." Condition (2) requires that the demand-signal noise $\boldsymbol{E}$ is small in operator norm relative to this threshold, while condition (3) requires that $\underline{\boldsymbol{D}}$ has some eigenvalues whose absolute value exceeds it. Together, conditions (2) and (3) ensure that the eigenvectors associated with large eigenvalues can be recovered from the noisy observation $\underline{\widehat{\boldsymbol{D}}}$; the key tool for this is the Davis–Kahan theorem, discussed in Section 6.1.1.

Condition (3) also requires that $\underline{\boldsymbol{q}}^0$ has a non-vanishing projection onto the eigenvectors of large eigenvalues. Together with (1), this ensures that the authority can align interventions well enough so that, at prevailing quantities, an intervention based on major eigenvectors has non-trivial expenditure. If, instead, the status-quo quantities were asymptotically orthogonal to the recoverable eigenspace, the identifiable structure of $\underline{\boldsymbol{D}}$ would be unrelated to the quantities, leading to negligible and/or very noisy welfare effects from interventions based on this structure.

5.2.1. *Significant structure: Example.* We now illustrate the concept of significant structure by presenting an instance of the block model of demand in Section 4 in which significant structure is present. For simplicity, this instance has the feature that the substitutability within each of the K categories is homogeneous, as is the complementarity among them. That is, for market state $\boldsymbol{\theta} \in \boldsymbol{\Theta}(n)$, the

Slutsky matrix is

$$D_{ij} = \begin{cases} -1 & \text{if } i = j, \\ \alpha_{\boldsymbol{\theta}} < 0 & \text{if } k(i) \neq k(j), \\ \omega_{\boldsymbol{\theta}} > 0 & \text{otherwise.} \end{cases}$$

The associated eigenvalues of $\boldsymbol{D}$ are

$$\begin{array}{ll} \lambda_3(\boldsymbol{\theta}) = -1 - \omega_{\boldsymbol{\theta}} & \text{with multiplicity } K(m-1) \\ \lambda_2(\boldsymbol{\theta}) = \lambda_3(\boldsymbol{\theta}) + m(\omega_{\boldsymbol{\theta}} - \alpha_{\boldsymbol{\theta}}) & \text{with multiplicity } (K-1) \\ \lambda_1(\boldsymbol{\theta}) = \lambda_2(\boldsymbol{\theta}) + mK\alpha_{\boldsymbol{\theta}} & \text{with multiplicity } 1. \end{array}$$

Property NSD is satisfied if and only if all these eigenvalues are negative, which is the case as long as $\lambda_2(\boldsymbol{\theta}) < 0$ for all $\boldsymbol{\theta}$. This imposes an upper bound on the number of products within each category, namely,

$$m < m(n) := \inf_{\boldsymbol{\theta} \in \boldsymbol{\Theta}(n)} \lfloor (1 + \omega_{\boldsymbol{\theta}})/(\omega_{\boldsymbol{\theta}} - \alpha_{\boldsymbol{\theta}}) \rfloor.$$

Consider an environment $\mathcal{E} := \{(\boldsymbol{\Theta}(n), \varphi(n))\}_n$ constructed by successively adding product categories of fixed size m, with $n = mK$ and fixed $m < m(n)$ along the sequence. Assuming that the environment satisfies Condition 1[19], we demonstrate that it exhibits significant structure.

First, quantities are observed without noise in this example, which takes care of the conditions involving quantity errors.[20] Next, when the number K of categories is large enough, the dominant eigenvalue (in absolute value) is $\lambda_1(\boldsymbol{\theta})$, so we take

$$b(n) = \inf_{\boldsymbol{\theta} \in \boldsymbol{\Theta}(n)} |\lambda_1(\boldsymbol{\theta})| = \inf_{\boldsymbol{\theta} \in \boldsymbol{\Theta}(n)} \big(1 - \omega_{\boldsymbol{\theta}}(m-1) + m(K-1)|\alpha_{\boldsymbol{\theta}}|\big).$$

This threshold grows linearly in n because cross-category complementarities are bounded away from zero.[21] Moreover, under the observation-error assumptions in Condition 1(b), standard random matrix bounds give $\|\boldsymbol{E}\| = O_{\mathrm{p},\boldsymbol{\Theta}}(\sqrt{n})$, and hence $\|\boldsymbol{E}\|/b(n) = o_{\mathrm{p},\boldsymbol{\Theta}}(1)$. (See, for example, Vershynin (2018); Bandeira and van Handel (2016).) Finally, for large K, the eigenvector associated with λ_1 is proportional to $\mathbf{1}$, and therefore the projection of $\boldsymbol{q}^0$ onto $\mathcal{L}(\underline{\boldsymbol{D}}, b(n))$ is proportional to $\sum_i q_i^0$ and bounded away from 0.[22]

Hence, this environment has significant structure. As we will see, significant structure allows us to estimate accurately, when n is sufficiently large, the eigenvector $\boldsymbol{u}^1$ associated with the true market state $\boldsymbol{\theta} \in \boldsymbol{\Theta}(n)$. Using the insights from Section 5.1, we can then design an intervention $\boldsymbol{\sigma} \propto \boldsymbol{u}^1$ and, since the associated $|\lambda_1|$ is very large (and converges to ∞ as n grows), this intervention leads to the outcome described in Proposition 2.

We have shown significant structure in the block model where demand interdependencies are homogeneous within a category and across categories. A generalization of these arguments applies to the environment of Proposition 2. The

[19] That is, cross-category complementarities are uniformly bounded away from zero, aggregate quantity satisfies $\sum_i q_i^0 > \delta\sqrt{n}$, and the observation errors E_{ij} are independent, mean-zero, with variances bounded uniformly over $\boldsymbol{\Theta}(n)$ and n.

[20] Formally, $\boldsymbol{\varepsilon} = \mathbf{0}$, so $\|\boldsymbol{\varepsilon}\|/\|\boldsymbol{q}^0\| = 0$, as required in part (1) of Definition 2.

[21] Under Condition 1(a)(1), we have $|\alpha_{\boldsymbol{\theta}}| \geq |\kappa| > 0$ uniformly, so $b(n) = \Theta(K) = \Theta(n)$ when $n = mK$ grows by increasing K with m fixed.

[22] When $\boldsymbol{u}^1 \propto \mathbf{1}$, we have $|\boldsymbol{u}^1 \cdot \boldsymbol{q}^0| = (\sum_i q_i^0)/\sqrt{n}$; Condition 1(a)(2) imposes $\sum_i q_i^0 > \delta\sqrt{n}$, which yields part (3) of Definition 2.

important properties are simply that the dominant eigenvalue grows, in absolute value, linearly in n, whereas the operator norm of the observation noise grows slower, only at rate $\sqrt{n}$. This minimum level of complementarities implies that as K increases, at least one low-dimensional and significant part of the complementarity pattern can be accurately estimated.

Whereas this section has illustrated plausible conditions on interactions for significant structure to be present, practically speaking it is important to explain how it can be verified in applications. We take up such tests in Section 8.2.

6. Main result

Our main result is that in an economy with significant structure, the authority can design interventions that substantially increase total surplus. To state the result we introduce one technical assumption, which simplifies the proof considerably.

Assumption 4 (Recoverable diagonal)**.** In the environment $\mathcal{E} := (\boldsymbol{\Theta}(n), \varphi(n))_n$, for all $(\boldsymbol{D}, \boldsymbol{q}^0) \in \boldsymbol{\Theta}(n)$ and for all n, there exist mappings $\Delta_n : \mathbb{R}^{n\times n} \to \mathbb{R}^n$ (defined for all n) such that $\max_i \left|\Delta_n(\boldsymbol{D}+\boldsymbol{E})_i - D_{ii}\right| = o_{\mathrm{p},\boldsymbol{\Theta}}(1)$.

This assumes that, for each state $\boldsymbol{\theta}$, the diagonal of $\boldsymbol{D}$ (the own-price effects) can be estimated consistently from the noisy observation $\widehat{\boldsymbol{D}}$. It allows the authority to perform calculations with the normalized version of the noisy observation $\underline{\widehat{\boldsymbol{D}}}$. For instance, the assumption holds when $\max_i |E_{ii}| = o_{\mathrm{p},\boldsymbol{\Theta}}(1)$. In this case we can take $\Delta_n(\boldsymbol{D}+\boldsymbol{E}) = \mathrm{diag}(\boldsymbol{D}+\boldsymbol{E})$. For a natural setting in which this assumption holds, see Appendix D.[23]

Theorem 1. Consider an environment with significant structure satisfying Assumptions 1–4. For every $\epsilon > 0$ and for every target expenditure $s > 0$, the following properties can be simultaneously achieved ϵ–robustly for sufficiently large n:

(i) The sum of marginal consumer and producer surplus gains $\dot{C} + \dot{P}$ is at least twice the marginal expenditure, up to a small multiplicative error: $\dot{C} + \dot{P} \geq (2-\epsilon)\dot{S}$.

(ii) The marginal effect on consumer surplus is vanishing: $|\dot{C}| < \epsilon$. Moreover, no individual consumer's surplus changes significantly—i.e., $|\dot{C}^h| < \epsilon$ for each household h.

(iii) The marginal expenditure is arbitrarily close to the target expenditure s, i.e., $|\dot{S} - s| < \epsilon$.

Part (i) states that the authority can robustly achieve at least two dollars of surplus per dollar spent. Part (ii) states that it is possible to achieve this while leaving consumer surplus essentially unchanged. In particular, under the robust intervention rule we construct, producer surplus increases by approximately two dollars per dollar spent, while consumer surplus changes only by a vanishing amount.[24] Finally, part (iii) says the authority can precisely target

[23] We believe the assumption can be relaxed without changing the possibility of robust interventions, but it does simplify arguments considerably.

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

the realized expenditure (and thus the total surplus impact) of the intervention. Comparing the theorem with Proposition 1, there is a robust intervention that, subject to a given target expenditure $\dot{S} = s$, achieves approximately the maximum total surplus possible subject to $\dot{C} \geq 0$. In Section 8.1 we will see that at least in some cases, this is the only welfare outcome that can be implemented robustly (see part (2) of Proposition 3).

We now provide intuition for the properties of the intervention. First recall that, by Lemma 1, the pass-through to prices is proportional to $1/(1 + |\lambda_\ell|)$, while the pass-through to quantities is proportional to $|\lambda_\ell|/(1 + |\lambda_\ell|)$. Combining this with Lemma 2, each expenditure component $(\boldsymbol{u}^\ell \cdot \underline{\boldsymbol{q}}^0)(\boldsymbol{u}^\ell \cdot \underline{\boldsymbol{\sigma}})$ contributes $1/(1 + |\lambda_\ell|)$ to consumer surplus, $2|\lambda_\ell|/(1 + |\lambda_\ell|)$ to producer surplus, and $|\lambda_\ell|/(1+|\lambda_\ell|)$ to net total surplus. When $|\lambda_\ell| \geq M$ for a large M, each dollar of intervention yields consumer-surplus incidence $O(1/M)$, while the producer-surplus incidence is $2 - O(1/M)$. It is critical to the first part that prices move only slightly, but this small movement creates a quantity effect sufficient to generate the gains in the second part.

The same pass-through reasoning applies to any individual household after replacing $\underline{\boldsymbol{q}}^0$ by that household's normalized status-quo quantity vector. Thus, we can upgrade the statement about consumer surplus to a uniform one across households (provided consumers are not compelled to make other transfers); this is especially valuable when targeting consumers in the product market for personalized transfers is not straightforward. Of course, the authority *may* separately collect producer surplus through fixed transfers and redistribute it to consumers.[25] Absent such transfers, the result yields a robust total-surplus intervention with negligible direct consumer-surplus incidence in the strong "for all households" sense we have described.

Remark 1. Our notion of ϵ–robustness guarantees these surplus properties can also be stated as bounds on *expected* performance conditional on each possible $\boldsymbol{\theta}$. The reason is that interventions can be readily adapted so that relevant welfare functionals become uniformly bounded; therefore, as we explain in Appendix A.4, the convergence-in-probability conclusions in Theorem 1 upgrade to expectation statements.

6.1. **Sketch of the proof.** To prove Theorem 1, we apply a statistical method that accurately identifies a subspace spanned by top eigenvectors of the Slutsky matrix from noisy observations. We then use the spectral price theory developed in Section 5.1 to show that interventions projecting exclusively onto this subspace possess the welfare properties stated in Theorem 1.

6.1.1. *Recovering the subspace of top eigenvectors: The Davis–Kahan theorem.* The key tool in our statistical exercise that leverages the property of significant structure is the Davis–Kahan theorem. Under the hypothesis that some eigenvalues of $\underline{\boldsymbol{D}}$ are large, standard eigenvalue perturbation bounds guarantee that, despite the noise in $\boldsymbol{E}$, the large eigenvalues of the normalized observed matrix are good approximations of the true large eigenvalues of $\underline{\boldsymbol{D}}$. In other words, the noise in $\boldsymbol{E}$ cannot cause the large eigenvalues of $\underline{\boldsymbol{D}}$ to become "mixed up" with the eigenvalues far away in the spectrum; see Figure 2 for an illustration. The

[25] For example, if the authority is a platform, possible mechanisms include commissions or entry fees charged to firms, along with rebates to consumers.

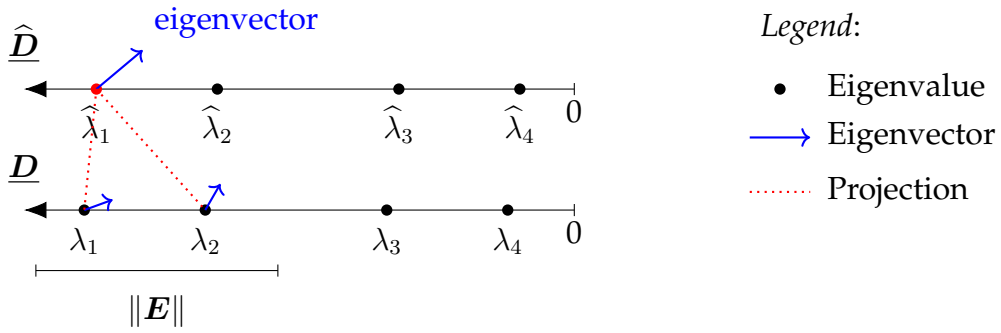


FIGURE 2. An illustration of the Davis–Kahan theorem: How eigenvectors of $\widehat{\underline{D}}$ (perturbed matrix) project onto eigenvectors of $\underline{D}$ (true matrix) with similar eigenvalues (off by at most $\|\boldsymbol{E}\|$). This relationship ensures that the subspace generated by the "top" (i.e., large-eigenvalue) eigenvectors of $\widehat{\underline{D}}$ is a good approximation of the subspace generated by the top eigenvectors of $\underline{D}$. For our economic problem this implies that interventions based on top eigenvectors of $\widehat{\underline{D}}$ yield high surplus pass-through, despite noise $\boldsymbol{E}$.

Davis–Kahan theorem then permits the recovery of eigenvectors. More precisely, these perturbation results have the following two central implications in our setting: if $\underline{D}$ has large eigenvalues then:

(i) the normalized matrix $\widehat{\underline{D}}$ of $\widehat{\boldsymbol{D}} = \boldsymbol{D} + \boldsymbol{E}$ has some eigenvalues that are themselves large;
(ii) the eigenvectors of $\widehat{\underline{D}}$ associated with such eigenvalues can be expressed (up to a small error) as linear combinations of eigenvectors of $\underline{D}$ with large eigenvalues.

6.1.2. *Intervention rules based on recovered eigenvectors.* To facilitate the illustration, we assume that the largest eigenvalue of $\underline{D}$ is sufficiently well-separated from all other eigenvalues by a "gap" much larger than $\|\boldsymbol{E}\|$. Under this condition, the Davis–Kahan theorem yields an even stronger implication: we can use $\widehat{\underline{D}}$ to recover a normalized eigenvector $\widehat{\boldsymbol{u}}^1$ that correlates almost perfectly with the corresponding eigenvector $\boldsymbol{u}^1$ of $\underline{D}$. This property is, for example, satisfied by the block model of demand in Section 5.2.1.

Lemma 1 and Lemma 2 together imply that the effect of intervention $\boldsymbol{\sigma}$ is

$$\dot{W}_{\boldsymbol{\sigma}} = \sum_{\ell=1}^{n} (\boldsymbol{u}^\ell \cdot \underline{\boldsymbol{q}}^0)(\boldsymbol{u}^\ell \cdot \underline{\boldsymbol{\sigma}}) \frac{|\lambda_\ell|}{1+|\lambda_\ell|} \quad \text{and} \quad \dot{S}_{\boldsymbol{\sigma}} = \sum_{\ell=1}^{n} (\boldsymbol{u}^\ell \cdot \underline{\boldsymbol{q}}^0)(\boldsymbol{u}^\ell \cdot \underline{\boldsymbol{\sigma}}). \tag{10}$$

Let us design an intervention $\underline{\boldsymbol{\sigma}} \propto \widehat{\boldsymbol{u}}^1$ based on the recovered information and choose the sign of $\underline{\boldsymbol{\sigma}}$ to ensure $(\widehat{\boldsymbol{u}}^1 \cdot \underline{\boldsymbol{q}}^0)(\widehat{\boldsymbol{u}}^1 \cdot \underline{\boldsymbol{\sigma}})$ is positive. Since, by the Davis–Kahan theorem, $\widehat{\boldsymbol{u}}^1$ is effectively equal to its true counterpart $\boldsymbol{u}^1$ with a very small error, from now on we will treat $\boldsymbol{u}^1$ as known. We can see from the equations above that $\dot{W}_{\boldsymbol{\sigma}}$ closely approximates $\dot{S}_{\boldsymbol{\sigma}}$, since a sufficiently large $|\lambda_1|$ implies $\frac{|\lambda_1|}{1+|\lambda_1|} \approx 1$. Moreover, if we know $\boldsymbol{u}^1 \cdot \underline{\boldsymbol{q}}^0$ and this differs from zero (which is a requirement of significant structure in the illustrative case where $\mathcal{L}(\underline{\boldsymbol{D}}, b(n))$ is one-dimensional), we can scale the intervention to be of the size that we desire, and achieve $\dot{S}_{\boldsymbol{\sigma}} = s$.

This argument contains some wishful thinking, however. When we arranged the sign of $\boldsymbol{\sigma}$ so that $\dot{S}_{\boldsymbol{\sigma}} > 0$, we did not consider that we only have a noisy observation $\widehat{\boldsymbol{q}}^0$ of $\boldsymbol{q}^0$. So part of the challenge of the proof is to manage the observation error that makes $\widehat{\boldsymbol{q}}^0$ different from $\boldsymbol{q}^0$, and to show that we can obtain a correct estimate of the sign with probability arbitrarily close to 1 as $n \to \infty$. If we fail to do this correctly, our intervention will decrease welfare with positive probability. This explains the necessity of controlling the magnitude of the noise in the quantity signal so that its projections on low-dimensional subspaces are $o_{\mathrm{p},\boldsymbol{\Theta}}(1)$. This is ensured by part (1) of Definition 2.

However, bounded noise alone is insufficient. If $\boldsymbol{u}^1 \cdot \underline{\boldsymbol{q}}^0$ is very small, there may be no hope for consistently recovering the true magnitude or sign of $\boldsymbol{u}^1 \cdot \underline{\boldsymbol{q}}^0$ from the signal $\widehat{\boldsymbol{u}}^1 \cdot \underline{\widehat{\boldsymbol{q}}}^0$. That is, even very small noise could still overwhelm a similarly small underlying mean projection $\boldsymbol{u}^1 \cdot \underline{\boldsymbol{q}}^0$. This unrecoverability would make it impossible to orient and scale our intervention appropriately, because the orientation procedure relies on such projections. Part (3) of the definition of significant structure prevents this problem by requiring that the projection of $\underline{\boldsymbol{q}}^0$ onto eigenvectors with large eigenvalues is bounded away from zero.

This discussion illustrates the following special case of our main result: If the largest eigenvalue of $\underline{\boldsymbol{D}}$ is well-separated from others and if $\underline{\boldsymbol{q}}^0 \cdot \boldsymbol{u}^1$ is not vanishingly small, then a subsidy profile proportional to $\boldsymbol{u}^1$ can, if it is suitably scaled, achieve all the properties of Theorem 1.

6.1.3. *Intervention rules based on recovered eigenvectors: General case.* The proof of the main result improves on this sketch in two ways. First, it does not rely only on the eigenspace spanned by $\boldsymbol{u}^1$. Instead, it uses a potentially much larger eigenspace of $\underline{\widehat{\boldsymbol{D}}}$. The general intervention projects $\underline{\widehat{\boldsymbol{q}}}^0$ onto $\mathcal{L}(\underline{\widehat{\boldsymbol{D}}}, b(n))$, the eigenspace of all eigenvectors of the observed $\underline{\widehat{\boldsymbol{D}}}$ with eigenvalues whose absolute value is at least $b(n)$. This makes it easier for the analog of $\widehat{\boldsymbol{u}}^1 \cdot \underline{\boldsymbol{q}}^0$ not to be too small, since the projection of $\underline{\boldsymbol{q}}^0$ onto a larger eigenspace will have a larger norm. Second, the general proof dispenses with assuming that any eigenvalues are well-separated. Instead, it handles any possible spectrum of $\underline{\boldsymbol{D}}$ subject to our maintained assumptions. This introduces considerable complexity, as it is no longer generally possible to recover any true eigenvector $\boldsymbol{u}^1$ with any accuracy. We instead work directly with a recovered eigenspace that generalizes the span of $\widehat{\boldsymbol{u}}^1$. We show that despite limited knowledge of individual true eigenvectors of $\underline{\boldsymbol{D}}$ underlying this space, we can use the fact that all of them have large eigenvalues to generalize our argument for showing that surplus expressions in (10) can be made very close and nonzero with a feasible intervention. This is where the arguments go beyond standard applications of the Davis–Kahan theorem.

7. Illustrating recoverability in a two-block example

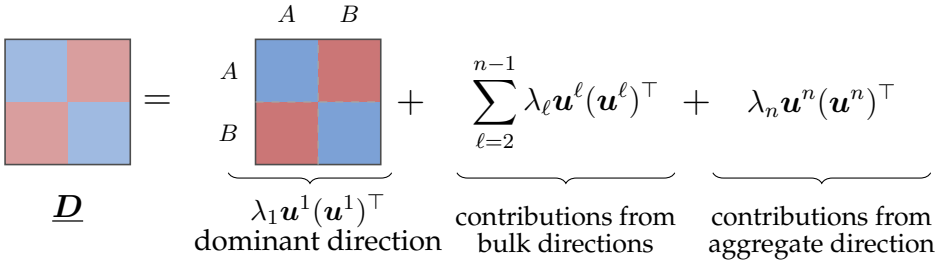
We develop a Monte Carlo experiment designed to demonstrate the central mechanism behind Theorem 1: significant structure enables recovery of welfare-relevant information from noisy observations of the demand system. Technical details supporting the simulation are in Appendix C.

7.1. **Demand system and noise in the experiment.** This example uses a block model in which the classification of goods into categories is unknown. The authority knows that there are two categories, that each category includes half of the products, and that within-category products are complements, while across-category products are substitutes. However, the authority does not know which products belong to the same category.

A concrete interpretation is a pair of competing product ecosystems: within an ecosystem, products such as a console, controllers, games, and subscription services are consumed together, while lower prices in one ecosystem draw demand away from the rival ecosystem. The two-block specification captures this reduced-form pattern by treating goods within an ecosystem as complements and goods across ecosystems as substitutes.

The sign pattern of substitutes and complements is different from the example presented in Section 4, where we had substitution within small categories and complementarity across categories. The configuration we use in this section makes it possible for the significant structure to be encoded in a rank-one matrix, and this makes the analysis more transparent.[26] We introduce a tunable parameter b governing the intensity with which this component is represented (and this parameter corresponds to the largest eigenvalue of the Slutsky matrix). Theorem 1 implies that when this eigenvalue is large enough, the authority can robustly increase total market surplus at the best rate possible. Our simulations will demonstrate this. They will also show that when this eigenvalue is small, the interventions presented here fail to achieve this goal, which provides intuition for the negative results.

Constructing a demand system. For each market state $\boldsymbol{\theta}$ we construct a normalized Slutsky matrix $\underline{\boldsymbol{D}}$ which has a block structure corresponding to this partition:



where:

Dominant direction (large eigenvalue). A market state $\boldsymbol{\theta}$ determines a partition of the products into two sets, A and B, of the same size. Given $\boldsymbol{\theta}$, the A-entries of $\boldsymbol{u}^1$ are $1/\sqrt{n}$ and the B-entries are $-1/\sqrt{n}$. We associate with $\boldsymbol{u}^1$ an eigenvalue of $\lambda_1 = -(1+b)$, where $b \geq 0$ is a parameter we vary. The "main" part of $\underline{\boldsymbol{D}}$ is a rank-one component $\lambda_1 \boldsymbol{u}^1(\boldsymbol{u}^1)^\top$ that creates a block structure in the Slutsky matrix, with a product's block membership determined by the sign of the corresponding entry of $\boldsymbol{u}^1$. For goods i and j in the *same* group (same sign in $\boldsymbol{u}^1$), we have $u_i^1 u_j^1 > 0$, so $\lambda_1 u_i^1 u_j^1 < 0$ (since $\lambda_1 < 0$). This means same-group goods

[26] When we have within-category substitutes, this is not possible because the sign pattern of a matrix $-\lambda \boldsymbol{u}\boldsymbol{u}^\mathsf{T}$ necessarily has negative entries (complements) on the diagonal blocks and positive entries (substitutes) on the off-diagonal blocks.

are complements. For goods in different groups (opposite signs), $u_i^1 u_j^1 < 0$, so $\lambda_1 u_i^1 u_j^1 > 0$—different-group goods are substitutes.

Aggregate direction (zero eigenvalue). The eigenvector $\boldsymbol{u}^n = \frac{1}{\sqrt{n}}\mathbf{1}$ represents uniform changes across all goods. We assign it eigenvalue $\lambda_n = 0$.

Bulk directions. We now turn to most of the eigenvectors—the so-called bulk. These $n-2$ eigenvectors receive eigenvalues chosen so that they are close to -1 and $\sum_\ell \lambda_\ell = -n$. Specifically, $\lambda_\ell = -1 + \frac{b-1}{n-2}$ for $\ell \notin \{1, n\}$. This ensures that the diagonal $\underline{D}_{ii} = -1$ is constant, as required of the normalized Slutsky matrix.

Status quo quantities. The normalized status-quo quantity vector is $\underline{\boldsymbol{q}}^0 = \boldsymbol{u}^n + 0.05\,\boldsymbol{u}^1$. Hence, $\underline{\boldsymbol{q}}^0$ projects primarily onto the aggregate direction $\boldsymbol{u}^n$ with a smaller projection onto the dominant direction $\boldsymbol{u}^1$.

Observation noise and the recovery threshold. In this simulation, we assume the authority observes $\underline{\boldsymbol{q}}^0$ without error, and observes $\widehat{\underline{\boldsymbol{D}}} = \underline{\boldsymbol{D}} + \boldsymbol{E}$, where $\boldsymbol{E}$ is a symmetric random matrix with $E_{ij} \sim \mathcal{N}(0,1)$ i.i.d. for $i < j$, $E_{ij} = E_{ji}$, and $E_{ii} = 0$, in line with Assumption 4. This is a Wigner matrix (with zeros on the diagonal), and classical results give $\|\boldsymbol{E}\| \approx 2\sqrt{n}$ with high probability. Thus, the uncertainty about $\boldsymbol{\theta}$ is which dominant direction $\boldsymbol{u}^1$ (equivalently, which A/B partition) obtains. Each Monte Carlo repetition corresponds to a fresh draw of this signal holding $\boldsymbol{\theta}$ fixed.

The significant structure condition (Definition 2) requires $|\lambda_1| \gg \|\boldsymbol{E}\|$, which is equivalent to $b \gg \sqrt{n}$. In this simulation, the recovery pattern changes sharply when $b/\sqrt{n}$ is near 1: below that range, the dominant eigenvalue is comparable to the noise and the leading eigenvector of $\widehat{\underline{\boldsymbol{D}}}$ is largely noise-driven, while for $b/\sqrt{n} \gg 1$ the leading eigenvector $\widehat{\boldsymbol{u}}$ aligns closely with the true dominant direction $\boldsymbol{u}^1$. The Davis–Kahan theorem is the general version of this last part and is used in the main positive result of the paper: recovery improves once the spike is sufficiently large relative to $\|\boldsymbol{E}\|$.

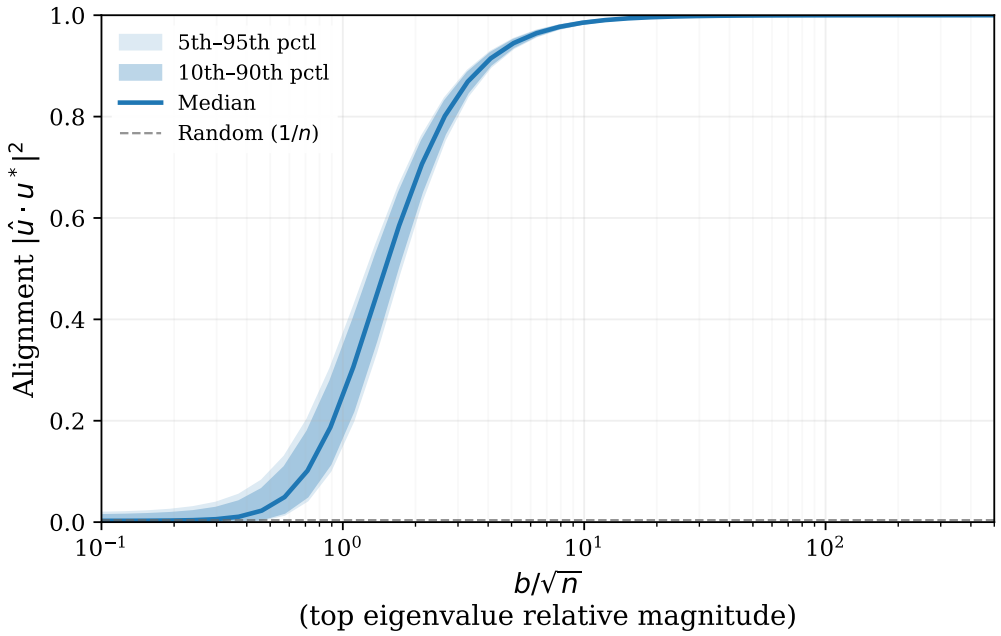
7.2. **The intervention rule.** Let $\widehat{\boldsymbol{u}}$ be the eigenvector of $\widehat{\underline{\boldsymbol{D}}}$ with the largest eigenvalue in absolute value. The authority implements the intervention

$$\underline{\boldsymbol{\sigma}} = \frac{s \cdot \widehat{\boldsymbol{u}}}{\widehat{\boldsymbol{u}} \cdot \underline{\widehat{\boldsymbol{q}}}^0},$$

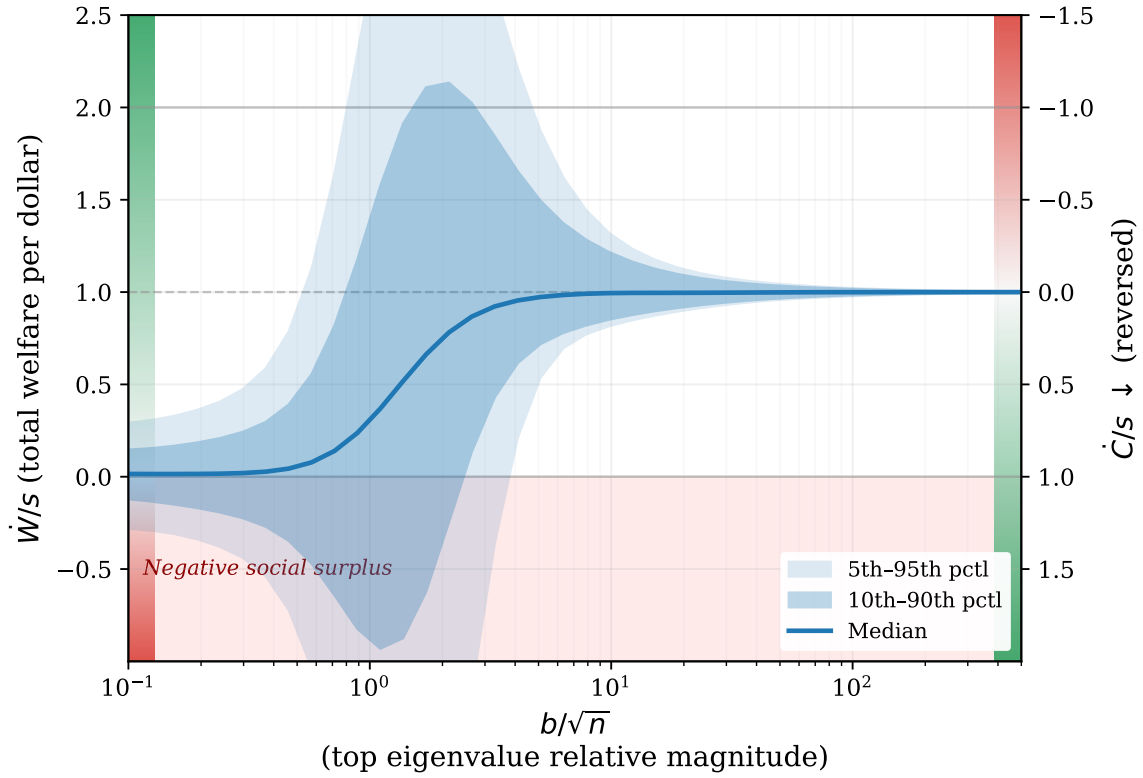
noting that this guarantees spending $\dot{S} = \underline{\boldsymbol{q}}^0 \cdot \underline{\boldsymbol{\sigma}}$ is positive; assume the authority scales the intervention to achieve target spending $s = 1$.

7.3. **Results.** We fix $n = 256$ and we vary the size of the dominant eigenvalue λ_1 by changing b. At each value of $b/\sqrt{n}$ we run 500 Monte Carlo draws of the noise matrix $\boldsymbol{E}$ and implement the intervention $\underline{\boldsymbol{\sigma}}$. For each draw we compute total surplus outcomes and the alignment of the dominant eigenvector $\widehat{\boldsymbol{u}}$ with the dominant eigenvector $\boldsymbol{u}^1$ measured by $|\boldsymbol{u}^1 \cdot \widehat{\boldsymbol{u}}|^2$. We report median outcomes and confidence intervals. Figure 3 summarizes the results.

Panel (A): Eigenvector recovery. This panel illustrates the sharp recovery threshold in the present simulation. Below the critical signal-to-noise ratio threshold of 1 ($b/\sqrt{n} < 1$), the median alignment of $\widehat{\boldsymbol{u}}$ and $\boldsymbol{u}^1$ is of order $1/n$, which is the value expected for two random unit vectors. Above the threshold, alignment



(A) Eigenvector recovery



(B) Total Surplus outcomes

FIGURE 3. **Recovery of $\underline{D}$ and welfare outcomes** ($n = 256$ and 500 draws per value of b). **(a)** Alignment $|\widehat{\boldsymbol{u}} \cdot \boldsymbol{u}^1|^2$ between recovered and true spike eigenvectors. Shaded bands: 5th–95th and 10th–90th percentiles; line: median. **(b)** Total surplus per dollar $\dot{W}/\dot{S}$ (left axis) and consumer surplus per dollar $\dot{C}/\dot{S}$ (reversed right axis). The shaded region marks negative social surplus.

rises rapidly toward the case of perfect alignment of 1, with narrowing bands indicating increasingly reliable recovery.

Panel (B): Welfare outcomes. Consider the spectral decomposition of total surplus:

$$\dot{W} = (\boldsymbol{u}^1 \cdot \underline{\boldsymbol{q}}^0)(\boldsymbol{u}^1 \cdot \underline{\boldsymbol{\sigma}}) \frac{|\lambda_1|}{1+|\lambda_1|} + \sum_{\ell=2}^{n-1} (\boldsymbol{u}^\ell \cdot \underline{\boldsymbol{q}}^0)(\boldsymbol{u}^\ell \cdot \underline{\boldsymbol{\sigma}}) \frac{|\lambda_\ell|}{1+|\lambda_\ell|} + (\boldsymbol{u}^n \cdot \underline{\boldsymbol{q}}^0)(\boldsymbol{u}^n \cdot \underline{\boldsymbol{\sigma}}) \frac{|\lambda_n|}{1+|\lambda_n|}.$$

Below the critical threshold, median total surplus per dollar $\dot{W}/\dot{S}$ is close to zero. The wide bands indicate high variability: some noise realizations happen to align with $\boldsymbol{u}^1$ by chance, yielding positive welfare, while others yield negative total surplus—outcomes in the shaded region represent net welfare destruction.

The reason is that lack of recoverability of $\boldsymbol{u}^1$ implies that $\underline{\boldsymbol{\sigma}}$ is based on a random direction and, therefore, the projection $\boldsymbol{u}^\ell \cdot \underline{\boldsymbol{\sigma}}$ is roughly equal across all ℓ. But the quantity vector $\underline{\boldsymbol{q}}^0$ loads almost entirely on $\boldsymbol{u}^n$, and $\boldsymbol{u}^n$ has pass-through weight $|\lambda_n|/(1+|\lambda_n|) = 0$ in our construction. Thus, failed recovery systematically traps spending in a low-pass-through direction and, so, expected total surplus approaches zero.

Above the threshold, $\dot{W}/\dot{S}$ rises quickly toward the theoretical maximum and the bands narrow dramatically. The intervention reliably delivers close to the optimal total surplus per dollar. As b increases, $\dot{W}/\dot{S}$ approaches 1: nearly one dollar of net total surplus per dollar of spending. This is equivalent to a 2-for-1 gain in gross total surplus $\dot{C}+\dot{P}$. This is why the significant structure condition matters: it ensures the authority can identify and target the high-pass-through directions that yield substantial welfare gains.

Finally, panel (B) also shows consumer surplus per dollar $\dot{C}/\dot{S}$ (right axis, reversed). Above threshold, $\dot{C}/\dot{S}$ approaches zero while $\dot{W}/\dot{S}$ approaches 1. This means nearly all the welfare gain accrues to producers ($\dot{P}/\dot{S} \approx 2$), confirming the prediction of Theorem 1: interventions targeting high-pass-through directions achieve maximum total surplus but allocate almost all the surplus increase to the supply side.

7.4. **What does successful recovery mean economically?** Figure 4 shows the ground-truth block structure $\lambda_1 \boldsymbol{u}^1 (\boldsymbol{u}^1)^\top$. Figure 5 shows the authority's recovered rank-one component $\widehat{\lambda}\widehat{\boldsymbol{u}}(\widehat{\boldsymbol{u}})^\top$ of $\underline{\widehat{\boldsymbol{D}}}$ for three signal-to-noise ratios. We use the same noise realization $\boldsymbol{E}$ across panels, varying only the value of the dominant eigenvalue λ_1.

At threshold ($b/\sqrt{n} = 1$), the estimated structure is weakly related to the true structure—the authority cannot tell which goods are substitutes versus complements. Above threshold, the two-block pattern emerges with increasing clarity. At $b/\sqrt{n} = 5$, the recovered structure closely matches the ground truth.

This visualization makes concrete what "significant structure" means for intervention design. The authority does not need to estimate every entry of $\underline{\boldsymbol{D}}$ accurately. Instead, it needs to recover the *global* pattern—which families of goods interact as substitutes versus complements—that determines where welfare-improving interventions are possible. When this pattern is sufficiently

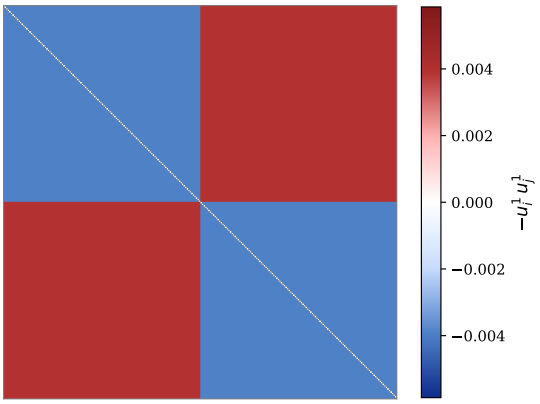


FIGURE 4. **Ground-truth block structure** ($n = 256$). The rank-one component $\lambda_1 \boldsymbol{u}^1(\boldsymbol{u}^1)^\top$ corresponding to the dominant eigenvalue of the true normalized Slutsky matrix $\underline{\boldsymbol{D}}$, with goods ordered by sign in $\boldsymbol{u}^1$. Blue: complements (same group); red: substitutes (different groups).

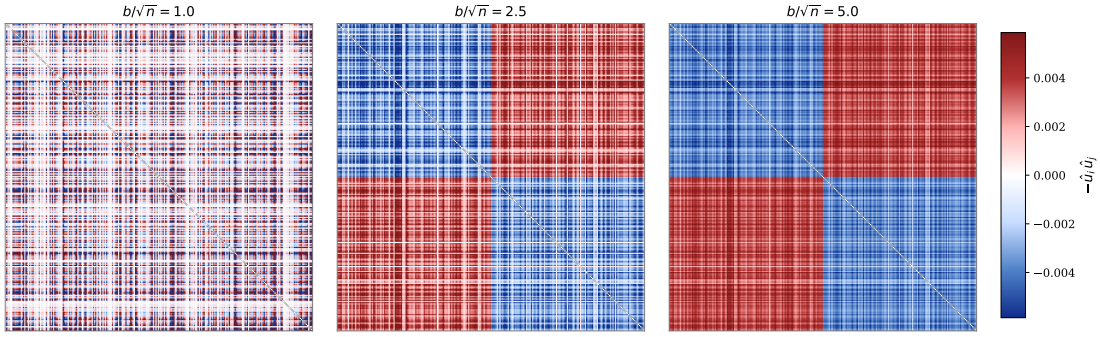


FIGURE 5. **Recovery of substitute/complement structure** ($n = 256$, same noise realization). Leading rank-one component $\widehat{\lambda}\widehat{\boldsymbol{u}}\widehat{\boldsymbol{u}}^\top$ of $\underline{\widehat{\boldsymbol{D}}}$, with goods ordered by true $\boldsymbol{u}^1$ signs. **Left** ($b/\sqrt{n} = 1$): Pure noise, no discernible structure. **Center** ($b/\sqrt{n} = 2.5$): Block structure emerges but with significant noise. **Right** ($b/\sqrt{n} = 5$): Near-perfect recovery of the two-block pattern.

pronounced relative to noise (significant structure), eigenvector-based methods extract it reliably.

8. DISCUSSION AND CONCLUDING REMARKS

We conclude with some observations about the scope of our analysis. Section 8.1 shows the limits of robust interventions. Section 8.2 discusses the theoretical foundation of significant structure and develops an empirical test for it. Section 8.3 formally maps our leading example, the block model of demand in Section 5.2.1, to a representative consumer with quadratic utilities. Section 8.4 discusses local linear demand. Section 8.5 relates our work to the literature on network games.

8.1. **The limits of robust interventions.** We have established that if demand has significant 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 has a precise distributional property as it boosts production with minimal price changes, resulting in firms capturing all efficiency gains.

The key condition of significant structure ensures sufficiently strong signal-to-noise ratio in the signal about the Slutsky matrix. This raises two questions. First, when this property is violated, so that the signal is weak relative to observation noise, can we find interventions that robustly increase total surplus? Second, under the assumption of significant structure, are the distributional properties of the robust interventions in our main result necessary for any robust intervention?

Proposition 3 (Limits of robust interventions)**.** Consider interventions with a positive expenditure by the authority, $\dot{S} > 0$.

(1) There exist environments satisfying Assumptions 1–4, but without significant structure, such that for any $\varepsilon \in (0, 1/2)$ there is no intervention rule that ε-robustly achieves $\dot{C} + \dot{P} \geq (1 + \varepsilon)\dot{S}$.
(2) There exist environments satisfying Assumptions 1–4, with significant structure, such that, for any $\kappa > 0$, no intervention rule robustly achieves $\dot{C} \geq \kappa \dot{S}$.

Whereas with significant structure we could construct an intervention returning almost two dollars in market surplus for each dollar invested, part (1) says that without significant structure, we may not be able to achieve any non-negligible return on the authority's spending. And part (2) says that guaranteeing any fixed positive share of spending accrues to consumer surplus is impossible even with significant structure.

Both parts are proved in Appendix B using Hadamard-basis constructions closely related to the two-block example in Section 7. In part (1), the market state is defined by a two-block classification of goods: a partition of the n products into two equal-sized subsets determining which pairs behave as complements versus substitutes. We use a special type of basis (a Hadamard basis) to index a large collection of such balanced classifications. Each basis vector has entries $\pm 1/\sqrt{n}$, with the signs corresponding to the assignment of goods to blocks.[27]

For part (1), we construct $\boldsymbol{D}$ so that the vector corresponding to the block partition that is realized has a large eigenvalue; however, all realizations of $\boldsymbol{D}$ are designed to generate the same signal. Any deterministic rule must therefore select the same intervention across these states. We prove that for any such intervention, there exists a state in which it is nearly orthogonal to the welfare-relevant spike direction, so $\dot{W}$ is necessarily small. This is akin to the phenomenon illustrated in Section 7 below the recovery threshold, where the leading eigendirection of $\widehat{\underline{\boldsymbol{D}}}$ is dominated by noise and eigenvector-based interventions

[27]The special feature of the Hadamard basis is that there are n such vectors that are all orthogonal to one another.

behave like random directions (see Figure 3a and Figure 3b for $b/\sqrt{n} < 1$). The proof here shows that, not only does the specific type of intervention studied there fail, but any possible intervention fails. In our terminology, this is an environment without significant structure: although the true demand system features a welfare-relevant two-block pattern, the signal is too weak relative to observation noise to recover which pattern obtains.[28]

For part (2), we instead ensure significant structure by making the "major" eigenspace fixed and strongly detectable through noise, so it is recoverable and high-pass-through welfare gains remain available. The key observation is that *consumer* surplus pass-through behaves very differently from producer-surplus pass-through (Lemmas 1 and 2). Along the large-eigenvalue directions that drive $\dot{W}$ (Equation (10)), prices barely move, so $\dot{C}$ is negligible. To obtain a fixed positive consumer-surplus incidence per dollar of expenditure uniformly, a rule must load on directions with smaller eigenvalues, where price pass-through is order one; we hide one such direction in the quantity vector while making the relevant quantity signals nearly indistinguishable, and with a proof similar to that of part (1) show that the desired lower bound on $\dot{C}/\dot{S}$ must fail in some state.

We now note a negative result that focuses on expected performance, which is often a focus of the frequentist statistical decision theory literature (Manski, 2004).

Remark 2. The construction used for Proposition 3(1) also yields the following implication for expected performance.

After normalizing any positive-expenditure rule so that $\dot{S}_{\boldsymbol{R}(\widehat{\boldsymbol{\theta}})} = 1$,

$$\inf_{\boldsymbol{\theta}\in\boldsymbol{\Theta}(n)} \mathbb{E}_{\varphi_{\boldsymbol{\theta}}(n)}\left[\dot{C} + \dot{P} - \dot{S}\right] \leq \varepsilon.$$

A proof is given in Appendix B.3.

8.2. **Theoretical foundations and empirical tests for significant structure.** Significant structure is a joint property concerning the norm of estimation errors in market quantities and $\boldsymbol{D}$, and the operator norm of $\boldsymbol{D}$. We now discuss foundations for and tests of the key parts of this assumption.

8.2.1. *Theoretical foundation.* In Section 4 we provide a block demand structure under which the norm of $\boldsymbol{D}$ grows linearly in n. In this class of demand environments, whenever the norm of the estimation errors grows in n at a slower rate, we obtain recoverability of eigenvectors associated with large eigenvalues. The results raise the question of whether estimation procedures that an authority can carry out are consistent with this requirement.

In Appendix D we formalize a natural estimation procedure under which $\|\boldsymbol{E}\| \in O_{\mathrm{p}}(\sqrt{n})$. The authority samples a distinct household for each product pair (i, j), and performs a demand experiment to obtain *independent, unbiased* estimates of how the demands for both products change when p_i is perturbed. For each product i, we obtain n distinct estimates of the own-price effect, and so we can recover the own-price demand derivative (D_{ii} for each i) by the law of

[28] Formally, the demand noise is of the same order as the eigenvalue scale that distinguishes the blocks, so the associated major eigendirections are not stably identified uniformly over states.

large numbers; thus we can satisfy Assumption 4. Turning to off-diagonal entries, standard random matrix theory implies that, under suitable assumptions, the operator norm of the estimation error matrix is $O_{\mathrm{p}}(\sqrt{n})$.

Estimation errors on demand derivatives may exhibit spatial or structural correlations: for example, errors in estimating cross-price elasticities between nearby product categories may be positively correlated. Recent advances in random matrix theory show that the key consequence we require—namely, that the spectral norm $\|\boldsymbol{E}\|$ is $o(n)$—holds under far more general dependence structures.[29]

8.2.2. *A test for significant structure.* In applications, significant structure is a substantive joint condition on (i) the scale of recoverable eigenvalues of the normalized Slutsky matrix and (ii) the magnitude of observation error. Since this condition is not directly observable, it is useful to have an empirical diagnostic. The proof of our main result implies a statistical test that can be used to assess whether significant structure is present.

The test works as follows: Take one estimate of the market state and construct the intervention recommended in our positive result. Then take an independent estimate of the market state, and *treating this estimate as if it were exactly the true market state*, assess the net total surplus effect from using the intervention constructed based on the first estimate. If significant structure is present, the predicted net total surplus increase per dollar of predicted spending should be close to 1.

The reason for this is that under significant structure, the information on the underlying Slutsky matrix $\boldsymbol{D}$ will accurately reveal eigenspaces corresponding to large eigenvalues. Along eigenspaces with large eigenvalues, the total surplus pass-through behavior of interventions is the same for the true matrix and for any perturbed version $\boldsymbol{D}+\boldsymbol{E}$ where $\boldsymbol{E}$ is drawn like our error matrices; that is what makes the recovery procedure so robust. Thus, an intervention that is chosen to have a good effect for one sample should perform well in another.

More formally: Suppose the analyst has $T \geq 4$ independent estimates of the market state (e.g., from disjoint subsamples); call them $(\widehat{\boldsymbol{D}}^{(t)}, \widehat{\boldsymbol{q}}^{0,(t)})_t$. Let $\widehat{\underline{\boldsymbol{\sigma}}}^{(t)}$ denote the intervention produced by our rule (see Section 7.2) from the estimate t, with target spending level $s = 1$:

$$\widehat{\underline{\boldsymbol{\sigma}}}^{(t)} = \frac{P_{\mathcal{L}(\widehat{\underline{\boldsymbol{D}}}^{(t)}, b(n))}\, \widehat{\underline{\boldsymbol{q}}}^{0,(t)}}{\left\| P_{\mathcal{L}(\widehat{\underline{\boldsymbol{D}}}^{(t)}, b(n))}\, \widehat{\underline{\boldsymbol{q}}}^{0,(t)} \right\|^2}$$

For two splits $t \neq t'$, the predicted welfare effect of the intervention designed on split t, evaluated at the state estimate from split t', is:

$$\widehat{\tau}^{(t \to t')} := \sum_{\ell=1}^{n} \left(\widehat{\boldsymbol{u}}^{\ell,(t')} \cdot \widehat{\underline{\boldsymbol{q}}}^{0,(t')}\right)\left(\widehat{\boldsymbol{u}}^{\ell,(t')} \cdot \widehat{\underline{\boldsymbol{\sigma}}}^{(t)}\right) \frac{|\widehat{\lambda}_\ell^{(t')}|}{1 + |\widehat{\lambda}_\ell^{(t')}|}, \tag{11}$$

[29] When correlations between entries decay polynomially with distance under a suitable metric, or when entries exhibit finite-range dependence (i.e., E_{ij} and E_{kl} are independent whenever indices are sufficiently far apart), the spectral norm of a symmetric $n \times n$ error matrix $\boldsymbol{E}$ satisfies $\|\boldsymbol{E}\| = O_{\mathrm{p}}(\sqrt{n})$; see Erdős et al. (2019) for slowly decaying correlations, Anderson and Zeitouni (2008) for finite-range dependence, and Reker (2022) for polynomially decaying metric correlations. These results provide the operator-norm control needed for condition (2) of Definition 2.

where $(\widehat{\lambda}_\ell^{(t')}, \widehat{\boldsymbol{u}}^{\ell,(t')})$ is the eigendecomposition of $\underline{\widehat{\boldsymbol{D}}}^{(t')}$ and the normalization is as in Section 2.7. We pair the T estimates arbitrarily. For simplicity, assume T is even, writing $T = 2Z$, and for each pair $z = 1, \ldots, Z$, define the pair statistic

$$\widehat{\tau}^{(z)} := \frac{1}{2}\Big(\widehat{\tau}^{(2z-1\to 2z)} + \widehat{\tau}^{(2z\to 2z-1)}\Big), \qquad \overline{\tau} := \frac{1}{Z}\sum_{z=1}^{Z}\widehat{\tau}^{(z)}.$$

Here the arrow notation $t \to t'$ means: construct the intervention using split t, but evaluate its predicted welfare effect using the independently estimated state from split t'. The statistic $\widehat{\tau}^{(z)}$ symmetrizes this cross-evaluation within each paired split.

We now formalize the prediction that if the true state $\boldsymbol{\theta}$ belongs to an environment with significant structure then $\widehat{\tau}^{(z)}$ is near one.

Let $\mu_n(\boldsymbol{\theta}) := \mathbb{E}[\widehat{\tau}^{(1)} \mid \boldsymbol{\theta}]$ denote the (conditional) population mean of the paired statistic when there are n products. We form a one-sided $(1-\alpha)$ lower confidence bound for $\mu_n(\boldsymbol{\theta})$ via the nonparametric bootstrap at the level of pairs: resample $\{\widehat{\tau}^{(z)}\}_{z=1}^{Z}$ with replacement many times (indexing these bootstrap pairings $r = 1, 2, \ldots, R$) compute the resampled means $\overline{\tau}^{(r)}$, and let $L_{1-\alpha}$ be the empirical α-quantile of $\overline{\tau}^{(r)}$. We say significant structure is not certified at tolerance ϵ whenever $L_{1-\alpha} < 1 - \epsilon$. Intuitively, a bound well below one indicates that predicted total surplus falls materially short of predicted expenditure in the resampling exercise. Appendix D.2 formalizes the coverage guarantee for this diagnostic, proving that certification occurs with high probability when significant structure is indeed present.

Illustration based on Section 7. We consider the environment constructed in Section 7 and consider a variety of values of $b/\sqrt{n}$. We generate $T = 500$ independent noisy estimates of demand and quantities, and for each bootstrap draw r, we pair them up and then compute the $\overline{\tau}^r$ outcome for that pairing. We repeat this many times with different pairings, take the 5^{th} percentile of the resulting empirical distribution (this is the curve in Figure 6), and reject when that number is below $1 - \epsilon$.

Thus, the test is quite demanding in terms of how the intervention has to perform when an out-of-sample market state is taken as ground truth. And yet our result in Appendix D.2 shows that when significant structure is present and "strong" in the sense of a large $b/\sqrt{n}$ ratio, it is very likely that the intervention does indeed perform very well out of sample, so that the test statistic falls above the line.

8.3. **Linear–quadratic utility and relationship with hedonic utility models.**

8.3.1. *Linear–quadratic utility.* Our leading example, the block model of demand in Section 5.2.1, is consistent with optimal behavior of a representative consumer with a standard quasilinear quadratic utility function (as in Singh and Vives (1984)). Suppose there are K product categories with m products in each category. The representative consumer's utility is:

$$U(\boldsymbol{q}) = \boldsymbol{a}^\mathsf{T}\boldsymbol{q} - \frac{1}{2}\boldsymbol{q}^\mathsf{T}\boldsymbol{H}\boldsymbol{q},$$

where $\boldsymbol{a} \in \mathbb{R}^n_+$, $\boldsymbol{H}$ is an $n \times n$ symmetric matrix with $H_{ii} = z > 0$, $H_{ij} = z_{\text{In}} > 0$ if $i \neq j$ belong to the same category and, otherwise, $H_{ij} = z_{\text{Ex}} < 0$. Hence, the

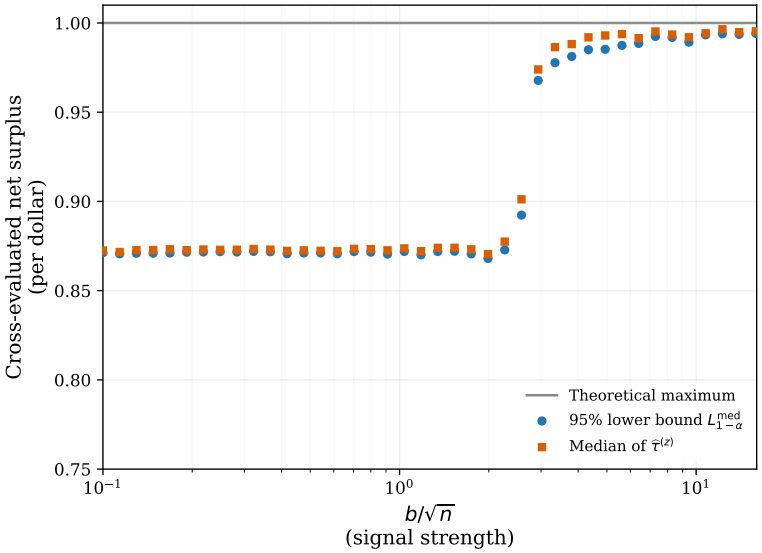


FIGURE 6. Bootstrap lower confidence bound for cross-evaluated total surplus. For each value of $b/\sqrt{n}$, the line gives the 5^{th} percentile of the empirical distribution of test statistics across the bootstrap experiments. At low signal strength this is well below one. At high signal strength this value approaches one, consistent with significant structure. At tolerance ϵ, values below $1-\epsilon$ indicate a failure to certify significant structure according to this diagnostic.

marginal utility of product i's consumption is decreasing in the consumption of products within the same category and increasing in the consumption of products in different categories. Given prices $\boldsymbol{p}$, the first-order condition of the maximization of $U - \boldsymbol{p} \cdot \boldsymbol{q}$ with respect to $\boldsymbol{q}$ provides the demand function:

$$\boldsymbol{q}(\boldsymbol{p}) = \boldsymbol{H}^{-1}(\boldsymbol{a} - \boldsymbol{p}),$$

and, differentiating it with respect to prices yields the Jacobian:

$$\boldsymbol{D} = \frac{\partial \boldsymbol{q}(\boldsymbol{p})}{\partial \boldsymbol{p}} = -\boldsymbol{H}^{-1}. \tag{12}$$

For each market state $\boldsymbol{\theta}$, we wish to find $(z^*(\boldsymbol{\theta}), z^*_{\text{In}}(\boldsymbol{\theta}), z^*_{\text{Ex}}(\boldsymbol{\theta}))$ so that the $\boldsymbol{D}$ in (12) corresponds to the Slutsky matrix in Section 5.2.1, which is equivalent to solving a system of three equations fully specified in Appendix E.

As we increase the number of product categories K we want the largest (in absolute value) eigenvalue of $\boldsymbol{D}$ grows linearly in K and that $\boldsymbol{D}$ is negative definite. This requires that, as we increase K, the following holds:

$$\lim_{K\to\infty} z^*_{\text{Ex}}(\boldsymbol{\theta})m(K-1) + (z^*(\boldsymbol{\theta}) + z^*_{\text{In}}(\boldsymbol{\theta})(m-1)) = 0^+. \tag{13}$$

The left-hand side is the corresponding eigenvalue of $\boldsymbol{H}$ in the all-ones direction (equivalently, the curvature of U in that direction is its negative). Hence, (13) requires that, as we expand product categories, this eigenvalue remains positive (preserving concavity) but approaches zero. Intuitively, this captures that increases of consumption in this direction remain valuable, corresponding to the idea of large-scale complementarities.

One might have thought that a near-singular Hessian causes some kind of instability or bad behavior of the equilibrium around such a point. Our results show that this is not the case, due to the way the Hessian maps into $\boldsymbol{D}$ and then into the system of equations governing the equilibrium. When all is said and done, cost perturbations in any direction have bounded pass-through to prices and quantities.

Finally, we note that equilibrium quantities and prices are well-behaved in an economy defined by (13). To see this, consider the case where $a_i = a$ and $c_i = c$ for all i; this ensures that the equilibrium is symmetric and takes the following form:

$$p = \frac{c}{1 + |\lambda_1(\boldsymbol{\theta})|} + \frac{|\lambda_1(\boldsymbol{\theta})|a}{1 + |\lambda_1(\boldsymbol{\theta})|} \text{ and } q = \frac{|\lambda_1(\boldsymbol{\theta})|}{1 + |\lambda_1(\boldsymbol{\theta})|}(a - c),$$

where $\lambda_1(\boldsymbol{\theta})$ is the dominant eigenvalue of the Slutsky matrix defined in Section 5.2.1. Note that equilibrium price and quantity depend on K only via $|\lambda_1(\boldsymbol{\theta})|$. As K goes to infinity $|\lambda_1(\boldsymbol{\theta})| \to \infty$ and $p \to a$ and $q \to (a - c)$. This is consistent with Condition 1(a)(2) imposed on quantity in the illustrative example of Section 4, since any uniform vector satisfies it after normalization.

8.3.2. *Relationship to hedonic models.* Recent work by Pellegrino (2021) and Ederer and Pellegrino (2021) uses a canonical oligopoly model to empirically quantify the evolution of market power. In particular, 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.

As in the linear–quadratic demand above, the Slutsky matrix $\boldsymbol{D}$ derived from the hedonic model can be expressed as (12). In this case, $\boldsymbol{H}$ is a convex combination of the identity matrix and the cosine similarity matrix of products' characteristics, which Hoberg and Phillips (2016) estimated for a large set of consumer goods using text data.[30] An implication of this is that all entries H_{ij} are positive. This means that the marginal utility of consuming product i is decreasing in product j's consumption, and this holds for all products (i, j). By assumption, the hedonic model does not allow for products that the consumer likes to consume together, e.g., a tennis racket and a tennis ball.

This turns out to impose a bound on $\|\underline{\boldsymbol{D}}\|$, and so significant structure cannot be realized, as we show in Appendix F.[31]

The intuition can be found in Section 5: blocks of substitutes do not create large eigenvalues. However, as that section also illustrated, significant structure will naturally emerge once we broaden our view of the market to allow for both complementary and substitute products. There is a straightforward economic reason for why such complementarities more readily produce significant structure: when one good's price decreases, all its complements can experience comparable nonvanishing increases in demand. This creates the clusters of nonvanishing entries in $\boldsymbol{D}$ that are the hallmark of significant structure.

In summary, direct complementarities seem practically important and can naturally yield the significant structure central to our results. We hope these

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

[31] We checked that using the data and the parameter values used by Pellegrino (2021), the leading eigenvalues of the Slutsky matrix are small.

observations will motivate further empirical research on the structure of large-scale oligopoly models with complementary goods.

8.4. **Nonlinear demand.** We have assumed that each good's demand is exactly linear in its own price in a neighborhood of the status quo equilibrium. This assumption has helped significantly with tractability; for example, it afforded a very clean characterization of the complete-information benchmark in Proposition 1.

This section clarifies the aspects of our analysis that extend with non-linear demand, and outlines aspects that need further development.

8.4.1. *Two conceptually distinct steps.* Our analysis rests on two results that play different roles:

(1) A *spectral decomposition of total surplus pass-through*: given any price perturbation $\dot{\boldsymbol{p}}$, the first-order total surplus effect depends only on the Slutsky matrix $\boldsymbol{D}$ and equilibrium markups at the status quo, via an interpretable spectral decomposition. This result is independent of the assumption that each good's demand is locally linear in its own price.
(2) A *mapping from subsidies to prices*: in the linear environment, equation (5) gives $\dot{\underline{\boldsymbol{p}}} = -(\boldsymbol{I} - \underline{\boldsymbol{D}})^{-1}\underline{\boldsymbol{\sigma}}$, so that the eigenbasis of $\underline{\boldsymbol{D}}$ diagonalizes both total surplus and the cost-to-price mapping. This result exploits the assumption that each good's demand is locally linear in its own price in a fundamental way.

The first step identifies which price movements are desirable; the second step shows how to implement them via per-unit subsidies. Local linearity is essential only for the second step.

8.4.2. *First-order welfare theory extends fully.* Because demand $\boldsymbol{q}(\boldsymbol{p})$ is differentiable, the first-order quantity response to any small price perturbation is $\dot{\boldsymbol{q}} = \boldsymbol{D}\dot{\boldsymbol{p}}$, without any linearity assumptions whatsoever. Indeed, the first-order welfare effect of a price perturbation is

$$\dot{W} = \sum_{\ell} \lambda_\ell \left[(\underline{\boldsymbol{p}}^0 - \underline{\boldsymbol{c}}^0) \cdot \boldsymbol{u}^\ell\right] \left(\dot{\underline{\boldsymbol{p}}} \cdot \boldsymbol{u}^\ell\right), \tag{14}$$

This is of interest in its own right. The expression shows that welfare gains depend on the alignment among three objects: the status-quo markup vector $(\underline{\boldsymbol{p}}^0 - \underline{\boldsymbol{c}}^0)$, the demand responsiveness encoded in $\underline{\boldsymbol{D}}$, and the direction of the price perturbation $\dot{\underline{\boldsymbol{p}}}$. The spectral form decomposes this into independent contributions from each eigendirection: a price movement in direction $\boldsymbol{u}^\ell$ contributes to total surplus in proportion to the eigenvalue λ_ℓ and the markup's loading on that direction.

Evaluating the Bertrand first-order condition at the status quo and normalizing gives

$$\underline{\boldsymbol{p}}^0 - \underline{\boldsymbol{c}}^0 = \underline{\boldsymbol{q}}^0, \tag{15}$$

so the total surplus formula becomes

$$\dot{W} = \underline{\boldsymbol{q}}^0 \cdot \underline{\boldsymbol{D}}\, \dot{\underline{\boldsymbol{p}}} = \sum_{\ell} \lambda_\ell \left(\underline{\boldsymbol{q}}^0 \cdot \boldsymbol{u}^\ell\right) \left(\dot{\underline{\boldsymbol{p}}} \cdot \boldsymbol{u}^\ell\right).$$

This is operationally advantageous: quantities are typically easier to measure than marginal costs, so the authority can evaluate potential price movements

using observable data. (Note, however, that the normalized quantities and prices require good estimates of $\partial_i q_i(\boldsymbol{p})/\partial p_i$ to perform the normalization—arguably still more accessible than precise measurements of markups.)

The spectral perspective—that large eigenvalue directions with substantial quantity loadings matter most for total surplus—is thus fully general and does not depend on local linearity.

8.4.3. *Price pass-through of cost changes does not extend directly.* The assumption that each good's demand is locally linear in its own price is used in our theory only in the mapping from subsidies to prices. Under this assumption, subsidies in the eigendirection $\boldsymbol{u}^\ell$ pass through exclusively to the price and quantity of that bundle, with coefficients that depend only on λ_ℓ (Lemma 1). This "diagonalization" is what allows us to design robust interventions by targeting cost interventions to the estimated top eigenspace.

When the assumption that each good's demand is locally linear in its own price fails, the price pass-through of a per-unit subsidy depends additionally on the relevant slice of the demand-curvature tensor (see, e.g., Miklos-Thal and Shaffer, 2021)

$$\left(\partial^2 q_i(\boldsymbol{p})/\partial p_i \partial p_k\right)_{i,k}.$$

The eigenbasis of $\boldsymbol{D}$ no longer diagonalizes the cost-to-price mapping, and interventions designed to move prices in a particular eigendirection may have significant effects in other directions.

8.4.4. *Paths forward.* Extending Theorem 1 to the case beyond the linear setting of Assumption 2 requires new techniques for one of two tasks: either (i) characterizing which cost perturbations achieve desired price movements when the price pass-through operator is more complicated, or (ii) designing alternative instruments that directly control prices.

For the first approach, the challenge is learning the mapping from cost perturbations to price movements. For example, under the most direct extension of the Bertrand equilibrium theory, predicting pass-through requires recovering both the Slutsky matrix and certain mixed partials in the Hessian of the demand function. Performing this recovery from noisy observations is a harder problem than the one our framework addresses.[32] There are many more numbers to recover, and a priori it is unclear how robust interventions should be generalized, since the linear algebra becomes more complicated. There is also the possibility that pass-through from costs to prices does not obey the equations of a Bertrand model.

A general approach here is to assume that prices move differentiably in costs but otherwise take the local mapping $\boldsymbol{L}$ in $\dot{\boldsymbol{p}} = \boldsymbol{L}\dot{\boldsymbol{c}}$ as a black box. The goal then is simply to recover nonzero vectors in the preimage of the subspace of price vectors that we know have predictable welfare effects. Here *compressed sensing* techniques are relevant—methods for recovering structured signals from limited measurements by exploiting sparsity or low-rank structure (see, e.g.,

[32] See, e.g., Miravete, Seim, and Thurk (2024) and Birchall, Mohapatra, and Verboven (2024) for recent work on some of the challenges introduced by uncertainty about the shape/curvature of demand in differentiated-product markets.

Foucart and Rauhut, 2013). We leave to future work the question of what assumptions on pass-through (whether with or without structural foundations) enable this type of robust intervention design.

For the second approach, one can sidestep the pass-through problem entirely by designing mechanisms that directly enforce target price perturbations, or incentivize them in a direct way (rather than by tailoring cost perturbations). Our spectral analysis of welfare effects remains useful for a planner who has power to control prices in this way: the Davis–Kahan analysis identifies which eigenspaces of $\widehat{\underline{D}}$ can be trusted. This allows the authority to know that targeting price movements in these directions yields reliable total surplus improvements. This is the most immediate way to use the ideas behind Theorem 1 in environments that do not satisfy Assumption 2.

8.5. **Games on networks.** We have studied robust interventions in an oligopoly pricing model, which can be seen as a network game with linear best replies (e.g., Ballester et al., 2006; Bramoullé et al., 2016) where the network summarizes the complement and substitute relationships among products. The methods we have developed can be extended to other settings where robust interventions have not been studied. For example, in a public goods setting, interventions would aim to realign private marginal returns with social marginal returns. The literature has developed tools to understand how to do this when the authority has precise information on the spillovers causing the underprovision of public goods.[33] However, we know little about designing 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 leveraged, including the symmetry and negative semidefiniteness of the strategic interactions matrix. So there are challenges to overcome in extending our results. We hope this paper stimulates research in these directions.

[33] See, for instance, Bramoullé et al. (2014) and Elliott and Golub (2019).

REFERENCES

AGUIAR, V. H. AND R. SERRANO (2017): "Slutsky matrix norms: The size, classification, and comparative statics of bounded rationality," *Journal of Economic Theory*, 172, 163–201.

ANDERSON, G. W. AND O. ZEITOUNI (2008): "A law of large numbers for finite-range dependent random matrices," *Communications on Pure and Applied Mathematics*, 61, 1118–1154.

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.

ATHEY, S. AND S. WAGER (2021): "Policy Learning With Observational Data," *Econometrica*, 89, 133–161.

AZAR, J. AND X. VIVES (2021): "General equilibrium oligopoly and ownership structure," *Econometrica*, 89, 999–1048.

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.

BANDEIRA, A. S. AND R. VAN HANDEL (2016): "Sharp nonasymptotic bounds on the norm of random matrices with independent entries," *Annals of Probability*, 44, 2479–2506.

BANERJEE, A., A. G. CHANDRASEKHAR, E. DUFLO, AND M. O. JACKSON (2013): "The diffusion of microfinance," *Science*, 341, 1236498.

BAQAEE, D. R. AND E. FARHI (2020): "Productivity and misallocation in general equilibrium," *The Quarterly Journal of Economics*, 135, 105–163.

BELHAJ, M. AND F. DEROIAN (2017): "Contracting on Networks," Mimeo.

BIRCHALL, C., D. MOHAPATRA, AND F. VERBOVEN (2024): "Estimating Substitution Patterns and Demand Curvature in Discrete-Choice Models of Product Differentiation," *Review of Economics and Statistics*.

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., A. GALEOTTI, AND B. ROGERS (2016): *The Oxford handbook of the economics of networks*, Oxford University Press.

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," *arXiv preprint arXiv:2210.10024*.

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.

CARROLL, G. (2019): "Robustness in mechanism design and contracting," *Annual Review of Economics*, 11, 139–166.

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.

CUMBUL, E. AND G. VIRÁG (2018): "Multilateral limit pricing in price-setting games," *Games and Economic Behavior*, 111, 250–273.

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=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.
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.
ERDŐS, L., T. KRÜGER, AND D. SCHRÖDER (2019): "Random matrices with slow correlation decay," *Forum of Mathematics, Sigma*, 7, e8.
FELDSTEIN, M. (1976): "On the theory of tax reform," *Journal of Public Economics*, 6, 77–104.
FOUCART, S. AND H. RAUHUT (2013): *A Mathematical Introduction to Compressive Sensing*, New York: Birkhäuser.
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., B. GOLUB, S. GOYAL, E. TALAMÀS, AND O. TAMUZ (2021): "Taxes and market power: A principal components approach," *arXiv preprint arXiv:2112.08153*.
GALEOTTI, A. AND S. GOYAL (2009): "Influencing the Influencers: a theory of strategic diffusion," *The Rand Journal of Economics*, 40, 509–532.
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.
KITAGAWA, T. AND A. TETENOV (2018): "Who Should Be Treated? Empirical Welfare Maximization Methods for Treatment Choice," *Econometrica*, 86, 591–616.
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.
MANSKI, C. F. (2004): "Statistical Treatment Rules for Heterogeneous Populations," *Econometrica*, 72, 1221–1246.
MARSHALL, A. (1890): *Principles of Economics*, New York: Macmillan.
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.
MIRAVETE, E. J., K. SEIM, AND J. THURK (2024): "Targeted Pricing and the Shape of Demand," Tech. rep., Working Paper, Yale University.
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. D. WHINSTON (2022): "Concentration thresholds for horizontal mergers," *American Economic Review*, 112, 1915–1948.

NORRIS, J. J. (2024): "Bounding high dimensional comparative statics," *Available at SSRN 4543870*.

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.

REKER, J. (2022): "On the operator norm of a Hermitian random matrix with correlated entries," *Random Matrices: Theory and Applications*, 11, 2250036.

ROSTEK, M. AND N. YODER (2023): "Reallocative auctions and core selection," *Mimeo*.

SINGH, N. AND X. VIVES (1984): "Price and Quantity Competition in a Differentiated Duopoly," *The Rand Journal of Economics*, 15, 546–554.

SYVERSON, C. (2019): "Macroeconomics and market power: Context, implications, and open questions," *Journal of Economic Perspectives*, 33, 23–43.

TIROLE, J. AND R. GUESNERIE (1981): "Tax reform from the gradient projection viewpoint," *Journal of Public Economics*, 15, 275–293.

VAN DER VAART, A. W. (1998): *Asymptotic Statistics*, Cambridge: Cambridge University Press.

VERSHYNIN, R. (2018): *High-Dimensional Probability: An Introduction with Applications in Data Science*, Cambridge: Cambridge University Press.

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.

WALD, A. (1950): *Statistical Decision Functions*, New York: Wiley.

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. **Proof of Part (2) of Proposition 1.** Since $\boldsymbol{D}$ is a non-diagonal matrix, the normalized $\underline{\boldsymbol{D}} \neq -\boldsymbol{I}$. Hence, there exists $\ell \neq \ell'$ with $\lambda_\ell \neq \lambda'_\ell$. For generic $\underline{\boldsymbol{q}}^0$, we also have that $(\boldsymbol{u}^\ell \cdot \underline{\boldsymbol{q}}^0)(\boldsymbol{u}^{\ell'} \cdot \underline{\boldsymbol{q}}^0) \neq 0$. 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 (8).

For any real number $\beta \neq 0$ and any $\dot{S}$, consider the following class of interventions

$$\underline{\boldsymbol{\sigma}}^0(\beta) = \beta\left((\boldsymbol{u}^1 \cdot \underline{\boldsymbol{q}}^0)\boldsymbol{u}^1 + \alpha(\boldsymbol{u}^2 \cdot \underline{\boldsymbol{q}}^0)\boldsymbol{u}^2\right),$$

where α is chosen so that $\underline{\boldsymbol{\sigma}}^0 \cdot \underline{\boldsymbol{q}}^0 = \dot{S}$; that is:

$$\alpha = \frac{\dot{S}/\beta - (\boldsymbol{u}^1 \cdot \underline{\boldsymbol{q}}^0)^2}{(\boldsymbol{u}^2 \cdot \underline{\boldsymbol{q}}^0)^2}.$$

Hence, as we vary β, we keep $\underline{\boldsymbol{\sigma}}^0(\beta) \cdot \underline{\boldsymbol{q}}^0 = \dot{S}$. Next, note that

$$\begin{aligned}
\dot{C} &= \sum_\ell \frac{1}{1+|\lambda_\ell|}(\boldsymbol{u}^\ell \cdot \underline{\boldsymbol{q}}^0)(\boldsymbol{u}^\ell \cdot \underline{\boldsymbol{\sigma}}^0(\beta)) = \frac{\beta}{1+|\lambda_1|}(\boldsymbol{u}^1 \cdot \underline{\boldsymbol{q}}^0)^2 + \frac{\alpha\beta}{1+|\lambda_2|}(\boldsymbol{u}^2 \cdot \underline{\boldsymbol{q}}^0)^2 \\
&= \frac{\beta}{1+|\lambda_1|}(\boldsymbol{u}^1 \cdot \underline{\boldsymbol{q}}^0)^2 + \frac{\dot{S} - \beta(\boldsymbol{u}^1 \cdot \underline{\boldsymbol{q}}^0)^2}{(\boldsymbol{u}^2 \cdot \underline{\boldsymbol{q}}^0)^2}\frac{1}{1+|\lambda_2|}(\boldsymbol{u}^2 \cdot \underline{\boldsymbol{q}}^0)^2 \\
&= \beta(\boldsymbol{u}^1 \cdot \underline{\boldsymbol{q}}^0)^2\left(\frac{1}{1+|\lambda_1|} - \frac{1}{1+|\lambda_2|}\right) + \frac{1}{1+|\lambda_2|}\dot{S}.
\end{aligned}$$

Hence, $\dot{C}$ evaluated at $\boldsymbol{\sigma}^0(\beta)$ depends linearly on β, which implies that, for a given $\dot{S}$, we can achieve any given $\dot{C}$ by choosing β appropriately.

A.2. **Proof of Proposition 2: Significant structure in the block model demand.** We show that the environment covered by Proposition 2 is a special case of the environment covered by Theorem 1. Since in the block model $D_{ii} = -1$, Assumptions 3 and 4 are satisfied. Thus, it remains to show that significant structure holds. The following Lemma is key for this.

Lemma 3 (Spectral radius and positive block–constant eigenvector)**.** Fix the state and let $\boldsymbol{D}$ satisfy the block structure with K product categories each of size m. Denote its spectral radius by $\rho(\boldsymbol{D}) = \max_j |\lambda_j(\boldsymbol{D})|$. Then:

(a) There exist constants $0 < c_1 \leq c_2 < \infty$, independent of K, such that

$$c_1 K \;\leq\; \rho(\boldsymbol{D}) \;\leq\; c_2 K \qquad \text{for all large enough} K.$$

(b) For all large enough K there exists an eigenpair $(\lambda, \boldsymbol{u})$ of $\boldsymbol{D}$ such that
- $\lambda < 0$ and $|\lambda| \geq c_1 K$;
- $\boldsymbol{u}_i > 0$ for all products i;
- $\boldsymbol{u}_i = \boldsymbol{u}_j$ whenever $k(i) = k(j)$ (the eigenvector is constant within categories).

A key assumption we use for Lemma 3 is part (a)(1) of Condition 1 which rules out the possibility that cross-category demand complementarity vanishes. Part (a) of Lemma 3 shows that, for each state, the spectral radius is of order K. Part (b) of Lemma 3 shows that there is a large (negative) eigenvalue of order K whose eigenvector is strictly positive and constant within product categories.

Lemma 3 shows that $b(n)$ goes to infinity at a rate of n when we increase n by increasing K, holding m constant. Part (b) of Condition 1 implies that the norm of the noise diverges at a rate of $\sqrt{n}$. Hence, part (2) of Definition 2 holds. Part (a)(2) of Condition 1 and part (b) of Lemma 3 imply part (3) of Definition 2 holds.

Proof of Lemma 3. Step 1: Reduction to a $K \times K$ matrix. Let

$$\mathcal{S} \{\boldsymbol{u} \in \mathbb{R}^n : \boldsymbol{u}_i = \boldsymbol{u}_j \text{ whenever } k(i) = k(j)\}$$

be the subspace of vectors that are constant within each category. For $\boldsymbol{u} \in \mathcal{S}$, let u_c denote its common value on category c. If i is a product in category c, then

$$(\boldsymbol{D}\boldsymbol{u})_i = \big(-1 + (m-1)\boldsymbol{B}_{cc}\big)u_c + \sum_{\ell \neq c} m\boldsymbol{B}_{c\ell}u_\ell,$$

which depends only on c. Hence $\boldsymbol{D}\mathcal{S} \subseteq \mathcal{S}$. On $\mathcal{S}$, $\boldsymbol{D}$ acts as the $K \times K$ symmetric matrix $\boldsymbol{A} = (\boldsymbol{A}_{c\ell})$ with

$$\boldsymbol{A}_{cc} = -1 + (m-1)\boldsymbol{B}_{cc}, \qquad \boldsymbol{A}_{c\ell} = m\boldsymbol{B}_{c\ell} \quad (c \neq \ell). \tag{16}$$

In particular, if $\boldsymbol{A}\boldsymbol{y} = \lambda\boldsymbol{y}$, then the lifted vector $\boldsymbol{u} \in \mathbb{R}^n$ defined by $\boldsymbol{u}_i := \boldsymbol{y}_{k(i)}$ satisfies $\boldsymbol{D}\boldsymbol{u} = \lambda\boldsymbol{u}$.

Step 2: Spectral-radius bounds. Let $\boldsymbol{v} := \mathbf{1}/\sqrt{K}$. Since $\boldsymbol{A}$ is symmetric, Rayleigh–Ritz implies

$$\lambda_{\min}\big(\boldsymbol{A}\big) \leq \boldsymbol{v}^\top \boldsymbol{A}\boldsymbol{v} = \frac{1}{K}\sum_{c,\ell=1}^{K} \boldsymbol{A}_{c\ell}.$$

By Property NSD and the normalization $D_{ii} = -1$, we have $|D_{ij}| \leq 1$ for all $i \neq j$, hence $\boldsymbol{B}_{cc} \leq 1$ and $\boldsymbol{A}_{cc} \leq m - 2$. Moreover, by part (a)(1) of Condition 1, for $c \neq \ell$ we have $\boldsymbol{B}_{c\ell} < \kappa < 0$, so $\boldsymbol{A}_{c\ell} < m\kappa$. Therefore,

$$\lambda_{\min}\big(\boldsymbol{A}\big) \leq (m-2) + m\kappa(K-1) \leq -c_1 K$$

for all sufficiently large K, for some $c_1 > 0$ independent of K. Since $\lambda_{\min}(\boldsymbol{A})$ is also an eigenvalue of $\boldsymbol{D}$, it follows that $\rho(\boldsymbol{D}) \geq |\lambda_{\min}(\boldsymbol{A})| \geq c_1 K$.

For the upper bound, again $|D_{ij}| \leq 1$ implies

$$\rho\big(\boldsymbol{D}\big) \leq \|\boldsymbol{D}\|_\infty := \max_i \sum_{j=1}^{n} |D_{ij}| \leq n = mK,$$

which proves part (a) with $c_2 := m$.

Step 3: A strictly positive block-constant eigenvector. Define $\boldsymbol{M} := m\boldsymbol{I} - \boldsymbol{A}$. Then $\boldsymbol{M}_{cc} = m - \boldsymbol{A}_{cc} \geq 2$ and, for $c \neq \ell$, $\boldsymbol{M}_{c\ell} = -\boldsymbol{A}_{c\ell} = -m\boldsymbol{B}_{c\ell} > 0$, so $\boldsymbol{M}$ is entrywise positive. By the Perron–Frobenius theorem, $\boldsymbol{M}$ has a largest eigenvalue $\mu = \rho(\boldsymbol{M})$ with an associated eigenvector $\boldsymbol{y} \in \mathbb{R}^K$ satisfying $\boldsymbol{y}_c > 0$ for all c. Since $\boldsymbol{M} = m\boldsymbol{I} - \boldsymbol{A}$, we have $\boldsymbol{A}\boldsymbol{y} = (m - \mu)\boldsymbol{y}$. Because the eigenvalues of $\boldsymbol{M}$ are $m - \lambda_j(\boldsymbol{A})$, we have $m - \mu = \lambda_{\min}(\boldsymbol{A})$. Thus $\boldsymbol{y}$ is an eigenvector of $\boldsymbol{A}$ associated with $\lambda_{\min}(\boldsymbol{A})$, which satisfies $\lambda_{\min}(\boldsymbol{A}) \leq -c_1 K$ by Step 2. Lifting $\boldsymbol{y}$ to $\boldsymbol{u} \in \mathbb{R}^n$ by $\boldsymbol{u}_i := \boldsymbol{y}_{k(i)}$ yields an eigenpair $(\lambda, \boldsymbol{u})$ of $\boldsymbol{D}$ with $\lambda = \lambda_{\min}(\boldsymbol{A}) < 0$, $|\lambda| \geq c_1 K$, $\boldsymbol{u}_i > 0$ for all i, and $\boldsymbol{u}$ constant within categories. This proves part (b). □

A.3. **Proof of Theorem 1.** First, we introduce a normalization that will be helpful in some of our asymptotic arguments, to avoid carrying around the norm of the quantity vector.

Assumption 5. $\|\boldsymbol{q}^0\| = 1$.

Remark 3. Normalization of $\boldsymbol{q}^0$. Our analysis is invariant to a common rescaling of quantities together with a compensating change in the unit of the numeraire. Fix any sequence of market states $\boldsymbol{\theta}(n) = (\boldsymbol{D}(n), \boldsymbol{q}^0(n)) \in \boldsymbol{\Theta}(n)$ in the environment, with associated baseline price vectors $\boldsymbol{p}^0(n)$. For any sequence $\kappa(n) > 0$, define transformed quantities by $\tilde{\boldsymbol{q}}^0(n) := \kappa(n)\boldsymbol{q}^0(n)$ and measure the numeraire in units that are $\kappa(n)$ times larger. Then the numerical price vectors $\boldsymbol{p}^0(n)$, budget sets, Hicksian demands, and normalized Slutsky matrices are unchanged, and welfare and expenditure derivatives (which depend on $(\boldsymbol{q}^0(n), \boldsymbol{\sigma}(n))$ only through inner products) are identical, in terms of the new numeraire, to those computed from $(\tilde{\boldsymbol{q}}^0(n), \tilde{\boldsymbol{\sigma}}(n))$. Taking $\kappa(n) = 1/\|\boldsymbol{q}^0(n)\|$ yields an observationally equivalent specification with $\|\tilde{\boldsymbol{q}}^0(n)\| = 1$ for all n, so the normalization entails no loss of generality. Under this rescaling, the content of the significant structure condition does not change because the rescaling constants cancel out of all the relevant comparisons.

We also assume the normalization $D_{ii} = -1$ for all i, which is also without loss of generality by Section 2.7. This saves on writing underlines throughout the proof.

Next, we introduce some notation. For any matrix $\boldsymbol{M}$, we define $\mathbb{\Lambda}(\boldsymbol{M}, \underline{\lambda})$ as the set of eigenvalues of $\boldsymbol{M}$ with absolute value at least $\underline{\lambda}$, and $\mathcal{L}(\boldsymbol{M}, \underline{\lambda})$ as the space spanned by corresponding eigenvectors.

For the rest of the proof we will fix an environment $\mathcal{E} = (\boldsymbol{\Theta}(n), \varphi(n))$ with significant structure and let $b(n)$ denote the associated sequence from Definition 2. Fix an arbitrary market state $\boldsymbol{\theta} = (\boldsymbol{D}, \boldsymbol{q}^0) \in \boldsymbol{\Theta}(n)$; by Definition 2, all bounds below are required to hold uniformly over such $\boldsymbol{\theta}$.

Under Definition 2, the signal about $\boldsymbol{D}$ can be written $\widehat{\boldsymbol{D}} = \boldsymbol{D} + \boldsymbol{E}$, where condition (2) of Definition 2 guarantees that, for every $\boldsymbol{\theta} \in \boldsymbol{\Theta}(n)$ and every $\eta > 0$,

$$\mathbb{P}_{\varphi_{\boldsymbol{\theta}}(n)}(\|\boldsymbol{E}\| > \eta\, b(n)) \to_n 0. \tag{17}$$

We now define three threshold sequences

$$\underline{M}(n) := \tfrac{1}{2}b(n), \quad \widehat{M}(n) := \tfrac{3}{4}b(n), \quad M(n) := b(n).$$

For each n:

$$\underline{M}(n) < \widehat{M}(n) < M(n),$$

and when n is sufficiently large the differences

$$M(n) - \underline{M}(n), \quad \widehat{M}(n) - \underline{M}(n) \quad \text{and} \quad M(n) - \widehat{M}(n)$$

are all of order $b(n)$. Since $\|\boldsymbol{E}\| = o_{\mathrm{p},\boldsymbol{\Theta}}(b(n))$ by (17), these differences dominate $\|\boldsymbol{E}\|$ in probability.

Now fix an arbitrary sequence $(\boldsymbol{D}(n), \boldsymbol{q}^0(n)) \in \boldsymbol{\Theta}(n)$. We use the following notation:

$$\underline{\Lambda}(n) := \mathbb{\Lambda}(\boldsymbol{D}; \underline{M}(n)), \quad \widehat{\Lambda}(n) := \mathbb{\Lambda}(\widehat{\boldsymbol{D}}; \widehat{M}(n)), \quad \Lambda(n) := \mathbb{\Lambda}(\boldsymbol{D}; M(n))$$
$$\underline{L}(n) := \mathcal{L}(\boldsymbol{D}, \underline{M}(n)), \quad \widehat{L}(n) := \mathcal{L}(\widehat{\boldsymbol{D}}, \widehat{M}(n)), \quad L(n) := \mathcal{L}(\boldsymbol{D}, M(n)).$$

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

Proof strategy: sandwiching eigenspaces. The three thresholds define three nested eigenspaces. The *core space* $L(n)$ (threshold M) contains the eigenvectors of $\boldsymbol{D}$ with eigenvalues whose absolute value is at least $b(n)$. This is the space that we know is nonempty; indeed, this is exactly what is guaranteed, uniformly across possible market states, by condition (3) of Definition 2. The *estimated space* $\widehat{L}(n)$ (threshold $\widehat{M}$, applied to the noisy matrix $\widehat{\boldsymbol{D}}$) is what we can actually compute from data; because $\|\boldsymbol{E}\| = o_{\mathrm{p},\boldsymbol{\Theta}}(b(n))$, the Davis–Kahan theorem guarantees that this space captures all directions in $L(n)$. The gap $M(n) - \widehat{M}(n) = \Theta(b(n))$ provides sufficient slack: any true eigenvalue above $M(n)$ produces an estimated eigenvalue above $\widehat{M}(n)$, and moreover the associated eigenvectors are almost fully in the span of eigenvectors of $\widehat{\boldsymbol{D}}$ with eigenvalues at least $\widehat{M}$.

The *buffer space* $\underline{L}(n)$ (threshold $\underline{M}$) plays a subtler role. We will design an intervention $\boldsymbol{\sigma}$ that lives in the estimated space $\widehat{L}(n)$—this is what we can compute from data. But to verify that this intervention actually works, we must analyze its welfare effects in the *true* eigenbasis of $\boldsymbol{D}$, which we do not observe. The buffer space provided by $\underline{L}(n)$ helps us achieve this: it is defined based on the true matrix, yet Davis–Kahan guarantees that $\widehat{L}(n)$ is almost entirely contained in $\underline{L}(n)$. And $\underline{L}(n)$, since it is defined using the true $\boldsymbol{D}$, has nice properties (in particular, being spanned by eigenvectors with large eigenvalues) that guarantee good behavior of the designed intervention. Here again we use Davis–Kahan to show that the differences between $\widehat{L}(n)$ and $\underline{L}(n)$ cause only mild errors that do not disrupt a good intervention. We now dive into the details.

We now define our intervention:

$$\boldsymbol{\sigma} = \frac{P_{\widehat{L}(n)} \widehat{\boldsymbol{q}}^0}{\|P_{\widehat{L}(n)} \widehat{\boldsymbol{q}}^0\|^2} \tag{18}$$

The expenditure of this intervention is

$$\dot{S} = \boldsymbol{\sigma} \cdot \boldsymbol{q}^0 = \frac{P_{\widehat{L}(n)} \widehat{\boldsymbol{q}}^0}{\|P_{\widehat{L}(n)} \widehat{\boldsymbol{q}}^0\|^2} \cdot \boldsymbol{q}^0 \tag{19}$$

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 $\widehat{\boldsymbol{D}}$. 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.

Lemma 4. The norm $\|P_{\underline{L}(n)}^\perp P_{\widehat{L}(n)}\|$ converges to 0 in probability as $n \to \infty$.

Proof. Standard eigenvalue perturbation bounds and the Davis–Kahan theorem guarantee 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 properties hold. First, every eigenvalue in the set $\Lambda(n)$, which is nonempty by construction (it always

contains at least the largest eigenvalue of $\boldsymbol{D}$), has a corresponding eigenvalue within distance $O(b(n))$, and therefore in $\widehat{\Lambda}(n)$. Second,

$$\|P_{\underline{L}(n)}^{\perp} P_{\widehat{L}(n)}\| \leq \frac{2\|\boldsymbol{E}\|}{g}, \tag{20}$$

where g, the "gap," is the minimum distance between some eigenvalue of $\widehat{\boldsymbol{D}}$ in $\widehat{\Lambda}(n)$ and some eigenvalue of $\boldsymbol{D}$ not contained in $\underline{\Lambda}(n)$. By our choice of thresholds, this gap is at least $\widehat{M}(n) - \underline{M}(n)$, which is a fixed positive fraction of $b(n)$ and therefore bounded below by a positive constant multiple of $b(n)$ for all large n. Combining this with $\|\boldsymbol{E}\| = o_{\mathrm{p},\boldsymbol{\Theta}}(b(n))$, we obtain that the right-hand side of (20) converges to zero in probability. Thus, increasing $N_1(\eta)$ if necessary, we conclude the statement of the lemma. □

Lemma 5. As $n \to \infty$, the expenditure derivative $\dot{S}$ converges in probability to 1.

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

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

By the Cauchy–Schwarz inequality,

$$\left|\dot{S} - 1\right| \leq \frac{\|P_{\widehat{L}(n)} \boldsymbol{\varepsilon}\|}{\|P_{\widehat{L}(n)} \widehat{\boldsymbol{q}}^0\|}.$$

Condition (3) of Definition 2 and Assumption 5 imply $\|P_{L(n)} \boldsymbol{q}^0\| \geq \delta$. By the same Davis–Kahan perturbation argument used above (now comparing $L(n)$ and $\widehat{L}(n)$ with gap $M(n) - \widehat{M}(n) = \Theta(b(n))$), we have $\|P_{\widehat{L}(n)}^{\perp} P_{L(n)}\| \xrightarrow{p} 0$. Moreover,

$$P_{L(n)} \boldsymbol{q}^0 = P_{L(n)} P_{\widehat{L}(n)} \boldsymbol{q}^0 + P_{L(n)} P_{\widehat{L}(n)}^{\perp} \boldsymbol{q}^0.$$

Therefore,

$$\delta \leq \|P_{L(n)} \boldsymbol{q}^0\| \leq \|P_{\widehat{L}(n)} \boldsymbol{q}^0\| + \|P_{L(n)} P_{\widehat{L}(n)}^{\perp}\| \, \|\boldsymbol{q}^0\| = \|P_{\widehat{L}(n)} \boldsymbol{q}^0\| + o_{\mathrm{p}}(1),$$

where the last equality uses Assumption 5. Hence

$$\|P_{\widehat{L}(n)} \boldsymbol{q}^0\| \geq \delta - o_{\mathrm{p}}(1).$$

Since $\widehat{\boldsymbol{q}}^0 = \boldsymbol{q}^0 + \boldsymbol{\varepsilon}$,

$$\|P_{\widehat{L}(n)} \widehat{\boldsymbol{q}}^0\| \geq \|P_{\widehat{L}(n)} \boldsymbol{q}^0\| - \|P_{\widehat{L}(n)} \boldsymbol{\varepsilon}\|.$$

So it remains to show that $\|P_{\widehat{L}(n)} \boldsymbol{\varepsilon}\|$ tends in probability to 0.

For this, we state a lemma.

Lemma 6. Fix an environment that has significant structure, subject to the normalization $\|\boldsymbol{q}^0\| = 1$. Let $V(n)$ be a sequence of linear subspaces of $\mathbb{R}^n$, possibly random but measurable with respect to $\widehat{\boldsymbol{D}}$, such that $\dim V(n)/n \xrightarrow{p} 0$. Then $\|P_{V(n)} \boldsymbol{\varepsilon}\|$ tends to 0 in probability.

Proof. If $\boldsymbol{\varepsilon} = \mathbf{0}$ almost surely, there is nothing to prove. Otherwise let $\sigma_n^2 := \mathrm{Var}(\varepsilon_i) \in (0, \infty)$ and define $Z_i := \varepsilon_i/\sigma_n$, so $Z := (Z_1, \ldots, Z_n)^\mathsf{T}$ has i.i.d. entries with mean 0 and variance 1. Since scaling cancels,

$$\frac{\|P_{V(n)}\boldsymbol{\varepsilon}\|^2}{\|\boldsymbol{\varepsilon}\|^2} = \frac{\|P_{V(n)}Z\|^2}{\|Z\|^2} = \frac{\frac{1}{n}\|P_{V(n)}Z\|^2}{\frac{1}{n}\|Z\|^2}.$$

Let $d(n) := \dim V(n)$. Because $V(n)$ is measurable with respect to $\widehat{\boldsymbol{D}}$ and $\boldsymbol{\varepsilon}$ is independent of $\boldsymbol{E}$ by our signal assumption, the vector Z is independent of $V(n)$. Therefore,

$$\mathbb{E}\left[\frac{1}{n}\|P_{V(n)}Z\|^2 \,\middle|\, V(n)\right] = \frac{1}{n}\,\mathrm{tr}\big(P_{V(n)}\big) = \frac{d(n)}{n}.$$

Since $0 \le d(n)/n \le 1$ and $d(n)/n \xrightarrow{p} 0$, we also have $\mathbb{E}[d(n)/n] \to 0$. Markov's inequality then gives, for every $\eta > 0$,

$$\mathbb{P}\left(\frac{1}{n}\|P_{V(n)}Z\|^2 > \eta\right) \le \frac{1}{\eta}\,\mathbb{E}\left[\frac{1}{n}\|P_{V(n)}Z\|^2\right] = \frac{1}{\eta}\,\mathbb{E}\left[\frac{d(n)}{n}\right] \to 0.$$

By the weak law of large numbers, $\frac{1}{n}\|Z\|^2 = \frac{1}{n}\sum_{i=1}^n Z_i^2 \xrightarrow{p} 1$. Slutsky's theorem implies $\|P_{V(n)}\boldsymbol{\varepsilon}\|/\|\boldsymbol{\varepsilon}\| \xrightarrow{p} 0$. Under significant structure and the normalization $\|\boldsymbol{q}^0\| = 1$, condition (1) of Definition 2 implies $\|\boldsymbol{\varepsilon}\| = O_{\mathrm{p},\boldsymbol{\Theta}}(1)$, and therefore $\|P_{V(n)}\boldsymbol{\varepsilon}\| \xrightarrow{p} 0$. □

Applying Lemma 6 with $V(n) = \widehat{L}(n)$, and using that[34] $\dim \widehat{L}(n)/n \xrightarrow{p} 0$, we obtain $\|P_{\widehat{L}(n)}\boldsymbol{\varepsilon}\| \xrightarrow{p} 0$. Combining this with what we stated just before the lemma,

$$\|P_{\widehat{L}(n)}\widehat{\boldsymbol{q}}^0\| \ge \|P_{\widehat{L}(n)}\boldsymbol{q}^0\| - \|P_{\widehat{L}(n)}\boldsymbol{\varepsilon}\|,$$

shows that $\|P_{\widehat{L}(n)}\widehat{\boldsymbol{q}}^0\| \ge \delta/2$ with probability $1 - o(1)$. Returning to the bound on $|\dot{S} - 1|$, we conclude that $\dot{S} \xrightarrow{p} 1$. □

Now, to prove the theorem, we use Lemma 2 to write:

$$\dot{W} = \underbrace{\sum_{\lambda_\ell \in \underline{\Lambda}(n)} (\boldsymbol{u}^\ell \cdot \boldsymbol{q}^0)(\boldsymbol{u}^\ell \cdot \boldsymbol{\sigma})\frac{|\lambda_\ell|}{1+|\lambda_\ell|}}_{\dot{W}_M} + \underbrace{\sum_{\lambda_\ell \notin \underline{\Lambda}(n)} (\boldsymbol{u}^\ell \cdot \boldsymbol{q}^0)(\boldsymbol{u}^\ell \cdot \boldsymbol{\sigma})\frac{|\lambda_\ell|}{1+|\lambda_\ell|}}_{\dot{W}_R}. \tag{21}$$

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

$$\dot{S} = \underbrace{\sum_{\lambda_\ell \in \underline{\Lambda}(n)} (\boldsymbol{u}^\ell \cdot \boldsymbol{q}^0)(\boldsymbol{u}^\ell \cdot \boldsymbol{\sigma})}_{\dot{S}_M} + \underbrace{\sum_{\lambda_\ell \notin \underline{\Lambda}(n)} (\boldsymbol{u}^\ell \cdot \boldsymbol{q}^0)(\boldsymbol{u}^\ell \cdot \boldsymbol{\sigma})}_{\dot{S}_R}. \tag{22}$$

The proof can be completed with two further lemmas.

Lemma 7. As $n \to \infty$, both $\dot{W}_R$ and $\dot{S}_R$ converge in probability to 0.

[34] Weyl's inequality implies that with probability $1 - o(1)$, every eigenvalue counted by $\widehat{L}(n)$ comes from a true eigenvalue of $\boldsymbol{D}$ whose absolute value exceeds $b(n)/2$. Since $\boldsymbol{D}$ is negative semidefinite with trace $-n$, there are at most $2n/b(n) = o(n)$ such eigenvalues.

Proof. This follows from Lemma 4 and the fact that by construction, $\boldsymbol{\sigma} \in \widehat{L}(n)$. □

Lemma 8. As $n \to \infty$,

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

Proof. Using the expressions above, write the difference

$$|\dot{W}_M - \dot{S}_M| \leq \sum_{\lambda_\ell \in \underline{\Lambda}(n)} \frac{1}{1+|\lambda_\ell|} \left|(\boldsymbol{u}^\ell \cdot \boldsymbol{q}^0)(\boldsymbol{u}^\ell \cdot \boldsymbol{\sigma})\right|.$$

Note that for all $\ell \in \underline{\Lambda}(n)$, $|\lambda_\ell| \geq \underline{M}(n)$, so $\frac{1}{1+|\lambda_\ell|} \leq \frac{1}{\underline{M}(n)}$. Therefore,

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

By Cauchy–Schwarz for each ℓ and then summing, we have

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

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

Since $\|P_{\underline{L}(n)}\boldsymbol{q}^0\| \leq \|\boldsymbol{q}^0\| = O(1)$ by Assumption 5, and $\|P_{\underline{L}(n)}\boldsymbol{\sigma}\| \leq \|\boldsymbol{\sigma}\| = O_{\mathrm{p}}(1)$ by the denominator lower bound in the proof of Lemma 5, the result follows. □

We can put everything together. Lemma 5 gives that $\dot{S} \xrightarrow{p} 1$. Combining this with (22) and Lemma 7 (which says that $\dot{S}_R \xrightarrow{p} 0$) we see that $\dot{S}_M \xrightarrow{p} 1$. Then using Lemma 8, we find that $\dot{W}_M \xrightarrow{p} 1$. Another application of Lemma 7 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,

$$\dot{C}^h = -\boldsymbol{q}^h \cdot \dot{\boldsymbol{p}} = -\sum_{\ell=1}^{n} (\boldsymbol{u}^\ell \cdot \boldsymbol{q}^h)(\boldsymbol{u}^\ell \cdot \dot{\boldsymbol{p}}) \tag{23}$$

Using Lemma 1:

$$\boldsymbol{u}^\ell \cdot \dot{\boldsymbol{p}} = -\frac{1}{1+|\lambda_\ell|}(\boldsymbol{u}^\ell \cdot \boldsymbol{\sigma}) \tag{24}$$

Then by an argument very similar to the proof of Lemma 7, we conclude that $\dot{C}^h \xrightarrow{p} 0$.

A.4. **Expectation guarantees.** Here we upgrade our in-probability guarantees on the performance of the intervention to expectation guarantees.

A bounded modification of the intervention rule

The intervention in the proof of Theorem 1 is (up to scaling to hit spending s) of the form

$$\boldsymbol{\sigma} = P_{\widehat{L}(n)}\widehat{\boldsymbol{q}}^0 / \|P_{\widehat{L}(n)}\widehat{\boldsymbol{q}}^0\|^2,$$

which can be large on realizations where $\|P_{\widehat{L}(n)}\widehat{\boldsymbol{q}}^0\|$ is very small. Under significant structure, these realizations are asymptotically negligible uniformly over states. For expectation statements, it is convenient to enforce boundedness deterministically.

Fix the target spending level $s > 0$, and let $\delta > 0$ be the projection lower bound in condition (3) of Definition 2 after the normalization $\|\boldsymbol{q}^0\| = 1$ from Assumption 5. Define the truncated rule

$$\boldsymbol{R}^{\mathrm{tr}}(\widehat{\boldsymbol{\theta}}) = s \cdot \frac{P_{\widehat{L}(n)}\widehat{\boldsymbol{q}}^0}{\max\left\{\|P_{\widehat{L}(n)}\widehat{\boldsymbol{q}}^0\|^2,\ (\delta/4)^2\right\}}.$$

Then for all signal realizations,

$$\|\boldsymbol{R}^{\mathrm{tr}}(\widehat{\boldsymbol{\theta}})\| \le \frac{4s}{\delta}. \tag{25}$$

Moreover, by the same argument used in the proof of Theorem 1, we have $\|P_{\widehat{L}(n)}\widehat{\boldsymbol{q}}^0\| \ge \delta/2$ with probability $1 - o(1)$ uniformly over $\boldsymbol{\theta}$, so the truncation event $\{\|P_{\widehat{L}(n)}\widehat{\boldsymbol{q}}^0\| < \delta/4\}$ has probability $o(1)$ uniformly over states. Therefore, $\boldsymbol{R}^{\mathrm{tr}}$ coincides with the original rule except on a vanishing-probability event, so replacing the rule by $\boldsymbol{R}^{\mathrm{tr}}$ does not change any ϵ–robustness conclusions.

Uniform boundedness of the welfare objects

Let $\boldsymbol{\sigma} = \boldsymbol{R}^{\mathrm{tr}}(\widehat{\boldsymbol{\theta}})$. Since $\|\boldsymbol{q}^0\| = 1$,

$$|\dot{S}_{\boldsymbol{\sigma}}| = |\boldsymbol{q}^0 \cdot \boldsymbol{\sigma}| \le \|\boldsymbol{\sigma}\|.$$

By Lemma 1, the pass-through multipliers that map $\boldsymbol{\sigma}$ into first-order price changes have absolute values in $[0, 1]$ in every eigenvector; equivalently, there is a uniform operator-norm bound $\|\dot{\boldsymbol{p}}_{\boldsymbol{\sigma}}\| \le \|\boldsymbol{\sigma}\|$. Hence $|\dot{C}_{\boldsymbol{\sigma}}| = |\boldsymbol{q}^0 \cdot \dot{\boldsymbol{p}}_{\boldsymbol{\sigma}}| \le \|\dot{\boldsymbol{p}}_{\boldsymbol{\sigma}}\| \le \|\boldsymbol{\sigma}\|$. Using $\dot{W}_{\boldsymbol{\sigma}} = \dot{P}_{\boldsymbol{\sigma}} + \dot{C}_{\boldsymbol{\sigma}} - \dot{S}_{\boldsymbol{\sigma}}$ and the identity $\dot{P}_{\boldsymbol{\sigma}} = 2(\dot{S}_{\boldsymbol{\sigma}} - \dot{C}_{\boldsymbol{\sigma}})$ from (8) there is a constant K such that

$$|\dot{S}_{\boldsymbol{\sigma}}|,\ |\dot{C}_{\boldsymbol{\sigma}}|,\ |\dot{P}_{\boldsymbol{\sigma}}|,\ |\dot{W}_{\boldsymbol{\sigma}}| \le K\|\boldsymbol{\sigma}\| \le K \cdot \frac{4s}{\delta} \qquad \text{almost surely, uniformly over } \boldsymbol{\theta}.$$

Thus each welfare error term appearing in Theorem 1 is uniformly bounded under $\boldsymbol{R}^{\mathrm{tr}}$.

Conclusion: High-probability guarantees imply expectation guarantees

Proposition 4 (Expectation upgrade)**.** Maintain the assumptions of Theorem 1 and fix spending $s > 0$. Replace the intervention rule in the proof of Theorem 1 by the bounded modification $\boldsymbol{R}^{\mathrm{tr}}$ above, and let $\boldsymbol{\sigma} = \boldsymbol{R}^{\mathrm{tr}}(\widehat{\boldsymbol{\theta}})$. Then uniformly over $\boldsymbol{\theta} \in \boldsymbol{\Theta}(n)$,

$$\mathbb{E}\big[|\dot{S}_{\boldsymbol{\sigma}} - s|\big] \to 0, \qquad \mathbb{E}\big[|\dot{C}_{\boldsymbol{\sigma}}|\big] \to 0, \qquad \mathbb{E}\big[|\dot{P}_{\boldsymbol{\sigma}} - 2s|\big] \to 0, \qquad \mathbb{E}\big[|\dot{W}_{\boldsymbol{\sigma}} - s|\big] \to 0.$$

Proof. As noted above, $\boldsymbol{R}^{\mathrm{tr}}$ differs from the original rule only on a vanishing-probability event uniformly over $\boldsymbol{\theta}$, so the same high-probability statements remain valid. For each displayed term, let $X_n(\boldsymbol{\theta})$ denote the corresponding error random variable (e.g. $\dot{S}_{\boldsymbol{\sigma}} - s$). The proof of Theorem 1 gives $X_n = o_{\mathrm{p},\boldsymbol{\Theta}}(1)$, and the previous subsection gives a uniform almost-sure bound $|X_n(\boldsymbol{\theta})| \le M$ for some finite M independent of n and $\boldsymbol{\theta}$. For such random variables, the L^1 convergence of X_n follows from the corresponding in-probability convergence statement proved in Theorem 1. □

Because all conclusions are uniform over $\boldsymbol{\theta}$, they also hold under any prior over states by iterated expectations.

Appendix B. Proof of Proposition 3

All arguments in this proof are in normalized coordinates. The constructions below have $D_{ii} = -1$, so underlined and unnormalized variables coincide. We use Hadamard constructions to hide, respectively, the welfare-relevant demand direction and the consumer-surplus-relevant quantity direction.

We use a common Hadamard setup. Work along the subsequence $n = 2^m$, and let $\{\boldsymbol{u}^1, \dots, \boldsymbol{u}^n\}$ be an orthonormal Hadamard basis with $\boldsymbol{u}^n = \mathbf{1}/\sqrt{n}$. Thus $(u_i^\ell)^2 = 1/n$ for every i and ℓ. Whenever eigenvalues $\lambda_1, \dots, \lambda_n$ satisfy $\sum_{\ell=1}^n \lambda_\ell = -n$, the matrix

$$\boldsymbol{D} = \sum_{\ell=1}^{n} \lambda_\ell \boldsymbol{u}^\ell (\boldsymbol{u}^\ell)^\mathsf{T}$$

has diagonal entries equal to -1. If all eigenvalues are nonpositive, then $\boldsymbol{D}$ is negative semidefinite. The bounded-own-price and recoverable-diagonal assumptions then hold immediately.

B.1. **Part 1: No significant structure.** The quantity vector in this construction is common across states and is observed without error. Because the inequalities in Proposition 3 are homogeneous in the intervention, any signal-measurable rule with positive expenditure can be rescaled, using the observed $\boldsymbol{q}^0$, to satisfy

$$\dot{S}_{\boldsymbol{\sigma}} = \boldsymbol{q}^0 \cdot \boldsymbol{\sigma} = 1.$$

Under this normalization, part (1) requires $\dot{W}_{\boldsymbol{\sigma}} \geq \varepsilon$, because $\dot{C} + \dot{P} = \dot{S} + \dot{W}$.

Let $z_n = \lfloor n^{1/4} \rfloor$ and define the common quantity vector

$$\boldsymbol{q}^0 = 2z_n \boldsymbol{u}^n + \sum_{k=1}^{z_n} \boldsymbol{u}^k.$$

This vector is strictly positive for all large n, and its only nonzero Hadamard coordinates are $\boldsymbol{u}^n \cdot \boldsymbol{q}^0 = 2z_n$ and $\boldsymbol{u}^k \cdot \boldsymbol{q}^0 = 1$ for $k \leq z_n$. The hidden state will identify one of these z_n quantity-loaded Hadamard directions as the high-pass-through direction. Because the signal does not reveal this index, any signal-measurable rule must choose the same intervention without knowing which of the z_n directions is valuable.

For each hidden state $j \in \{1, \dots, z_n\}$, set

$$\tau_n := \frac{1}{z_n}\left(\frac{n^{1/2}}{1+n^{1/2}} + (z_n - 1)\frac{n^{-2}}{1+n^{-2}}\right),$$

so τ_n is the average welfare pass-through weight across the hidden directions in a fixed state: one direction has pass-through close to 1, and the other $z_n - 1$ directions have pass-through close to 0. Thus $\tau_n \leq 1/z_n + n^{-2} \to 0$. Let the eigenvalues of $\boldsymbol{D}^{(j)}$ be

$$\lambda_j^{(j)} = -n^{1/2}, \qquad \lambda_k^{(j)} = -n^{-2} \quad \text{for } k \leq z_n,\ k \neq j,$$

$$\lambda_n^{(j)} = -\frac{\tau_n}{1-\tau_n},$$

and, for all $\ell \notin \{1, \dots, z_n, n\}$,

$$\lambda_\ell^{(j)} = \frac{-n + n^{1/2} + (z_n - 1)n^{-2} + \tau_n/(1-\tau_n)}{n - z_n - 1}.$$

The eigenvalue in the common direction $\boldsymbol{u}^n$ is chosen so that its pass-through weight is also τ_n, since $|\lambda_n^{(j)}|/(1 + |\lambda_n^{(j)}|) = \tau_n$.

The last value makes the trace equal to $-n$ and converges to -1. Hence all eigenvalues are negative for large n.

The signal is deterministic and common across states:

$$\widehat{\boldsymbol{D}} = -\boldsymbol{I}, \qquad \widehat{\boldsymbol{q}}^0 = \boldsymbol{q}^0.$$

Thus a rule chooses the same normalized intervention $\boldsymbol{\sigma}$ in every state. Since $\boldsymbol{q}^0$ is orthogonal to all remaining Hadamard directions, the normalization gives

$$(\boldsymbol{u}^n \cdot \boldsymbol{q}^0)(\boldsymbol{u}^n \cdot \boldsymbol{\sigma}) + \sum_{k=1}^{z_n} (\boldsymbol{u}^k \cdot \boldsymbol{q}^0)(\boldsymbol{u}^k \cdot \boldsymbol{\sigma}) = 1.$$

By Equation (10), in state j,

$$\dot{W}_{\boldsymbol{\sigma}}^{(j)} = \frac{n^{1/2}}{1+n^{1/2}}(\boldsymbol{u}^j \cdot \boldsymbol{q}^0)(\boldsymbol{u}^j \cdot \boldsymbol{\sigma}) + \frac{n^{-2}}{1+n^{-2}} \sum_{\substack{k=1 \\ k \neq j}}^{z_n} (\boldsymbol{u}^k \cdot \boldsymbol{q}^0)(\boldsymbol{u}^k \cdot \boldsymbol{\sigma}) + \tau_n (\boldsymbol{u}^n \cdot \boldsymbol{q}^0)(\boldsymbol{u}^n \cdot \boldsymbol{\sigma}).$$

Averaging over $j = 1, \dots, z_n$ and using the definition of τ_n yields

$$\frac{1}{z_n} \sum_{j=1}^{z_n} \dot{W}_{\boldsymbol{\sigma}}^{(j)} = \tau_n \left[\sum_{k=1}^{z_n} (\boldsymbol{u}^k \cdot \boldsymbol{q}^0)(\boldsymbol{u}^k \cdot \boldsymbol{\sigma}) + (\boldsymbol{u}^n \cdot \boldsymbol{q}^0)(\boldsymbol{u}^n \cdot \boldsymbol{\sigma}) \right] = \tau_n.$$

Therefore some state j satisfies $\dot{W}_{\boldsymbol{\sigma}}^{(j)} \leq \tau_n$. Since $\tau_n \to 0$, for every fixed $\varepsilon \in (0, 1/2)$ and all sufficiently large n, the normalized target $\dot{W}_{\boldsymbol{\sigma}} \geq \varepsilon$ fails conditional on that state (with probability 1, since the signal is deterministic).

It remains only to verify that this environment lacks significant structure. In state j, the observation error $\boldsymbol{E}^{(j)} = -\boldsymbol{I} - \boldsymbol{D}^{(j)}$ has operator norm of order $n^{1/2}$. If a candidate threshold $b(n) \to \infty$ is not asymptotically larger than $n^{1/2}$, condition (2) of Definition 2 fails. If $b(n)/n^{1/2} \to \infty$, then the large-eigenvalue space is eventually empty, so condition (3) fails. Hence significant structure does not hold.

B.2. **Part 2: Significant structure present.** We now construct an environment with significant structure in which the large-eigenvalue direction is perfectly recoverable, but the quantity signal does not reveal a small component of quantities that is relevant for consumer surplus.

Let $\alpha_n = n^{-1}$. Work with the same Hadamard basis as above. Set

$$\eta_n := n^{-2}, \qquad \lambda_\ell = -\eta_n \quad (\ell = 1, \dots, n-1), \qquad \lambda_n = -n + (n-1)\eta_n,$$

and define

$$\boldsymbol{D} := \sum_{\ell=1}^{n} \lambda_\ell \boldsymbol{u}^\ell (\boldsymbol{u}^\ell)^\mathsf{T}.$$

The eigenvalues sum to $-n$, so the diagonal entries of $\boldsymbol{D}$ are -1. All eigenvalues are negative. Thus $\boldsymbol{D}$ is negative semidefinite, the bounded-own-price assumption holds, and the diagonal is exactly recoverable.

The state space consists of the quantity vectors

$$\boldsymbol{q}^{j,s} := \boldsymbol{u}^n + s\alpha_n \boldsymbol{u}^j, \qquad j \in \{1, \dots, n-1\}, \quad s \in \{-1, 1\}.$$

Because $\boldsymbol{u}^n = \mathbf{1}/\sqrt{n}$ and every coordinate of $\boldsymbol{u}^j$ is $\pm 1/\sqrt{n}$, all quantities are strictly positive:

$$q_i^{j,s} \in \{(1-\alpha_n)/\sqrt{n}, (1+\alpha_n)/\sqrt{n}\}.$$

The signal about $\boldsymbol{D}$ is exact, $\widehat{\boldsymbol{D}} = \boldsymbol{D}$. The quantity signal is

$$\widehat{\boldsymbol{q}}^0 = \boldsymbol{q}^{j,s} + \boldsymbol{\varepsilon},$$

where, conditional on the state (j, s), the coordinates ε_i are i.i.d. $N(0, 1/n)$ and are independent of the matrix observation noise. Moreover,

$$\|\boldsymbol{q}^{j,s}\| = \sqrt{1+\alpha_n^2}, \qquad \|\boldsymbol{\varepsilon}\|^2 = (1/n)\chi_n^2.$$

Hence $\|\boldsymbol{\varepsilon}\| = O_{\mathrm{p},\boldsymbol{\Theta}}(\|\boldsymbol{q}^0\|)$.

The environment has significant structure. Take $b(n) = n/2$. Since $\lambda_n = -n + o(1)$ and all other eigenvalues have absolute value $\eta_n = o(1)$, the large-eigenvalue space is

$$\mathcal{L}(\boldsymbol{D}, b(n)) = \mathrm{span}\{\boldsymbol{u}^n\}$$

for all large n. The matrix observation noise is zero, so condition (2) of Definition 2 holds. Finally,

$$\left\| P_{\mathcal{L}(\boldsymbol{D},b(n))} \boldsymbol{q}^{j,s} \right\| = |\boldsymbol{u}^n \cdot \boldsymbol{q}^{j,s}| = 1 \geq \frac{1}{\sqrt{1+\alpha_n^2}} \|\boldsymbol{q}^{j,s}\|,$$

so condition (3) holds with any fixed $\delta < 1$ for all large n.

It remains to show that a fixed positive consumer-surplus incidence per dollar of expenditure cannot be robustly guaranteed. The intuition is that consumer surplus can be created only through a hidden direction $\boldsymbol{u}^j$, whose sign the quantity signal cannot resolve. We make this precise in two steps: at any realized signal, whatever intervention the rule selects, the hidden direction can take the sign that opposes it, producing a state in which the rule fails—generating either no positive expenditure or a consumer-surplus incidence per dollar below the target level. The two such candidate states differ only in this sign and are statistically almost indistinguishable, so the rule fails with probability bounded away from zero.

Fix any $\kappa \in (0, 1/2)$.[35] Say a rule *fails* in a state if its expenditure there is nonpositive, or is positive with $\dot{C} < \kappa \dot{S}$. Consider an arbitrary rule and a realized signal, and let $\boldsymbol{\sigma}$ be the selected intervention, with components

$$a := \boldsymbol{u}^n \cdot \boldsymbol{\sigma} \qquad \text{and} \qquad b_j := \boldsymbol{u}^j \cdot \boldsymbol{\sigma}$$

along the recoverable direction $\boldsymbol{u}^n$ and a hidden direction $\boldsymbol{u}^j$. Since $\boldsymbol{q}^{j,s}$ is orthogonal to every other eigendirection, Lemmas 1 and 2 give

$$\dot{S}_{\boldsymbol{\sigma}}^{j,s} = a + s\alpha_n b_j, \qquad \dot{C}_{\boldsymbol{\sigma}}^{j,s} = \frac{a}{1+|\lambda_n|} + \frac{s\alpha_n b_j}{1+\eta_n}.$$

[35] This proves the result for every $\kappa > 0$: if $\kappa \geq 1/2$, failure to robustly achieve the inequality for any fixed $\kappa' \in (0, 1/2)$ implies failure to robustly achieve the stronger inequality for κ.

Fix $j \le n-1$ and pick s with $sb_j \le 0$, so the hidden direction opposes the intervention: $s\alpha_n b_j = -\alpha_n|b_j|$. Then $\dot{S}^{j,s}_{\boldsymbol{\sigma}} = a - \alpha_n|b_j|$, and since $\alpha_n|b_j| \ge 0$ and $(1+\eta_n)^{-1} \ge (1+|\lambda_n|)^{-1}$,

$$\dot{C}^{j,s}_{\boldsymbol{\sigma}} = \frac{a}{1+|\lambda_n|} - \frac{\alpha_n|b_j|}{1+\eta_n} \le \frac{a-\alpha_n|b_j|}{1+|\lambda_n|} = \frac{\dot{S}^{j,s}_{\boldsymbol{\sigma}}}{1+|\lambda_n|}.$$

If $\dot{S}^{j,s}_{\boldsymbol{\sigma}} \le 0$, the rule delivers no positive expenditure; if $\dot{S}^{j,s}_{\boldsymbol{\sigma}} > 0$, then $\dot{C}^{j,s}_{\boldsymbol{\sigma}}/\dot{S}^{j,s}_{\boldsymbol{\sigma}} \le 1/(1+|\lambda_n|) < \kappa$ for all large n, since $|\lambda_n| \to \infty$. Either way the rule fails in state (j,s). Writing $F_{j,s}$ for the set of signals at which the rule fails in state (j,s), we have shown that

$$F_{j,1} \cup F_{j,-1} = \widehat{\boldsymbol{\Theta}} \qquad \text{for all large } n.$$

We now translate this pointwise ambiguity into the i.i.d. signal experiment. In state (j,s) the quantity signal is distributed as $P_{j,s} = N(\boldsymbol{u}^n + s\alpha_n \boldsymbol{u}^j,\, n^{-1}\boldsymbol{I})$, so

$$D_{\mathrm{KL}}(P_{j,1} \,\|\, P_{j,-1}) = \frac{1}{2}(2\alpha_n \boldsymbol{u}^j)^{\mathsf{T}}(n\boldsymbol{I})(2\alpha_n \boldsymbol{u}^j) = 2n\alpha_n^2 = \frac{2}{n},$$

and Pinsker's inequality gives $\|P_{j,1} - P_{j,-1}\|_{\mathrm{TV}} \le n^{-1/2}$. Because $F_{j,1} \cup F_{j,-1} = \widehat{\boldsymbol{\Theta}}$,

$$P_{j,1}(F_{j,1}) + P_{j,-1}(F_{j,-1}) \ge 1 - \|P_{j,1} - P_{j,-1}\|_{\mathrm{TV}} \ge 1 - n^{-1/2}.$$

Hence some state in the pair $\{(j,1),(j,-1)\}$ has failure probability at least $(1 - n^{-1/2})/2$, which is bounded away from zero. Therefore, no intervention rule can robustly guarantee any fixed positive consumer-surplus incidence per dollar of expenditure.

B.3. **Expected-performance implication.** We record the expected-performance statement used in the remark following Proposition 3. Consider the environment from part (1). The signal is deterministic and common across the hidden states. Thus any rule chooses the same intervention $\boldsymbol{\sigma}$ in every hidden state. If the rule selects positive expenditure, the homogeneity of the criterion lets us normalize $\dot{S}_{\boldsymbol{\sigma}} = \boldsymbol{q}^0 \cdot \boldsymbol{\sigma} = 1$. The averaging calculation above gives

$$\frac{1}{z_n}\sum_{j=1}^{z_n} \dot{W}^{(j)}_{\boldsymbol{\sigma}} = \tau_n, \qquad \tau_n \to 0.$$

Therefore some state j satisfies $\dot{W}^{(j)}_{\boldsymbol{\sigma}} \le \tau_n$. Since the signal is deterministic in this state,

$$\mathbb{E}_{\varphi_{\boldsymbol{\theta}^{(j)}}(n)}\left[\dot{C}_{\boldsymbol{R}(\widehat{\boldsymbol{\theta}})} + \dot{P}_{\boldsymbol{R}(\widehat{\boldsymbol{\theta}})} - \dot{S}_{\boldsymbol{R}(\widehat{\boldsymbol{\theta}})}\right] = \dot{W}^{(j)}_{\boldsymbol{\sigma}} \le \tau_n.$$

Since $\dot{W} = \dot{C} + \dot{P} - \dot{S}$, it follows that

$$\inf_{\boldsymbol{\theta} \in \boldsymbol{\Theta}(n)} \mathbb{E}_{\varphi_{\boldsymbol{\theta}}(n)}\left[\dot{C} + \dot{P} - \dot{S}\right] \le \tau_n.$$

For every fixed $\varepsilon > 0$, $\tau_n \le \varepsilon$ for all sufficiently large n, which proves the displayed claim in the remark.

APPENDIX C. HADAMARD BASIS AND SIMULATION CONSTRUCTION

This appendix formalizes the Hadamard-basis construction used in the Monte Carlo illustration in Section 7. The goal is to define an orthonormal basis with entries $\pm 1/\sqrt{n}$, build a Slutsky matrix with a single "spike" eigenvalue, and obtain a transparent block structure once goods are ordered by the spike eigenvector. This makes it easy to separate signal strength from noise and to interpret the welfare incidence plots.

Hadamard basis. We work along the subsequence $n = 2^m$ so that an $n \times n$ Hadamard matrix exists. Define the Sylvester recursion

$$\boldsymbol{H}_1 = \begin{pmatrix} 1 \end{pmatrix}, \qquad \boldsymbol{H}_{2n} = \begin{pmatrix} \boldsymbol{H}_n & \boldsymbol{H}_n \\ \boldsymbol{H}_n & -\boldsymbol{H}_n \end{pmatrix}.$$

Then $\boldsymbol{H}_n \boldsymbol{H}_n^\top = n\,\boldsymbol{I}$ and the normalized columns $\boldsymbol{u}^\ell = \boldsymbol{H}_n[:, \ell]/\sqrt{n}$ form an orthonormal basis $\{\boldsymbol{u}^1, \ldots, \boldsymbol{u}^n\}$ with entries $\pm 1/\sqrt{n}$. We index so that the uniform vector is $\boldsymbol{u}^n = \mathbf{1}/\sqrt{n}$.

Eigenvalue spectrum and demand matrix. To match the construction in the main text, we take the spike direction to be $\boldsymbol{u}^1$ and choose parameters $\lambda > 0$ and $\varepsilon > 0$. Define eigenvalues

$$\mu_n = -\varepsilon, \qquad \mu_1 = -(1+\lambda), \qquad \mu_{\text{bulk}} = \frac{-n + \varepsilon + (1+\lambda)}{n-2},$$

and set $\mu_\ell = \mu_{\text{bulk}}$ for all $\ell \notin \{1, n\}$. This guarantees $\sum_{\ell=1}^n \mu_\ell = -n$, so the diagonal entries of

$$\underline{\boldsymbol{D}} = \sum_{\ell=1}^{n} \mu_\ell\, \boldsymbol{u}^\ell (\boldsymbol{u}^\ell)^\top$$

equal -1 exactly (since $(u_i^\ell)^2 = 1/n$ for all i, ℓ), and $\underline{\boldsymbol{D}}$ is negative semidefinite. The spike eigenvalue $\mu_1 = -(1+\lambda)$ controls signal strength; the aggregate eigenvalue $\mu_n = -\varepsilon$ is near zero; and the bulk compensates to satisfy the trace condition.

Status quo quantities. Let $v^1 = \sqrt{n}\,\boldsymbol{u}^1 \in \{\pm 1\}^n$ and define

$$\underline{\boldsymbol{q}}^0 = \mathbf{1} + \eta\, v^1 = \sqrt{n}\,(\boldsymbol{u}^n + \eta\, \boldsymbol{u}^1),$$

so $q_i^0 \in [1-\eta, 1+\eta]$ and $\underline{\boldsymbol{q}}^0$ has a small but nonzero projection on the spike direction. This choice isolates the spike's contribution to welfare pass-through while keeping quantities strictly positive.

Block structure interpretation. Ordering goods by the sign of u_i^1 reveals the two-block structure of the outer product $\boldsymbol{u}^1(\boldsymbol{u}^1)^\top$: within-block entries are positive and across-block entries are negative. Multiplying by $\mu_1 < 0$ reverses those signs in the rank-one term $\mu_1 \boldsymbol{u}^1(\boldsymbol{u}^1)^\top$, so within-block entries are negative and across-block entries are positive. In other words, the Hadamard-basis construction is a two-block model up to a permutation of goods; sorting by the sign of u^1 makes the outer product exactly a 2×2 block matrix with constant entries $\pm 1/n$. When the recovered eigenvector $\widehat{\boldsymbol{u}}$ aligns with $\boldsymbol{u}^1$, the estimated outer product $\widehat{\boldsymbol{u}}\widehat{\boldsymbol{u}}^\top$ displays the same block pattern; below the recovery threshold it appears scrambled.

Signal and intervention. The authority observes $\widehat{\underline{\boldsymbol{D}}} = \underline{\boldsymbol{D}} + \underline{\boldsymbol{E}}$ with $\underline{\boldsymbol{E}} = \boldsymbol{E}$ a Wigner matrix (zero diagonal, i.i.d. $\mathcal{N}(0,1)$ off-diagonal). Let $\widehat{\boldsymbol{u}}$ be the leading

eigenvector of $-\widehat{\underline{\boldsymbol{D}}}$ and orient it so that $\widehat{\boldsymbol{u}} \cdot \underline{\boldsymbol{q}}^0 > 0$. The intervention

$$\underline{\boldsymbol{\sigma}} = s \cdot \frac{\widehat{\boldsymbol{u}}}{\widehat{\boldsymbol{u}} \cdot \underline{\boldsymbol{q}}^0}$$

spends s and yields welfare and incidence outcomes computed via the eigendecomposition formula in Equation (10). In this Wigner-noise simulation, the key threshold is the ratio $\lambda/\sqrt{n}$, which governs eigenvector recovery and therefore the transition in the welfare-incidence panels.

Appendix D. Sampling errors: Theoretical foundations and empirical tests

D.1. **An explicit sampling procedure.** The definition of significant structure (Definition 2), as well as the recoverable diagonal assumption (Assumption 4), impose conditions on the distribution of the error matrix $\underline{\boldsymbol{E}}$. We present a sampling procedure and an associated estimator for the normalized demand matrix $\underline{\boldsymbol{D}}$ that satisfies these conditions. As in Section 2.7, we do not assume any normalization to begin with (since the authority cannot assume the market is already normalized.)

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 D_{ij} are uniformly bounded, i.e., $|D_{ij}| \leq A < \infty$ for some positive constant A. Also recall that part (a) of Assumption 3 requires that the true diagonal entries satisfy uniform bounds: there exist universal constants $0 < d_{\min} \leq d_{\max} < \infty$ such that $-D_{ii} \in [d_{\min}, d_{\max}]$ for all i. The constants A, $d_{\min}$, and $d_{\max}$ 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 unbiased estimates of the relevant entries. For off-diagonal entries $i \neq j$, we write

$$\widehat{D}_{ij}^{h(i,j)} = D_{ij} + F_{ij}^{h(i,j)},$$

where the $F_{ij}^{h(i,j)}$ are mean-zero, independent across unordered pairs $\{i, j\}$, uniformly bounded, and satisfy $\mathrm{Var}(F_{ij}^{h(i,j)}) \leq \overline{V}$ for some constant $\overline{V}$ that does not depend on i, j, or n. Dropping the superscript for notational convenience, we define the corresponding pooled off-diagonal estimators by

$$\widehat{D}_{ij} = D_{ij} + F_{ij}, \qquad i \neq j,$$

where the F_{ij} inherit the same distributional properties.

Turning to estimates of D_{ii}, for each ordered pair (i, j) we write

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

We assume that the $G_{ij}^{h(i,j)}$ are mean-zero, independent across (i, j), uniformly bounded, and independent of $F_{i'j'}$. We also assume that the variances of the $G_{ij}^{h(i,j)}$ are bounded by $\overline{V}$, uniformly in i, j, and n. Averaging over j, we define

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

For notational convenience, we now define the diagonal entries of the matrix $\boldsymbol{F}$ by $F_{ii} := G_{ii}$. Thus we can write

$$\widehat{\boldsymbol{D}} = \boldsymbol{D} + \boldsymbol{F},$$

where $\boldsymbol{F}$ is symmetric with independent, mean-zero entries on and above the diagonal, uniformly bounded support, and

$$\mathrm{Var}(F_{ij}) \le \overline{V} \quad \text{for all } i, j, \qquad \mathrm{Var}(F_{ii}) \le \frac{\overline{V}}{n}.$$

For $i > j$ we impose symmetry by setting $F_{ij} = F_{ji}$.[36] It will be convenient below to keep using the notation G_{ii} in those parts of the argument where the averaging over j and the variance bound $\mathrm{Var}(G_{ii}) \le \overline{V}/n$ play a direct role.

Next, as in Section 2.7, let $\widehat{\boldsymbol{\Gamma}}$ be the diagonal matrix whose (i,i) diagonal entry is $\widehat{\Gamma}_{ii} = 1/\sqrt{-\widehat{D}_{ii}}$, and construct

$$\underline{\widehat{\boldsymbol{D}}} = \widehat{\boldsymbol{\Gamma}}\,\widehat{\boldsymbol{D}}\,\widehat{\boldsymbol{\Gamma}}.$$

Turning to true values (as opposed to hatted estimators), recall the analogous definition $\Gamma_{ii} = 1/\sqrt{-D_{ii}}$ and that

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

is the true normalized Slutsky matrix. Write $\underline{\boldsymbol{E}} = \underline{\widehat{\boldsymbol{D}}} - \underline{\boldsymbol{D}}$. By independence across j and the bound on $\mathrm{Var}(G_{ij}^{h(i,j)})$, we have

$$\mathrm{Var}(G_{ii}) \le \frac{\overline{V}}{n},$$

a fact we will use throughout.

Lemma 9. Under the sampling procedure described, the associated error matrix satisfies

$$\|\underline{\boldsymbol{E}}\| = O_{\mathrm{p}}(n^{1/2}) \quad \text{as } n \to \infty.$$

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

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

Recall from the sampling construction that $\widehat{\boldsymbol{D}} = \boldsymbol{D} + \boldsymbol{F}$, where $\boldsymbol{F}$ is symmetric with independent, mean-zero entries on and above the diagonal and diagonal entries $F_{ii} = G_{ii}$. Substituting $\widehat{\boldsymbol{\Gamma}} = \boldsymbol{\Gamma} + \boldsymbol{\Phi}$ and this decomposition of $\widehat{\boldsymbol{D}}$, we have:

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

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

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

[36] Since $\boldsymbol{D}$ is a symmetric matrix, we view the primitive parameters as being the entries F_{ij} on and above the diagonal, which are independent.

Our goal is to show that the spectral norm of $\underline{\boldsymbol{E}}$ grows at most at rate $\sqrt{n}$ in probability:

$$\|\underline{\boldsymbol{E}}\| = O_{\mathrm{p}}(\sqrt{n}).$$

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

$$\|\underline{\boldsymbol{E}}\| \leq \|\boldsymbol{\Gamma}\,\boldsymbol{F}\,\boldsymbol{\Gamma}\| + \|\boldsymbol{\Gamma}\,\boldsymbol{D}\,\boldsymbol{\Phi}\| + \|\boldsymbol{\Phi}\,\boldsymbol{D}\,\boldsymbol{\Gamma}\| + \|\boldsymbol{\Gamma}\,\boldsymbol{F}\,\boldsymbol{\Phi}\| + \|\boldsymbol{\Phi}\,\boldsymbol{F}\,\boldsymbol{\Gamma}\| + \|\boldsymbol{\Phi}\,\boldsymbol{D}\,\boldsymbol{\Phi}\| + \|\boldsymbol{\Phi}\,\boldsymbol{F}\,\boldsymbol{\Phi}\|.$$

We will bound the typical size of each term on the right-hand side. The terms $\boldsymbol{\Phi}\,\boldsymbol{D}\,\boldsymbol{\Gamma}$ and $\boldsymbol{\Gamma}\,\boldsymbol{D}\,\boldsymbol{\Phi}$ are transposes of each other, and similarly for $\boldsymbol{\Phi}\,\boldsymbol{F}\,\boldsymbol{\Gamma}$ and $\boldsymbol{\Gamma}\,\boldsymbol{F}\,\boldsymbol{\Phi}$, so their norms can be bounded in the same way; in what follows we bound one representative of each pair and absorb the resulting factor of 2 into the constants. 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: $\|\boldsymbol{\Gamma}\,\boldsymbol{F}\,\boldsymbol{\Gamma}\|$. Since $\boldsymbol{\Gamma}$ is a diagonal matrix with entries $\Gamma_{ii} = 1/\sqrt{-D_{ii}}$, and the $-D_{ii}$ are bounded away from zero and infinity by assumption, there exist constants $\Gamma_{\min}, \Gamma_{\max} > 0$ such that:

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

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

Consider the matrix $\boldsymbol{K} = \boldsymbol{\Gamma}\,\boldsymbol{F}\,\boldsymbol{\Gamma}$ with entries $K_{ij} = \Gamma_{ii} F_{ij} \Gamma_{jj}$. The variances of K_{ij} satisfy:

$$\mathrm{Var}(K_{ij}) = \Gamma_{ii}^2 \Gamma_{jj}^2 \,\mathrm{Var}(F_{ij}) \leq \Gamma_{\max}^4 \overline{V}.$$

By standard results on the spectral norm of symmetric random matrices with independent, mean-zero, uniformly bounded entries (see, for example, Vershynin, 2018; Bandeira and van Handel, 2016), there exists a constant $C_1 > 0$ such that

$$P\big(\|\boldsymbol{K}\| > C_1\sqrt{n}\big) \to 0 \quad \text{as } n \to \infty,$$

that is, $\|\boldsymbol{K}\| = O_{\mathrm{p}}(\sqrt{n})$.

Second term: $\|\boldsymbol{\Gamma}\,\boldsymbol{D}\,\boldsymbol{\Phi}\|$. Since $\boldsymbol{\Phi} = \widehat{\boldsymbol{\Gamma}} - \boldsymbol{\Gamma}$ and $\widehat{\Gamma}_{ii} = 1/\sqrt{-\widehat{D}_{ii}}$, we use a Taylor expansion around $-D_{ii}$ to approximate[37] 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:

$$\mathrm{Var}(\Phi_{ii}) \leq \left(\frac{1}{2}\Gamma_{\max}^3\right)^2 \mathrm{Var}(G_{ii}) \leq \frac{C}{n} \tag{26}$$

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

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

[37] Let $f(x) = x^{-1/2}$ for $x > 0$. The derivative is $f'(x) = -\frac{1}{2}x^{-3/2}$. Since $\widehat{D}_{ii} = D_{ii} + G_{ii}$, we have $-\widehat{D}_{ii} = (-D_{ii}) - G_{ii}$. Expanding $f((-D_{ii}) - G_{ii})$ around $x = -D_{ii}$ yields $f((-D_{ii}) - G_{ii}) = f(-D_{ii}) + f'(-D_{ii})(-G_{ii}) + O(G_{ii}^2)$. Therefore, $\Phi_{ii} = \widehat{\Gamma}_{ii} - \Gamma_{ii} = f(-\widehat{D}_{ii}) - f(-D_{ii}) = \frac{1}{2}\Gamma_{ii}^3 G_{ii} + O(G_{ii}^2)$, where $\Gamma_{ii} = f(-D_{ii}) = 1/\sqrt{-D_{ii}}$.

Therefore, the squared Frobenius norm is:

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

Taking expectations:

$$\mathbb{E}[\|\boldsymbol{\Gamma}\,\boldsymbol{D}\,\boldsymbol{\Phi}\|_F^2] \le \Gamma_{\max}^2 A^2 \sum_{i,j} \mathbb{E}[\Phi_{jj}^2] \le C n^2 \cdot \frac{1}{n} = Cn,$$

where in the penultimate step we have used Equation (26). Using the fact that the Frobenius norm bounds the spectral norm and applying Markov's inequality, we obtain, for any $K > 0$,

$$P\big(\|\boldsymbol{\Gamma}\,\boldsymbol{D}\,\boldsymbol{\Phi}\| > K\sqrt{n}\big) \le P\big(\|\boldsymbol{\Gamma}\,\boldsymbol{D}\,\boldsymbol{\Phi}\|_F > K\sqrt{n}\big) \le \frac{\mathbb{E}[\|\boldsymbol{\Gamma}\,\boldsymbol{D}\,\boldsymbol{\Phi}\|_F^2]}{K^2 n} \le \frac{C}{K^2},$$

uniformly in n, which is exactly the definition of $\|\boldsymbol{\Gamma}\,\boldsymbol{D}\,\boldsymbol{\Phi}\| = O_{\mathrm{p}}(\sqrt{n})$.

Third term: $\|\boldsymbol{\Gamma}\,\boldsymbol{F}\,\boldsymbol{\Phi}\|$. The entries of $\boldsymbol{\Gamma}\,\boldsymbol{F}\,\boldsymbol{\Phi}$ are:

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

The squared Frobenius norm is:

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

Taking expectations and using the Cauchy–Schwarz inequality:

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

Therefore, by the same argument as for the second term (using the bound on $\mathbb{E}[\|\boldsymbol{\Gamma}\,\boldsymbol{F}\,\boldsymbol{\Phi}\|_F^2]$ and Markov's inequality), we obtain $\|\boldsymbol{\Gamma}\,\boldsymbol{F}\,\boldsymbol{\Phi}\| = O_{\mathrm{p}}(\sqrt{n})$.

Fourth term: $\|\boldsymbol{\Phi}\,\boldsymbol{D}\,\boldsymbol{\Phi}\|$. The entries of $\boldsymbol{\Phi}\,\boldsymbol{D}\,\boldsymbol{\Phi}$ are:

$$(\boldsymbol{\Phi}\,\boldsymbol{D}\,\boldsymbol{\Phi})_{ij} = \Phi_{ii} D_{ij} \Phi_{jj}.$$

The squared Frobenius norm is:

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

Taking expectations:

$$\mathbb{E}[\|\boldsymbol{\Phi}\,\boldsymbol{D}\,\boldsymbol{\Phi}\|_F^2] \le A^2 \sum_{i,j} \mathbb{E}[\Phi_{ii}^2]\mathbb{E}[\Phi_{jj}^2] \le C n^2 \left(\frac{C}{n}\right)^2 = C.$$

By Markov's inequality, this implies $\|\boldsymbol{\Phi}\,\boldsymbol{D}\,\boldsymbol{\Phi}\| = O_{\mathrm{p}}(1)$.

Fifth term: $\|\boldsymbol{\Phi}\,\boldsymbol{F}\,\boldsymbol{\Phi}\|$. The entries of $\boldsymbol{\Phi}\,\boldsymbol{F}\,\boldsymbol{\Phi}$ are:

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

The squared Frobenius norm is:

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

Taking expectations and using the Cauchy–Schwarz inequality,

$$\mathbb{E}[\|\boldsymbol{\Phi}\,\boldsymbol{F}\,\boldsymbol{\Phi}\|_F^2] \le \overline{V}\sum_{i\neq j}\mathbb{E}[\Phi_{ii}^2]\mathbb{E}[\Phi_{jj}^2] + \sum_i \mathbb{E}[\boldsymbol{\Phi}_{ii}^4] = \overline{V}\Big(\sum_i \mathbb{E}[\Phi_{ii}^2]\Big)^2 \le C.$$

By Markov's inequality, this implies $\|\boldsymbol{\Phi}\,\boldsymbol{F}\,\boldsymbol{\Phi}\| = O_{\mathrm{p}}(1)$.

Diagonal entries and the recoverable diagonal condition. We finally verify that, in this sampling scheme, the diagonal of the true Slutsky matrix is recoverable from $\widehat{\boldsymbol{D}}$. Define the recovery map by $\Delta_n(\boldsymbol{M})_i := M_{ii}$. Then

$$\Delta_n(\widehat{\boldsymbol{D}})_i - D_{ii} = \widehat{D}_{ii} - D_{ii} = G_{ii} = \frac{1}{n}\sum_j G_{ij}^{h(i,j)}.$$

Each G_{ii} is therefore an average of independent, mean-zero, uniformly bounded random variables. Therefore $\max_i |\Delta_n(\widehat{\boldsymbol{D}})_i - D_{ii}| = o_{\mathrm{p}}(1)$, which is exactly the recoverable diagonal condition in this setting. □

This construction therefore provides a sampling scheme under which the associated error matrix satisfies $\|\underline{\boldsymbol{E}}\| = O_{\mathrm{p}}(\sqrt{n})$. In particular, in any environment where $\|\underline{\boldsymbol{D}}\|$ grows faster than $n^{1/2}$ (as in the block-model example in Section 5.2.1), condition (2) of Definition 2 is satisfied. Furthermore, the final paragraph of the proof shows that the diagonal of the true Slutsky matrix is consistently recoverable from $\widehat{\boldsymbol{D}}$, so the recoverable diagonal condition (Assumption 4) also holds under this sampling scheme.

D.2. **Bootstrap test for recoverable structure.** This subsection formalizes a bootstrap test for significant structure suggested by our theory.

Setup and estimand. Fix n and a market state $\boldsymbol{\theta} = (\boldsymbol{D}, \boldsymbol{q}^0)$. Suppose the analyst can form $T \ge 2$ independent estimates of the market state by sampling households in disjoint subsamples:

$$\widehat{\boldsymbol{D}}^{(t)} = \boldsymbol{D} + \boldsymbol{E}^{(t)}, \qquad \widehat{\boldsymbol{q}}^{0,(t)} = \boldsymbol{q}^0 + \boldsymbol{\varepsilon}^{(t)}, \quad t = 1,\dots,T,$$

with $(\boldsymbol{E}^{(t)}, \boldsymbol{\varepsilon}^{(t)})$ independent across t conditional on $\boldsymbol{\theta}$. Let $\underline{\widehat{\boldsymbol{D}}}^{(t)}$ and $\underline{\widehat{\boldsymbol{q}}}^{0,(t)}$ denote the corresponding normalized objects (as in Section 2.7), and define the estimated major eigenspace

$$\widehat{\mathcal{L}}^{(t)} := \mathcal{L}\big(\underline{\widehat{\boldsymbol{D}}}^{(t)}, \widehat{M}(n)\big),$$

where $b(n)$ is the threshold from Definition 2 and $\widehat{M}(n) := \frac{3}{4}b(n)$ (as in the proof of Theorem 1). Let $\underline{\widehat{\boldsymbol{\sigma}}}^{(t)}$ be the intervention implied by our rule when fed the t-th estimate: project $\underline{\widehat{\boldsymbol{q}}}^{0,(t)}$ onto $\widehat{\mathcal{L}}^{(t)}$, orient the sign so that spending is positive, and rescale to hit target spending $s = 1$. For $t \neq t'$, define the directional predicted welfare effect $\widehat{\tau}^{(t\to t')}$ by Equation (11). Assume T is even and write $T = 2Z$. Pair the splits arbitrarily (say $2z-1$ with $2z$), and for each pair $z = 1,\dots,Z$ define

$$\widehat{\tau}^{(z)} := \frac{1}{2}\Big(\widehat{\tau}^{(2z-1\to 2z)} + \widehat{\tau}^{(2z\to 2z-1)}\Big).$$

(If T is odd, discard one split; this is immaterial asymptotically.) Define the estimand

$$\mu_n(\boldsymbol{\theta}) := \mathbb{E}\big[\widehat{\tau}^{(1)} \mid \boldsymbol{\theta}\big],$$

where $\widehat{\tau}^{(1)}$ denotes the generic pair statistic.

Null hypothesis: significant structure. For a tolerance $\epsilon \in (0,1)$, we take the null hypothesis to be that the economy has significant structure in the sense of Definition 2. Concretely, the null is

$$H_0^{\mathrm{SS}} : \ \boldsymbol{\theta} \text{ lies in an environment with significant structure (Definition 2).}$$

Under H_0^{SS} our main theorem (Theorem 1) implies that the intervention constructed from one split and evaluated on an independent split has predicted welfare per dollar close to one, i.e.,

$$\mu_n(\boldsymbol{\theta}) \to 1 \quad \text{uniformly over } \boldsymbol{\theta} \in \boldsymbol{\Theta}(n),$$

as $n \to \infty$. Thus, for any fixed $\epsilon \in (0,1)$ and all sufficiently large n, H_0^{SS} implies the inequality $\mu_n(\boldsymbol{\theta}) \geq 1 - \epsilon$. The bootstrap test below targets this implication.

A scaling normalization and a denominator bound. By Remark 3, we may normalize the underlying environment so that $\|\boldsymbol{q}^0\| = 1$. That, along with condition (1) of Definition 2, implies $\|\widehat{\boldsymbol{q}}^{0,(t)}\| = O_{\mathrm{p},\boldsymbol{\Theta}}(1)$. Moreover, Assumptions 3 and 4 imply that the diagonal normalization map used to pass from $\widehat{\boldsymbol{q}}^{0,(t)}$ to $\underline{\widehat{\boldsymbol{q}}}^{0,(t)}$ has operator norm uniformly $O_{\mathrm{p},\boldsymbol{\Theta}}(1)$. Hence $\|\underline{\widehat{\boldsymbol{q}}}^{0,(t)}\| = O_{\mathrm{p},\boldsymbol{\Theta}}(1)$. Write $\widehat{\boldsymbol{v}}^{(t)} := P_{\widehat{\mathcal{L}}^{(t)}}\, \underline{\widehat{\boldsymbol{q}}}^{0,(t)}$. Under the rule, the rescaling denominator satisfies

$$\widehat{\boldsymbol{v}}^{(t)} \cdot \underline{\widehat{\boldsymbol{q}}}^{0,(t)} = \underline{\widehat{\boldsymbol{q}}}^{0,(t)\top} P_{\widehat{\mathcal{L}}^{(t)}}\, \underline{\widehat{\boldsymbol{q}}}^{0,(t)} = \|\widehat{\boldsymbol{v}}^{(t)}\|^2.$$

Accordingly,

$$|\widehat{\tau}^{(t\to t')}| \leq \|\underline{\widehat{\boldsymbol{q}}}^{0,(t')}\|\, \|\underline{\widehat{\boldsymbol{\sigma}}}^{(t)}\| = \frac{\|\underline{\widehat{\boldsymbol{q}}}^{0,(t')}\|}{\|\widehat{\boldsymbol{v}}^{(t)}\|},$$

by Cauchy–Schwarz and because $\frac{|\widehat{\lambda}_\ell^{(t')}|}{1+|\widehat{\lambda}_\ell^{(t')}|} \leq 1$ for each ℓ. Therefore, whenever $\|\widehat{\boldsymbol{v}}^{(t)}\|$ is bounded away from zero, the directional cross-fit statistics $\widehat{\tau}^{(t\to t')}$ (hence also the pair statistics $\widehat{\tau}^{(z)}$) are $O_{\mathrm{p},\boldsymbol{\Theta}}(1)$. Under the significant-structure null, Lemma 4 and the proof of Lemma 5 imply that there exists a constant $\kappa > 0$ (depending only on the environment constants in Definition 2, not on $\boldsymbol{\theta}$) such that

$$\inf_{\boldsymbol{\theta} \in \boldsymbol{\Theta}(n)} \mathbb{P}\big(\|\widehat{\boldsymbol{v}}^{(t)}\| \geq \kappa \,\big|\, \boldsymbol{\theta}\big) \to 1 \quad \text{as } n \to \infty.$$

Bootstrap lower confidence bound and one-sided certification. Let $\overline{\tau} := \frac{1}{Z}\sum_{z=1}^{Z} \widehat{\tau}^{(z)}$. Define the nonparametric bootstrap as follows: conditional on the observed $\widehat{\tau}^{1:Z}$, draw weights $(W_1, \ldots, W_Z) \sim \mathrm{Multinomial}(Z; 1/Z, \ldots, 1/Z)$ and set $\overline{\tau}^* := \frac{1}{Z}\sum_{z=1}^{Z} W_z\, \widehat{\tau}^{(z)}$. Let $\widehat{Q}_{1-\alpha}$ be the $(1-\alpha)$ conditional quantile of $\overline{\tau}^* - \overline{\tau}$ given $\widehat{\tau}^{1:Z}$, and define the one-sided basic bootstrap lower confidence bound $L_{1-\alpha} := \overline{\tau} - \widehat{Q}_{1-\alpha}$. We use the flag $L_{1-\alpha} < 1 - \epsilon$: values $L_{1-\alpha} \geq 1 - \epsilon$ certify $\mu_n(\boldsymbol{\theta}) \geq 1 - \epsilon$ at one-sided confidence level $1 - \alpha$.

Proposition 5 (Conditional asymptotic coverage)**.** For any fixed $\alpha \in (0,1)$,

$$\mathbb{P}(\mu_n(\boldsymbol{\theta}) \geq L_{1-\alpha} \,|\, \boldsymbol{\theta}) \to 1 - \alpha \qquad \text{as } Z \to \infty.$$

Thus, $[L_{1-\alpha}, \infty)$ is an asymptotically valid one-sided $(1-\alpha)$ confidence set for $\mu_n(\boldsymbol{\theta})$. In particular, observing $L_{1-\alpha} \geq 1 - \epsilon$ provides one-sided $(1-\alpha)$ certification for $\mu_n(\boldsymbol{\theta}) \geq 1 - \epsilon$.

Proof sketch. Conditional on $\boldsymbol{\theta}$, the boundedness of all the random variables involved guarantees that a CLT is applicable to $\overline{\tau}$. Under the same moment condition, the nonparametric bootstrap for the sample mean is consistent (van der Vaart, 1998, Chapter 23): conditional on the data, $\sqrt{Z}(\overline{\tau}^* - \overline{\tau})$ converges in probability to the same Gaussian limit as $\sqrt{Z}(\overline{\tau} - \mu_n(\boldsymbol{\theta}))$. Therefore the conditional $(1-\alpha)$-quantile $\widehat{Q}_{1-\alpha}$ consistently estimates the $(1-\alpha)$-quantile of $\overline{\tau} - \mu_n(\boldsymbol{\theta})$, yielding the stated one-sided coverage. □

Appendix E. Microfoundation of block model of demand

Consider the standard quasilinear quadratic utility function:

$$U(\boldsymbol{q}) = \boldsymbol{a}^\top \boldsymbol{q} - \frac{1}{2}\boldsymbol{q}^\top \boldsymbol{H}\boldsymbol{q},$$

where $\boldsymbol{a} \in \mathbb{R}^n_+$ and $\boldsymbol{H}$ is an $n \times n$ symmetric matrix.

For a given price vector $\boldsymbol{p} \in \mathbb{R}^n_+$, the first-order condition for maximizing $U(\boldsymbol{q}) - \boldsymbol{p} \cdot \boldsymbol{q}$ with respect to $\boldsymbol{q}$ yields the demand function:

$$\boldsymbol{q}(\boldsymbol{p}) = \boldsymbol{H}^{-1}(\boldsymbol{a} - \boldsymbol{p}).$$

Differentiating it with respect to prices yields the Jacobian:

$$\boldsymbol{D} = \frac{\partial \boldsymbol{q}}{\partial \boldsymbol{p}} = -\boldsymbol{H}^{-1}.$$

Strict concavity of the utility of the representative consumer is equivalent to $\boldsymbol{D}$ being negative definite, i.e., all the eigenvalues of $\boldsymbol{D}$ are strictly negative. This is equivalent to all eigenvalues of $\boldsymbol{H}$ being strictly positive. Moreover, having an eigenvalue of $\boldsymbol{D}$ that grows large (in absolute value) with the market size requires having an eigenvalue of $\boldsymbol{H}$ that goes to 0 as the market grows large.

E.1. **Imposing block structure.** Assume now that $\boldsymbol{H}$ has K blocks each of size m and that:

$$H_{ij} = \begin{cases} z > 0, & i = j, \\ z_{\text{In}} > 0, & i \neq j \text{ and } \left\lfloor \frac{i-1}{m} \right\rfloor = \left\lfloor \frac{j-1}{m} \right\rfloor \quad \text{(same block)}, \\ z_{\text{Ex}} < 0, & \text{otherwise} \quad \text{(different blocks)}. \end{cases}$$

Hence, the marginal utility from consuming products is decreasing within the same category and increasing across categories. The matrix $\boldsymbol{D} = -\boldsymbol{H}^{-1}$ has three distinct eigenvalues:

$$\begin{aligned} \lambda_3 &= -\frac{1}{z - z_{\text{In}}} && \text{with multiplicity } K(m-1) \\ \lambda_2 &= -\frac{1}{z + z_{\text{In}}(m-1) - m z_{\text{Ex}}} && \text{with multiplicity } (K-1) \\ \lambda_1 &= -\frac{1}{z + z_{\text{In}}(m-1) + m(K-1) z_{\text{Ex}}} && \text{with multiplicity } 1 \end{aligned}$$

Note that $-\boldsymbol{H}$ is the Hessian of U. Hence each eigenvalue of $\boldsymbol{D} = -\boldsymbol{H}^{-1}$ is the reciprocal of an eigenvalue of the Hessian of U. Therefore, eigenvalue $1/\lambda_i$ measures the curvature of the utility function U when we change consumption in the direction of the associated eigenvector $\boldsymbol{u}^i$ of the Hessian of U. Concavity

requires that these eigenvalues are all negative. In particular, $\lambda_1 \leq 0$ ensures that the marginal utility of increasing quantities in the all-ones direction is decreasing. The corresponding eigenvalue of $\boldsymbol{H}$ in the all-ones direction is the sum of three terms:

- $z > 0$: The contribution to curvature of increasing consumption of product i.
- $(m-1)z_{\text{In}} > 0$: The contribution to curvature of increasing quantity of other products in the block.
- $m(K-1)z_{\text{Ex}} < 0$: The contribution to curvature of increasing quantity of products in other blocks.

Thus, while $\lambda_1 \leq 0$ is required to keep concavity in this direction, $1/\lambda_1 \to 0^-$ requires adding sufficient interblock interactions as we grow the market to flatten the utility function in this direction—so that, in the limit, the utility function scales linearly with quantity in this direction.

E.1.1. *From $\boldsymbol{D}$ induced by the representative consumer to $\boldsymbol{D}$ specified directly in Section 5.2.1.* In Section 5.2.1 we defined a Slutsky $\boldsymbol{D}$ associated with state $\boldsymbol{\theta}$ to be such that $D_{ii} = -1$, $D_{ij} = \alpha_{\boldsymbol{\theta}} < 0$ if i and j belong to different product categories, and $D_{ij} = \omega_{\boldsymbol{\theta}} > 0$ if i and j belong to the same category. Such a $\boldsymbol{D}$ has three distinct associated eigenvalues

$$\begin{array}{ll} \lambda_3(\boldsymbol{\theta}) = -1 - \omega_{\boldsymbol{\theta}} & \text{with multiplicity } K(m-1) \\ \lambda_2(\boldsymbol{\theta}) = \lambda_3(\boldsymbol{\theta}) + m(\omega_{\boldsymbol{\theta}} - \alpha_{\boldsymbol{\theta}}) & \text{with multiplicity } (K-1) \\ \lambda_1(\boldsymbol{\theta}) = \lambda_2(\boldsymbol{\theta}) + mK\alpha_{\boldsymbol{\theta}} & \text{with multiplicity } 1 \end{array}$$

So the equivalence of the $\boldsymbol{D}$ generated by the representative consumer and the example outlined in Section 5.2.1 requires defining, for each state $\boldsymbol{\theta}$, a triple $(z^*(\boldsymbol{\theta}), z^*_{\text{In}}(\boldsymbol{\theta}), z^*_{\text{Ex}}(\boldsymbol{\theta}))$ so that the three distinct eigenvalues in the two approaches coincide. This is equivalent to solving, state by state, the following system:

$$-1 = \omega_{\boldsymbol{\theta}} - \frac{1}{z^*(\boldsymbol{\theta}) - z^*_{\text{In}}(\boldsymbol{\theta})},$$

$$\omega_{\boldsymbol{\theta}} - \alpha_{\boldsymbol{\theta}} = \frac{z^*_{\text{In}}(\boldsymbol{\theta}) - z^*_{\text{Ex}}(\boldsymbol{\theta})}{[z^*(\boldsymbol{\theta}) - z^*_{\text{In}}(\boldsymbol{\theta}) + m(z^*_{\text{In}}(\boldsymbol{\theta}) - z^*_{\text{Ex}}(\boldsymbol{\theta}))](z^*(\boldsymbol{\theta}) - z^*_{\text{In}}(\boldsymbol{\theta}))},$$

$$\alpha_{\boldsymbol{\theta}} = \frac{z^*_{\text{Ex}}(\boldsymbol{\theta})}{[z^*(\boldsymbol{\theta}) - z^*_{\text{In}}(\boldsymbol{\theta}) + m(z^*_{\text{In}}(\boldsymbol{\theta}) - z^*_{\text{Ex}}(\boldsymbol{\theta}))][z^*(\boldsymbol{\theta}) - z^*_{\text{In}}(\boldsymbol{\theta}) + m(z^*_{\text{In}}(\boldsymbol{\theta}) - z^*_{\text{Ex}}(\boldsymbol{\theta})) + nz^*_{\text{Ex}}(\boldsymbol{\theta})]}.$$

E.2. **Equilibrium.** In what follows we fix the true market state $\boldsymbol{\theta}$. We want to illustrate that limiting economies where $|\lambda_1(\boldsymbol{\theta})|$ goes to infinity are consistent with well-behaved equilibrium outcomes. For this we consider the simple case where $a_i = a$ and $c_i = c$ for all i; this simplification, together with the block model structure, guarantees the existence of a symmetric equilibrium. Recall that demand is

$$\boldsymbol{q} \;=\; -\boldsymbol{D}\,(\boldsymbol{a} - \boldsymbol{p}).$$

Let $p_j = p$ for all $j \neq i$ and let p_i be the price of firm i. Then, using the definition of $\boldsymbol{D}$, firm i's demand is

$$q_i(p_i, p) = -\Big[-(a - p_i) + \omega_{\boldsymbol{\theta}} \sum_{\substack{j \text{ same block} \\ j \neq i}} (a - p) + \alpha_{\boldsymbol{\theta}} \sum_{k \text{ other blocks}} (a - p)\Big]$$

$$= -\Big[-(a - p_i) + \omega_{\boldsymbol{\theta}}\,(m - 1)(a - p) + \alpha_{\boldsymbol{\theta}}\, m(K - 1)(a - p)\Big].$$

Recall that $\lambda_1(\boldsymbol{\theta}) = -1 + \omega_{\boldsymbol{\theta}}(m - 1) + \alpha_{\boldsymbol{\theta}} m(K - 1)$. Hence,

$$q_i(p_i, p) = -\Big[-(a - p_i) + (\lambda_1(\boldsymbol{\theta}) + 1)\,(a - p)\Big]$$

$$= -\Big[p_i - p + \lambda_1(\boldsymbol{\theta})(a - p)\Big].$$

The profit function is $\pi_i(p_i, p) = (p_i - c)\, q_i(p_i, p)$. The first-order condition is:

$$\frac{\partial \pi_i(p_i, p)}{\partial p_i} = 0 \implies q_i(p_i, p) - (p_i - c) = 0.$$

Imposing symmetry $p_i = p$,

$$q_i(p, p) = |\lambda_1(\boldsymbol{\theta})|(a - p)$$

so equilibrium price and quantity are

$$p = \frac{|\lambda_1(\boldsymbol{\theta})|}{1 + |\lambda_1(\boldsymbol{\theta})|} a + \frac{c}{1 + |\lambda_1(\boldsymbol{\theta})|} \quad \text{and} \quad q = \frac{|\lambda_1(\boldsymbol{\theta})|}{1 + |\lambda_1(\boldsymbol{\theta})|}(a - c)$$

which depends on K only via $|\lambda_1(\boldsymbol{\theta})|$. Under the assumption of significant structure, $|\lambda_1(\boldsymbol{\theta})|$ goes to ∞ as n increases across states and so p goes to a and q goes to $a - c$.

Appendix F. Hedonic Models and Significant Structure

The (un-normalized) Slutsky matrix in Pellegrino (2021) is $-\boldsymbol{B}^{-1}$, where $\boldsymbol{B} = \boldsymbol{I} + \alpha(\boldsymbol{\Sigma} - \boldsymbol{I})$. The matrix $\boldsymbol{\Sigma}$ is positive semidefinite because it can be written as $\boldsymbol{V}^\top \boldsymbol{V}$, where the columns of $\boldsymbol{V}$ are the characteristic vectors of various products.

It follows from this that all eigenvalues of $\boldsymbol{B}$ are real numbers bounded below by $1 - \alpha$, and all eigenvalues of $-\boldsymbol{B}^{-1}$ are at most $1/(1 - \alpha)$ in magnitude. Pellegrino uses the value $\alpha = 0.12$, which prevents any eigenvalue from exceeding 1.14. Once we normalize to obtain $\underline{\boldsymbol{D}}$ the eigenvalues of the resulting matrix are somewhat different, but they can still be bounded by a constant by elaborating this argument. Numerically we see that the normalization makes little difference.