

A Market Interpretation of Treatment Effects*

by Robert Minton and Casey B. Mulligan[†]
November 2024

Abstract

Markets, likened to an invisible hand, often appear to contradict econometric assumptions that rule out spillovers of one person’s treatment on another’s outcomes. This paper provides a simple statistical framework highlighting that controls are indirectly affected by the treatment through the market. Further, the effect of the treatment on the treated reveals only part of the consequence for the treated of treating the entire market. When combined with economic theory, our framework leads to a new application of Marshall’s Laws of Derived Demand that relates econometric estimates of treatment effects in the marketplace to the substitution and scale effects of demand theory. We show how treatment-effect estimators can diverge – both in magnitude and direction – from the causal effects of treatment on the treated or counterfactual policies treating all market participants. The framework shows how the consequences of targeted treatments reveal the effects of marketwide treatments, and the role of market frictions in that inference. Examples from labor, public finance, economic geography, development, and the macro literature on the “missing intercept” are provided.

* Kevin Murphy’s Price Theory lectures and further recommendations strongly influenced the preparation of this paper. We thank seminar participants at the Federal Reserve Board, the University of Chicago and the participants of the 2024 Price Theory Summer Camp. We also appreciate comments from Jim Heckman, Jack MountJoy, Evan Munro, Allen Sanderson, Alex Torgovitsky, Giuseppe Forte, Josh Gross, João Pugliese, Alex Tordjman, Harald Uhlig.

[†] Affiliations and contact: Federal Reserve Board, Robert.j.minton@frb.gov and University of Chicago, c-mulligan@uchicago.edu, respectively. The views in this paper are those of the authors and do not represent those of the Board of Governors of the Federal Reserve System or the Federal Reserve System.

I. Introduction

Markets, likened to an invisible hand, often appear to contradict econometric assumptions that rule out spillovers of one person’s treatment on another’s outcomes. Beyond coordinating activities across the supply and demand sides of the market, the invisible hand also coordinates activities *within* the supply and demand sides of the market—even when participants do not directly interact. Spillovers are thus a hallmark of markets.

Our paper provides a model of market spillovers to investigate the relationship between treatment-effect estimators and counterfactual outcomes resulting from more broadly applied treatments. We show that traditional estimators do not measure the effect of scaling up treatments, even when adjusted for the spillover of the treatment on controls. Although not relying on maximization, the framework has the Hicks-Marshall Laws of derived demand as a special case in which price treatments have scale and substitution effects on quantity outcomes. We then formalize how effects of scaling up the treatment can be extrapolated from treatment-control comparisons in the presence of market frictions.

We begin by illustrating why the market spillovers resulting from a marketwide treatment are likely significant even when the spillovers from a targeted treatment are negligible by some metrics. In this example, market demand is price inelastic. Some suppliers experience a productivity treatment, while others are untreated. The additional output from the treated drives down the market price that all suppliers receive. The revenue growth of treated relative to untreated suppliers is positive, but the revenue effect of treating all suppliers is negative. In this sense, the treatment effect has the opposite sign, besides having the wrong magnitude.

We then develop a simple statistical framework that allows for market spillovers and distinguishes targeted treatments from marketwide treatments. Treatment-control comparisons, which we term “difference-in-differences” (hereafter, DiD), differ from meaningful effects such as the treatment effect on the treated (hereafter, ToT) and the effect of a marketwide treatment (hereafter, the “scale effect”). Indeed, the difference between scale and DiD can serve as a definition of market spillovers. We establish that knowledge of ToT and DiD is insufficient to compute the scale effect.

ToT and DiD differ by the spillover effect of the targeted treatment on the control group. The scale effect and the ToT differ by the counterfactual spillover effect on the treated of additionally treating all the controls. The former spillover effect can differ substantially from the latter, particularly when the treated group is “small” relative to the control group. Formally, ToT is a weighted average of the scale effect and DiD. The weight is unknown, although in many applications it is expected to be related to the fraction of the market that is treated.

ToT and DiD become less informative about the scale effect as the share of treated units falls to zero. In the limit of a zero treated share, ToT and DiD coincide, differing from the scale effect by an arbitrary degree. Research designs with a small treated share are underpowered for the purposes of estimating market spillovers.

We then use demand theory to derive additional restrictions on the relationship between DiD, ToT, and scale. Prices are the treatments and quantities are outcomes. Market spillovers become cross-price effects on demand. A first result is that, although the treated goods are potentially heterogeneous, Hick’s Composite Commodity theorem says that their aggregate can be treated as a single good due to the commonality of their treatment. This holds regardless of whether treated and controls are the same in the baseline, as is the goal of a randomized experiment, or have systematic baseline differences as in natural experiments. Second, DiD measures the degree of substitution in preferences between treatments and controls, regardless of the fraction of the market that is treated or the magnitude of market spillovers. Put differently, DiD contains no information about the scale effect, which is the degree of substitution with goods outside the market where treated and controls participate. Third, the Hicks-Marshall Laws of derived demand show ToT to be an expenditure-share weighted average of the scale effect and DiD.

Although we focus on market spillovers within the demand or supply-side of the market, another role of prices in equilibrium models is to equalize quantities supplied and demanded. This function is left implicit in our analysis because it is already familiar in econometrics, particularly regarding the simultaneous feedback between supply and demand schedules. However, we show that our results generalize to a simultaneous equilibrium analysis: the scale and DiD effects we identify are both eigenvalues of a matrix of direct and indirect treatment effects on the same side of the market. Therefore, any arithmetic operations on treatment-effects matrices – such as combining demand and supply effects in a simultaneous system – translate into the same scalar operations on their respective scale effects and DiDs.

Neither the statistical nor demand-theoretic frameworks by themselves indicate how the treatment effect of a targeted treatment provides any information about the consequences of a marketwide treatment. We show how they are related in models of market frictions. The first case completely segregates treatment and controls in unrelated markets. DiD and ToT are equivalent because control outcomes are unaffected by spillovers. However, if there are any untreated in the market with the treated, DiD and ToT are different from the scale effect, understood as the effect of treating the entire market where the targeted treatment was administered. However, the scale effect can be inferred from the outcomes of a targeted treatment if, additionally, outcomes are measured for some of the untreated participants in the market. We analyze a second case that lacks any rigid barrier between markets, but spillovers diminish with economic distance from the treated. By considering comparison groups at varying distances from the treated, DiD, ToT, and scale can each be inferred from the outcomes of a targeted treatment. In both cases, the scale effect is outside the range spanned by DiD and ToT.

Finally, we provide several additional examples in which acknowledging equilibrium effects profoundly changes the interpretation of DiD estimates. These include a variety of models with time and region fixed effects, a model of the effects of targeted wage subsidies, and a discussion of why the “missing intercept” literature in macroeconomics relates to the difference between ToT and DiD rather than their differences from the scale effect.

Section II covers our illustration of market spillovers. Section III covers our statistical framework and its economic interpretation in terms of Marshall’s laws of derived demand. Section

IV covers our guidance on how to recover scale effects. Section V covers additional examples. Section VI concludes.

I.A Related Literature

Econometric results on causal research designs, along with recent extensions in the literature, often rely on the assumption of “no spillovers.”¹ “Spillovers” and “peer effects” are treated in microeconometrics as advanced, albeit interesting, topics that primarily arise when there are “externalities” (Angrist & Pischke, 2008; Athey & Imbens, 2017). Attempts to relax this assumption entail structure on how treatment spills over to the controls (Manski, 1993)—a structure which could be economic or statistical. The statistical approach reviewed in Huber (2023) might allow spillovers from observations within predetermined clusters but not from observations outside those clusters (Sobel, 2006; Hong & Raudenbush, 2006; Hudgens & Halloran, 2008). It might also allow spillovers that are decreasing with network distance (Viviano et al. 2024) or geographical distance (Butts 2023). Our contribution aligns with the economic approach, which we view as lacking in the general frameworks more recently available in statistics. We provide closed-form results for interpreting quantity or price comparisons, showing how these estimates relate to broader treatment effects on the entire market. Our approach focuses on spillovers mediated by market forces as opposed to spillovers through externalities (such as the urban knowledge spillovers in Jacobs (1969) or spillovers of medical treatment in Miguel and Kremer (2004)).

The analysis of specific market-based spillovers is extensive and spans many fields. In urban economics, for example, Glaeser and Gottlieb (2009) assess the benefits of easy labor mobility across firms within cities. See also Banzhaf (2021). In labor economics, Monte, Redding, and Rossi-Hansberg (2018) find evidence that commuting is an important adjustment mechanism for localized labor demand shocks. Crépon et al. (2013) find that gains to unemployed job seekers of job placement assistance can be offset by displacement effects for those who did not receive the program. Cautioning against “inattention to the market consequences of the [programs evaluated],” Heckman, Lochner, and Taber (1999) provide an equilibrium model for evaluating both behavior and welfare effects of tuition subsidies and other public policies. Heckman, LaLonde, and Smith (1999) conclude that “the costs of ignoring indirect [equilibrium] effects may be substantial.”² In development economics, Egger et al. (2022) find that transfer payments in one village can affect outcomes in nearby villages. Cunha, De Giorgi, and Jayachandran (2015) and Muralidharan, Niehaus, and Sukhtankar (2017) also find that market spillovers are important. As discussed in our Section V, public economics acknowledges that the introduction of state-specific cigarette taxes may affect the wholesale price of cigarettes faced by all states, a broader market response not captured by analyses comparing retail price changes in different states.

¹ See, for example, de Chaisemartin and d’Haultfoeuille (2020); Goldsmith-Pinkham, Sorkin and Swift (2020); and Borusyak, Hull and Jaravel (2022). Sometimes “no spillovers” is called “no interference” or is wrapped into the broader notion of the “stable unit treatment value assumption.”

² See also Heckman and Pinto (2024), who cite the “simultaneous causality” inherent in market clearing as a reason why the “treatment-control paradigm” is too narrow.

Munro et al.'s (2021) observation that “the interference pattern produced by marketplace price effects is dense and simultaneously affects all units, so cluster- or sparsity-based methods are not applicable” is consistent with the view of market equilibrium taken in this paper. They refer to a “Global Treatment Effect (GTE)” which we link to the “scale effect” from price theory. Both our paper and theirs treat this as “a meaningful policy-relevant counterfactual of treating all individuals in the [market] compared to treating no individuals in the [market].” Their “average direct effect” is analogous to what we call the “difference-in-differences” estimator.

Our paper differs from Munro et al. by emphasizing the economic interpretation of treatment effects and enabling statistical analyses in observational settings of market spillovers when randomized experiments are unavailable. Tools from economic theory point to several results such as our eigenvalue characterization of various treatment-control comparisons as built on just two “fundamentals:” scale and substitution effects. They also point to the role of market frictions in using information from a targeted treatment to estimate a counterfactual scale effect. The economic-theory approach emphasizes that market spillovers are not a nuisance, but rather of intrinsic interest.

Muralidharan and Niehaus (2017) highlight the dual nature of spillovers, saying that “comparing outcomes for (randomly) treated and untreated neighbors will yield a doubly biased estimate of the average impacts of treating both, since it ‘nets out’ spillovers from the treated to the untreated and also fails to capture the effects of spillovers that would have occurred from the untreated to the treated had the former been treated as well” (p. 109). Spillovers are a reason that experimenters should consider conducting their experiments on a large scale, they say. Our findings call that advice into question if the spillovers operate through a frictionless market. Our approach agrees that the effect on the treated of a widespread treatment is closer to the scale effect than the effect of a targeted treatment. The problem is that the effects on controls also grow with the size of the treatment group. Without market frictions, treatment effects estimators can be worse than “biased” – they may contain no information about the scale effect, regardless of the size of the experiment.³

Sraer and Thesmar (2023) revisit the literature on the misallocation of production factors. In that context, they acknowledge that firm-specific treatments affect untreated firms through market prices. They define three potentially distinct effects that, in our terminology, correspond to *DiD*, *ToT*, and the effect of treating all firms. At the same time, they note that some meaningful outcomes are themselves differences among firms. With the commonly assumed log-linear policy functions, their firm differences are free of any indirect effects operating through the market. *DiD*, *ToT*, and scale coincide for such outcome measures. An estimate of one is a valid estimate of the other two.

A macroeconomics literature recognizes that cross-sectional outcome comparisons do not reveal the full effect of fiscal policies. For example, households not receiving a transfer are affected by the transfers going to other households. Wolf (2023) refers to the absence of market spillover from cross-sectional comparisons as a “missing intercept” problem. He uses a general

³ Muralidharan and Niehaus also point out that experimental results may reveal features of the effectiveness of the experimental organization. Our focus is on connections between treatment effects and the tastes and technology in the market being studied.

equilibrium model (that is, an equilibrium consists of multiple prices and quantities) to explore solutions. The point of our Section III is that equilibrium spillovers are relevant in microeconomics too, even in partial equilibrium (only one price equilibrating). We also note that the missing intercept relates to the difference between DiD and ToT as defined in microeconomics. Our setup further highlights the distinction between these two effects and the scale effect.

Drawing from market equilibrium analysis, we emphasize that the treated and untreated experience scale and substitution effects in different combinations.⁴ This analytical approach has parallels with Heckman and Vitlaycil's (2005) expression of estimators as combinations of "marginal treatment effects," each of which refers to a specific type of individual. To focus on the price theoretic components, this paper considers only limited heterogeneity, namely treated versus untreated and in-market versus out-of-market. It emphasizes market connections, with counterfactual treatment regimes understood as additional distinct combinations of scale and substitution effects.

II. An illustration of equilibrium spillovers

The industry in our illustration has many suppliers of a homogeneous product. The demand curve is $Q = D(P)$, where P is the output price and Q is industry-aggregate quantity. With suppliers producing a homogeneous product, consumers choose the lowest price, which is matched by any supplier that intends to sell any output. All production inputs are held fixed firm-by-firm, while a "treated" subset of the firms experiences an increase in productivity by a factor of $e^t > 1$. The remaining firms in the industry are referred to as "controls" because they are not "exposed" to the productivity treatment, except indirectly through the market. Although our setting is simplified for clarity, the illustration is not far from important examples such as farming, where capital and fertilizer improvements (i.e., productivity in this example) enable some farmers to produce more crops from a given amount of land than others.

Industry-level effects are illustrated in Figure 1a while firm-level effects are shown in Figure 1b. For the purposes of illustration, we assume that the treated and untreated have the same baseline output and revenue, which is why the first vertical line in Figure 1b simultaneously shows the supply of any firm in the baseline and untreated firms (if any) with has treated competitors. The second vertical line is the supply of any treated firm.

⁴ Much quantitative work in both micro- and macro-economics treats scale and substitution parameters as constants.

Figure 1a. Industry-level effects of productivity treatments

Each supplier provides a fixed input amount. Treatment increases firm productivity.

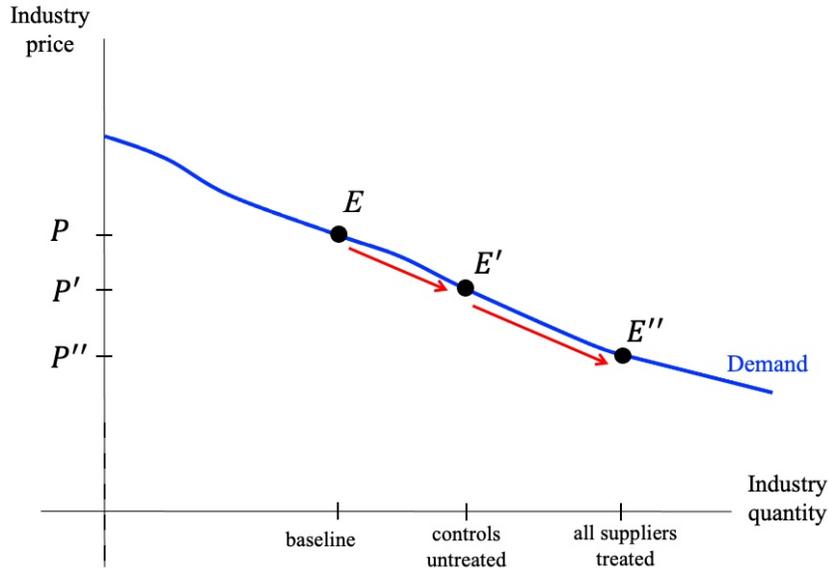
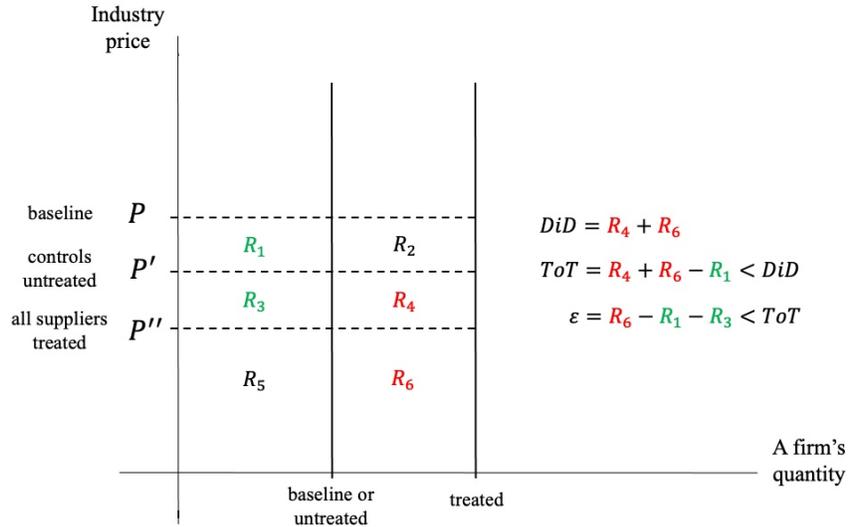


Figure 1b. Firm-revenue effects of productivity treatments

Each supplier provides a fixed input amount. Treatment increases firm productivity.



Note: DiD is the sum of the two red revenue amounts. The two green revenue amounts represent market spillovers.

Treating suppliers alters the market equilibrium. The point E in Figure 1a shows the baseline equilibrium. E' is the equilibrium when a subset of suppliers is treated. Because market supply has shifted outward, the market price is lower. Imagine using a DiD analysis at the firm level to study the effects of the productivity increase on revenues. The gap between the revenues of treated and untreated is represented by the areas $R_4 + R_6$ in Figure 1b. This treatment effect is also DiD (the difference between the average revenue change for treated suppliers and the average change for those untreated) because the two types of suppliers have the same revenue in the baseline. DiD is different from ToT , understood as the amount of revenue that the treated gain

relative to the baseline, because ToT includes the revenue loss (R_3 in magnitude) from the reduced price. Although the DiD for revenue must be positive, the ToT for revenue would be negative if price falls enough.

E'' is the equilibrium when the entire market is treated. Of course, E'' has a lower price than E' does because the productivity treatment is more widespread. The scale effect ε for revenue is defined as the difference between revenue under the marketwide treatment and revenue in the baseline. In Figure 1b, that is represented as the area difference $R_6 - R_1 - R_3$. ε is less than ToT and therefore even more likely to have the opposite sign as DiD . More precisely, ε and DiD have opposite signs if and only if market demand is price inelastic.

Through market competition, the price charged by any one supplier is largely determined by the productivity of competing suppliers. To put it another way, DiD “correctly” shows that supplier-specific productivity growth has little supplier-specific effect on price, but without clearly indicating the much larger price effects, and more negative revenue effects, of industry-wide productivity growth.

A statistician might say that the “control group is contaminated” because the productivity growth of the treated suppliers is “spilling over” to the untreated through competition for consumers. However, the divergence between DiD and ε is not mitigated by reducing the share treated. The empirical challenge is that equilibrium spillovers, such as the price effects illustrated in Figure 1a, are of intrinsic interest – they are part of the scale effect ε – but are differenced out by DiD and treatment-effect estimators. Equilibrium effects can exceed the direct effects, as they must in Figure 1a if demand is price inelastic and revenue is the outcome metric. Our purpose here is not to discourage treatment-control comparisons, even those with contaminated control groups, but rather to use price theory to help understand what they measure and how the findings can be applied to other settings.

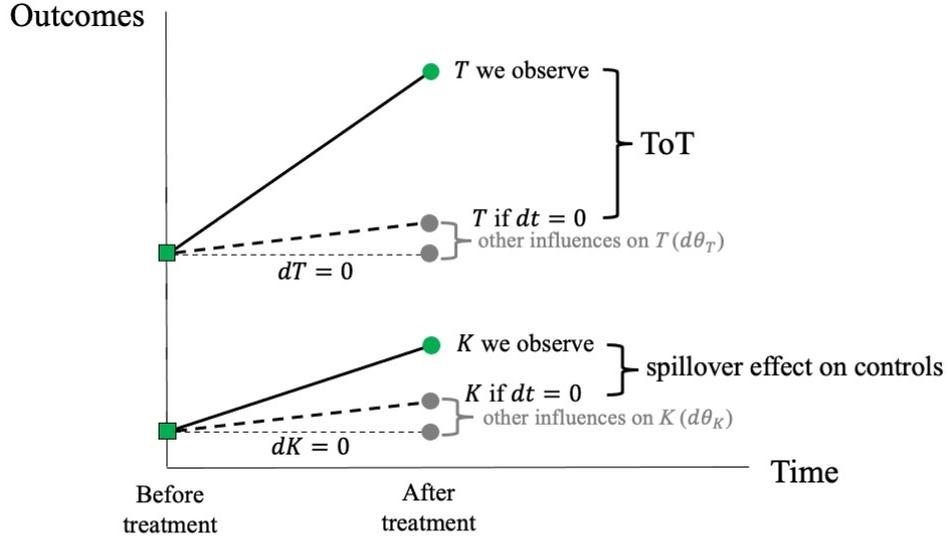
III. Treatments and controls according to Marshall’s laws

III.A. A vector representation of market spillovers

To begin the formal analysis, we consider a population of agents that are designated either as treatments or controls. Their population shares are denoted τ and $1-\tau$, respectively. Their outcomes are denoted T and K , respectively. The treatment, which directly affects the treated but not the controls, is denoted t . We let k denote a comparable shock that directly affects the controls but not the treated. The mappings from the two treatments to outcomes are denoted $T(t,k;\theta_T)$ and $K(t,k;\theta_K)$, where θ_T and θ_K denote other factors that influence outcomes for treated and controls.

Figure 2 illustrates, in the time dimension, the effects of a treatment $dt = 1$. The familiar parts of the diagram are that (i) other factors influence both the T and K outcomes over time and (ii) the effect of the treatment on the treated (ToT) is the difference between the final outcome for the treated and what that outcome would be without treatment. Given our emphasis on market connections between T and K , our Figure 2 also allows for an effect of $dt = 1$ on the outcome for the controls.

Figure 2. Treatments and Market Spillovers
in the time dimension



Notes: For the purposes of Figure 2, the controls are untreated ($dk = 0$) but in the same market. Observations reflect a treatment of $dt = 1$ directly affecting the treatment group.

To focus on equilibrium interpretations of DiD estimates, we maintain the parallel trends assumption that $d\theta_T = d\theta_K$ and have the same marginal effects on each K and T . We refer to this assumption as “parallel trends with respect to omitted variables” (PTOV).⁵ In other words, under PTOV and without any treatments ($dt = dk = 0$), the treated and control groups experience the same outcome changes $dT = dK$. PTOV requires that the heavy dashed lines in Figure 2 be parallel.

With PTOV, the relevant four first partial derivatives of the outcome mapping are represented as a two-by-two matrix S :

$$S = \begin{pmatrix} S_{Tt} & S_{Tk} \\ S_{Kt} & S_{Kk} \end{pmatrix} = \begin{pmatrix} \partial T / \partial t & \partial T / \partial k \\ \partial K / \partial t & \partial K / \partial k \end{pmatrix} \quad (1)$$

The first entry in S is the effect s_{Tt} of a unit treatment t on the treated, which is commonly known as ToT as labeled as such in Figures 1 and 2. The final entry s_{Kk} is the analog of ToT for the controls. The off-diagonal elements reflect spillovers, sometimes known as indirect effects of treatments. The spillover effect shown in Figure 2 is s_{Kt} .

S 's first column difference and first row sum are central to our interpretation of difference-in-differences. We therefore establish the following definitions:

$$DiD \equiv s_{Tt} - s_{Kt} \quad (2)$$

$$\varepsilon \equiv s_{Tt} + s_{Tk} \quad (3)$$

⁵ Up to sampling error, PTOV is guaranteed if treatments are randomly assigned.

Definition (2) is our representation of difference-in-differences (more literally, a difference in treatment derivatives) under the aforementioned parallel-trends assumption. *DiD* subtracts the effect, measured per unit t , of the treatment t on controls from its effect on the treated. The definition (3) refers to the “scale effect,” which is the effect on the treated group of applying the treatment uniformly across the entire population, or what we call “the entire market.” The scale effect ε is often the parameter of interest.

The case of no spillovers has S as a diagonal matrix ($s_{Tk} = 0 = s_{Kt}$), with no difference between *DiD* and ε or *DiD* and *ToT*. While not ruling out the zero-spillover case, the purpose of this paper is to link the off-diagonal elements to the diagonal elements and to results from price theory. More generally, the difference between the scale effect and *DiD* is the sum of the spillover elements of S : $\varepsilon - DiD = s_{Tk} + s_{Kt}$.⁶

Another restriction on the S matrix also resembles parallel trends and drives many of our results. Specifically, administering the treatment uniformly to both the treated group and control group should not affect the difference between their outcomes:

$$s_{Tt} + s_{Tk} = s_{Kt} + s_{Kk} \quad (4)$$

We refer to assumption (4) as “Parallel Trends for Parallel Treatments” (PTPT). At this point, PTOV is distinct from PTPT. Proposition 1 establishes that the familiar procedure of dividing differential outcome changes by a treatment differential yields *DiD* if and only if the PTPT assumption (4) holds.

PROPOSITION 1 (Differential and parallel treatments). Assume parallel trends for omitted variables (PTOV) and that dk is neither 0 nor equal to dt . Then the PTPT assumption (4) is equivalent to (5):

$$DiD = \frac{dT - dK}{dt - dk} \quad (5)$$

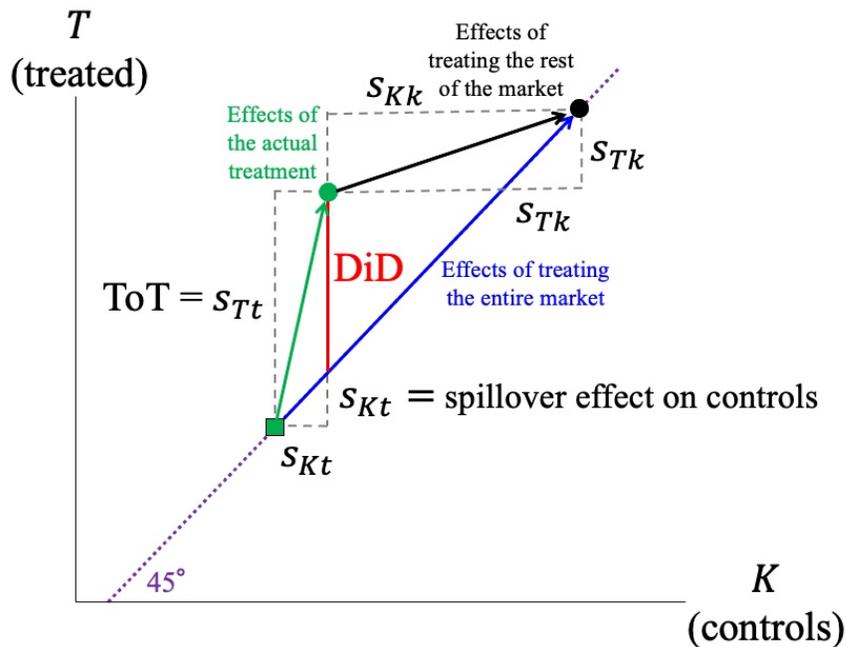
Proof. To obtain an expression for the numerator in (5), totally differentiate $T(t,k;\theta_t) - K(t,k;\theta_k)$. With PTOV eliminating the θ terms, the numerator is *DiD* ($dt-dk$) plus the product of dk and the difference between the LHS and RHS of (4). With $dk \neq 0$, the RHS of (5) differs from *DiD* if and only if equation (4) is satisfied. QED

A corollary to Proposition 1 is that, with PTPT, equation (5) corresponds to the *DiD* defined in (2) regardless of whether treatments are solely for the treatment group ($dt \neq 0 = dk$) or solely for the control group ($dt = 0 \neq dk$). If treatment and control groups are identical in the baseline, as is the design of randomized controlled trials, then equation (5) becomes the ratio of the outcome gap $T-K$ to the treatment gap $t-k$.

⁶ See also Munro et al.’s (2021) expression of a “global treatment effect” as the sum of “direct” and “indirect” treatment effects.

Figure 3 illustrates the model (1)-(4) for the case that $\varepsilon > DiD$, showing all four elements of S .⁷ The axes measure outcomes for controls and treatments. The square point represents the baseline, showing outcomes absent any treatment. The green vector is the first column of S , showing the effects on both groups of treating only the treated. As shown, that vector is not vertical but has a slope greater than 45 degrees, which indicates that t has a spillover effect, although one that is less than the direct effect on the treated group. Unsurprisingly, the DiD (red segment) measures the distance between the treatment effect and the 45-degree line.

Figure 3. Geometry of Market Spillovers
in the outcomes dimension, with the scale effect exceeding DiD



The black vector, which is the second column of S , shows the effect of subsequently treating the rest of the market. The sum of the two arrows follows the 45-degree line if and only if the PTPT assumption (4) holds. The vertical and horizontal dimension of their sum is ε .

Figure 4 illustrates a case with the same scale effect as Figure 3, but with DiD closer to zero. In contrast, Figure 5's case has the same DiD as Figure 3 but no scale effect (treating the controls “undoes” the effects of t).⁸ As such, it also shows an instance of $\varepsilon < DiD$.

⁷ It also shows $DiD > 0$, although for what follows the sign of each ε and DiD is less important than the sign of their difference.

⁸ The area of the triangle shown in Figures 2-4 is half of the magnitude of $DiD * \varepsilon$.

Figure 4. Market does not easily deviate from 45°

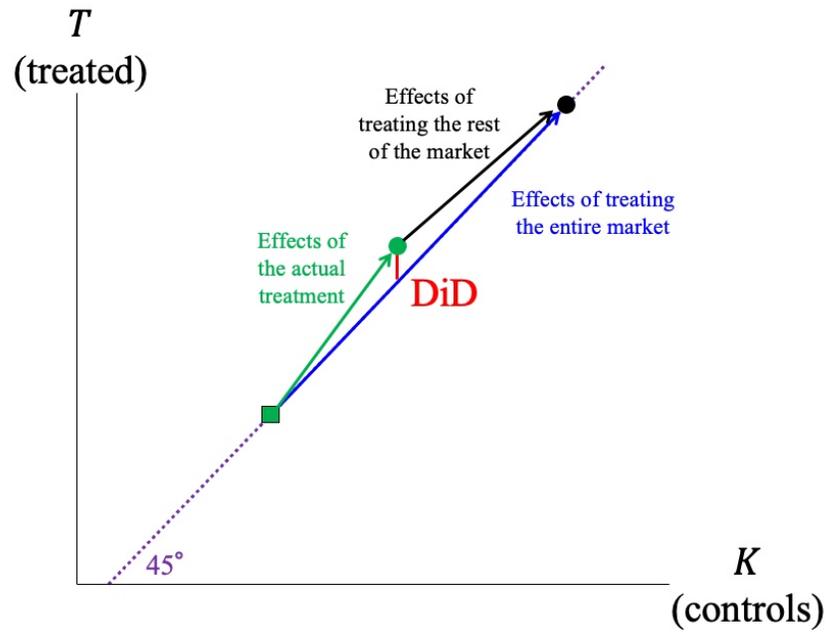
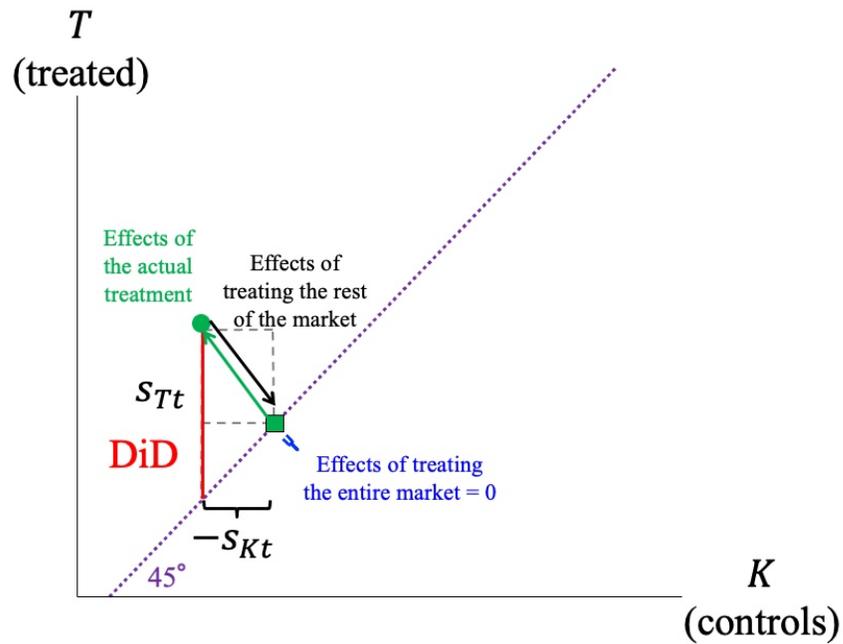


Figure 5. Scale effect = 0 \ll DiD



Let λ denote the share of the between-group spillover effects of a full market treatment $dt = dk$ that would be experienced by the control group.

$$\lambda \equiv \frac{S_{Kt}}{S_{Kt} + S_{Tk}} = \frac{S_{Kt}}{\varepsilon - DiD} \quad (6)$$

The symmetry of our discussion of treated and controls suggests that λ would be closely related to τ , if not identical to it, because a larger treatment group is expected to have a greater effect on the controls than treating a comparatively group would. However, until we say more about the units of K , T , k , and t (see subsection III.B and following), a precise relationship between λ and τ cannot be specified. Regardless, the common intuition that small-scale treatments ($\tau \approx 0$) have near-zero spillover effects on the controls can be represented by assuming ($\lambda \approx 0$).

PROPOSITION 2 (Treatment effects decomposition). If the PTPT assumption (4) holds, then the treatment effects matrix S can be written in terms of DiD , ε , and λ , as defined in (2), (3), and (6):

$$S = \begin{pmatrix} \lambda\varepsilon + (1 - \lambda)DiD & (1 - \lambda)(\varepsilon - DiD) \\ \lambda(\varepsilon - DiD) & (1 - \lambda)\varepsilon + \lambda DiD \end{pmatrix} \quad (7)$$

Proof. The share defined by (6) distributes the sum of the spillover terms, already established to be $\varepsilon - DiD$, between its two components as shown in (7). From (2), the S_{Ti} term is the sum $DiD + S_{Kt}$ and therefore what is shown in (7). PTPT requires that $S_{Kk} = DiD + S_{Tk}$, which is the result shown in (7). QED

The definitions and axiom (2)-(6) allow the diagonal of S to be expressed entirely in terms of weighted averages of DiD and ε , using λ and $1-\lambda$ as weights. The off-diagonal “spillover” elements are the difference between the scale effect and DiD , scaled by either λ or $1-\lambda$. The direction of the spillover effects can therefore be understood as a comparison between the scale effect and DiD . The expression (7) and the intuition about signing market spillovers are familiar from price theory, where they are known as Marshall’s Laws of Derived Demand. This connection is made explicit in section III.B.

Expression (7) also formalizes that, when λ is interpreted as the treated share of the market τ and is small, the ToT is approximately equivalent to DiD and is uninformative about the scale effect ε .

The eigenvalues of S are simple, of intrinsic interest, and useful for establishing additional results.

COROLLARY. Under the PTPT assumption (4), the eigenvalues of S are DiD and ε . The matrix sum (product) of two matrices each of the form (7) itself has the form (7), with one eigenvalue that is the sum (product) of the two component $DiDs$ and another eigenvalue that is the sum (product) of the two ε s, respectively.

Proof. Use (7), which requires (4), to calculate the eigenvalues. QED

Even though (so far) the matrix S has three degrees of freedom, its eigenvalues are independent of the spillover share λ . Each eigenvalue is of intrinsic interest because DiD is commonly measured while ε represents a meaningful counterfactual. As a result, any arithmetic operations on treatment-effects matrices – such as combining two different S matrices, one for each of the demand and supply effects in a simultaneous system – translate into the same operations on their respective scale effects and $DiDs$.

The inverse operation on S is particularly informative because the spillover-share parameter is retained:

$$\begin{pmatrix} \lambda\varepsilon + (1 - \lambda)DiD & (1 - \lambda)(\varepsilon - DiD) \\ \lambda(\varepsilon - DiD) & (1 - \lambda)\varepsilon + \lambda DiD \end{pmatrix}^{-1} = \begin{pmatrix} \frac{\lambda}{\varepsilon} + \frac{1 - \lambda}{DiD} & (1 - \lambda)\left(\frac{1}{\varepsilon} - \frac{1}{DiD}\right) \\ \lambda\left(\frac{1}{\varepsilon} - \frac{1}{DiD}\right) & \frac{1 - \lambda}{\varepsilon} + \frac{\lambda}{DiD} \end{pmatrix} \quad (8)$$

Suppose we have results for quantity outcomes of price treatments. Equation (8) translates those into results for price outcomes of quantity treatments.

III.B. Derived-demand interpretations of treatment effects

When the treatments are prices and outcomes are quantities, or vice versa, demand theory is another lens through which to compare DiD and scale effects as defined above. In contrast to the illustration featured in Figure 1, the demand analysis that follows allows for (but does not require) heterogeneous treatment effects at the individual level and imperfect substitution in demand among the various quantities. Our first result is that treated outcomes and control outcomes can be interpreted as composite commodities. This interpretation allows the DiD and compensated scale effects to be precisely linked to Allen elasticities of substitution in preferences. The magnitude of DiD proves to be a shadow elasticity of substitution. A final result is that compensated scale effect summarizes everything about substitution that cannot be discovered with DiD , and vice versa.

Consider an empirical procedure that identifies a group of G suppliers, selecting a fraction of them to be “treated” and using the remainder as a comparison group. The outcome of interest is equilibrium group-average quantity, expressed in logs. The components of each average are denoted as vectors \hat{T} and \hat{K} . T and K represent the per-capita log average quantities, respectively, supplied by the two groups and ultimately consumed by downstream consumers. The selection procedure could be random, although our main results do not require that T be representative of G overall. The quantities T and K could be labor employed in two different states, as in the minimum wage literature. They could be capital in two different industries. In other applications, T and K might represent sales of distinct retail products, sales or employment firms in the same industry that differ by size or location, or different sectors of the economy. The model is flexible in accommodating these cases.

The purpose of the comparison is to infer something about the characteristics of market demand. We refer to the “treatments” t and k as log prices on the demand side of the market and

refer to this interpretation as demand-price treatments with quantity outcomes. The dt and dk notation can be understood as a shorthand for a simultaneous-equation analysis of exactly how the T and K suppliers must be treated to result in demand price changes dt and dk . We expand those details later after we establish results connecting ε and DiD to consumer preferences.

Final consumers have convex preferences $u(\hat{T}, \hat{K}, \hat{X})$ over $G+N$ consumption choices, where N is the dimension of the vector \hat{X} of the “outside” goods supplied by neither the treated nor the controls. The t treatment increases prices of the components of \hat{T} by the common factor e^t . The k treatment applies to the components of \hat{K} , with a price factor of e^k . At this point, we do not necessarily assume that treated and controls are weakly separable from each other or the outside goods, which is why u has $G+N$ distinct arguments.

Composite Commodity Theorem. The only thing that the components of \hat{T} necessarily have in common with each other is the treatment t , and similarly for the components of \hat{K} . Nevertheless, that is enough to represent the treated with a composite commodity and the controls with a second composite commodity. As Sir John Hicks (1939/1975, p. p. 50) put it “when the relative prices of a group of commodities can be assumed to remain unchanged, they can be treated as a single commodity.” Formally, market responses to the treatment can be understood in terms of a three-dimensional utility function:

$$u(z_T, z_K, z_X; \hat{P}_T, \hat{P}_K, \hat{P}_X) \equiv \max_{\{\hat{T}, \hat{K}, \hat{X}\}} \hat{u}(\hat{T}, \hat{K}, \hat{X}) \quad s. t. \quad (9)$$

$$\hat{P}_T \cdot \hat{T} \leq z_T \wedge \hat{P}_K \cdot \hat{K} \leq z_K \wedge \hat{P}_X \cdot \hat{X} \leq z_X$$

where the scalars z_T , z_K and z_X in (9) are weighted sums of T , K , and X quantities, respectively, that would be demanded without any treatment.⁹ Note that the marginal rate of substitution between components of \hat{T} , between components of \hat{K} , or across those groups can depend on the value of \hat{X} . We include the price vectors $\{\hat{P}_T, \hat{P}_K, \hat{P}_X\}$ in the definition of u to be clear that those must be held constant in what follows, otherwise u becomes a different function of the three aggregates.

The scalars z_T and z_K represented weighted sums, whereas e^t and e^k are unweighted averages. According to demand theory, summing various quantities is of little economic meaning unless they have the same prices. Equivalently, the quantities of various goods are more meaningfully summed after each has units set so that they have the same price per unit. This is the principle of quantity indexes. Henceforth, we show results for the weighted sums and leave it as an exercise to the reader to determine how treatments affect $(\ln z_T - T)$ and $(\ln z_K - K)$ in specific applications.

Assuming that the treatments have no effect on relative prices among the treated and no effect on relative prices among the controls, the definition (9) can characterize market equilibrium responses to the treatments as shown in (10):

⁹ If there were more than the two treatment values t and k , then application of the composite commodity theorem results in more than two aggregates. Namely, a supplier would be aggregated with all others receiving the same treatment.

$$\max_{\{z_T, z_K, z_X\}} u(z_T, z_K, z_X; \hat{P}_T, \hat{P}_K, \hat{P}_X) \quad s. t. \quad (10)$$

$$e^t z_T + e^k z_K + z_X \leq M$$

The model (10) describes a prototype three-dimensional demand system. Its solution is three Marshallian demand functions for the three quantity indexes z_T , z_K and z_X , each as a function of income M and three prices, with the usual properties.¹⁰ It also defines three Hicksian demand functions that are related to the Marshallian system by the usual Slutsky correspondence.

In this setting, PTPT requires that (a) both z aggregates respond in the same proportion to any common treatment $dk = dt$, and (b) each aggregate's response is equally composed between substitution and income effects. These conditions are satisfied if, for example, utility over (z_T, z_K, z_X) can equivalently be written as utility over z_X and a homothetic aggregator of z_T and z_K , a common weak-separability assumption. These require that both aggregates have the same income elasticity. By definition, for any nonzero dt , PTPT requires Hicksian and Marshallian demand functions to satisfy:

$$\left. \frac{d \ln z_T}{dt} \right|_{dk=dt, d\hat{P}=0} = \left. \frac{d \ln z_K}{dt} \right|_{dk=dt, d\hat{P}=0} \quad (11)$$

where $d\hat{P} = 0$ is our shorthand for holding constant all three \hat{P} vectors.¹¹ In summary, this framework restricts preferences with PTPT (equation (11)) while restricting price treatments to be common among the treated and common among the controls.

Let η_{TK}^H denote compensated elasticity of z_T with respect to the price of z_K and likewise for any other pair of $\{T, K, X\}$. PTPT (equation (11)) requires the substitution matrix to satisfy $\eta_{TT}^H + \eta_{TK}^H = \eta_{KT}^H + \eta_{KK}^H$ in addition to the usual homogeneity and, in levels or substitution elasticities, symmetry conditions. Let ω_T , ω_K , and ω_X denote the expenditure shares corresponding to each of the three expenditure terms in the model (10)'s budget constraint, respectively. The shares sum to one across T , K , and X .

The off-diagonal terms in equation (1) are the spillovers. For the demand model (10) and (11), those terms have the economic interpretation of cross-price elasticities η_{TK}^H and η_{KT}^H , respectively. As already noted, the weight of scale in ToT is the share λ of the spillovers that is contamination of the controls. In the demand model that share is $\lambda = \eta_{KT}^H / (\eta_{TK}^H + \eta_{KT}^H)$. By Hicksian symmetry, that share can also be written as $\omega_T / (\omega_T + \omega_K)$. In words, a Hicksian ToT is the expenditure-share weighted average of scale ε and DiD , whose determinants are further revealed in the propositions that follow:

¹⁰ The same solution could be obtained by maximizing $\hat{u}(\hat{T}, \hat{K}, \hat{X})$ with respect to the $G+N$ quantities subject to the budget constraint $e^t \hat{P}_T \cdot \hat{T} + e^k \hat{P}_K \cdot \hat{K} + \hat{P}_X \cdot \hat{X} \leq M$, and then using the price vectors $\{\hat{P}_T, \hat{P}_K, \hat{P}_X\}$ to form the quantity indexes. See also Deaton and Muellbauer (1980, p. 121), although our paper presents much work in the treatment-control paradigm as counterexamples to their assumption that “the usefulness of this theorem in constructing commodity groupings for empirical analysis is likely to be somewhat limited.”

¹¹ Recall that the price vector for the T (K) goods is $e^t \hat{P}_T$ ($e^k \hat{P}_K$), respectively.

$$ToT = \frac{\omega_T}{\omega_T + \omega_K} \varepsilon + \frac{\omega_K}{\omega_T + \omega_K} DiD \quad (12)$$

The components of \hat{T} can have different income elasticities, although Engel aggregation of (9) requires that the share-weighted average of their income elasticities coincides with the common income elasticity of z_T and z_K . The treatment effect for treated good i is therefore the product of $\partial \ln z_T / \partial t$ and the allocation of z_T among its components based their relative income elasticities and substitution patterns with the outside good. In other words, the model (9) - (10) allows for rich heterogeneity in treatment effects, although such heterogeneity is not required. The treated goods can also differ in terms of the rate at which they substitute for various control-group goods or with outside goods because these differences are, beyond the income elasticity ratio already noted, irrelevant for predicting responses to t or k .

Let AES_{TK} denote the Allen elasticity of substitution between z_T and z_K , which is η_{TK}^H / s_K . Defining likewise for all other pairs of $\{T, K, X\}$, we have notation for the symmetric 3x3 matrix of Allen elasticities of substitution (AES). The shadow elasticity of substitution between z_T and z_K , as defined by McFadden (1963), is defined in terms of the AES and expenditure shares and denoted σ_{TK} :

$$\sigma_{TK} \equiv -\frac{\omega_T \omega_K}{\omega_T + \omega_K} [(AES_{TT} - AES_{TK}) + (AES_{KK} - AES_{TK})] > 0 \quad (13)$$

Proposition 3 links the treatment and scale effects to these two concepts of substitution from demand theory:

PROPOSITION 3 (Treatment and scale effects linked to Allen substitution). With PTPT, the comparative statics for the Hicksian demand functions can be expressed in terms of expenditure shares and Allen elasticities of substitution (AES):¹²

$$\left. \frac{d \ln z_T / z_K}{dt - dk} \right|_{d\hat{P}=0=dt-dk} = (AES_{TT} - AES_{TK})\omega_T = -\sigma_{TK} < 0 \quad (14)$$

$$\varepsilon^H \equiv \left. \frac{d \ln z_T}{dt} \right|_{d\hat{P}=0=dt-dk=du} = \frac{\omega_X^2}{\omega_T + \omega_K} AES_{XX} < 0 \quad (15)$$

Proof. In elasticity form, the differentials of the Hicksian demand functions are:

$$d \ln z_T = \omega_T AES_{TT} dt + \omega_K AES_{TK} dk \wedge d \ln z_K = \omega_T AES_{TK} dt + \omega_K AES_{KK} dk \quad (16)$$

¹² With \hat{P}_X held constant, $\hat{P}_X \cdot \hat{X}$ is itself a composite commodity with well-defined own- and cross-price Allen elasticities of substitution. Recall that Allen elasticities of substitution can be defined from price derivatives of the cost function corresponding to u .

where we hold constant the utility level and all three \hat{P} vectors and use the symmetry of the Allen elasticity of substitution. PTPT requires that the difference between the two dt coefficients in (16) has the same magnitude, and opposite sign, as the difference between the two dk coefficients. The former difference is therefore the DiD featured in the proposition. The aforementioned PTPT restriction on the coefficients also implies that the former difference is $-\sigma_{TK}$.

The scale effect shown in (15) is the sum of the first two coefficients in (16). Homogeneity and symmetry of the compensated demand for z_T allows that sum to be rewritten as $s_X AES_{XT}$. Homogeneity and symmetry for the other goods, together with PTPT, equate $s_X AES_{XT}$ to the expression shown in (16). QED

COROLLARY (Treatment effect linked to shadow substitution). The difference-in-differences estimator is negative, with magnitude equal to the shadow elasticity of substitution σ_{TK} between the T and K goods in $u()$.

Equation (14) is an expression for the DiD of the logs of the quantity indexes z_T and z_K . The ε^H defined in Proposition 3 is a compensated scale effect in the sense that it describes the compensated responses of both $\ln z_T$ and $\ln z_K$ to a common treatment. If the scale effect in equation (12) is interpreted as a compensated scale effect, then the equation describes a compensated ToT . By the Slutsky equation, the uncompensated scale effect is $\varepsilon^M = \varepsilon^H - (s_T + s_K)\eta$, where η is the income elasticity of demand common to z_T and z_K .¹³ If the scale effect in equation (12) is interpreted as an uncompensated scale effect, then the equation describes a uncompensated ToT .¹⁴

A comparison of (14) and (15) is the starkest demonstration that the treatment-control comparison $-\sigma_{TK}$ reveals something entirely different about the structure of demand than a common treatment $dt = dk$ does. They refer to different entries in the AES matrix. Although the matrix is restricted by symmetry and homogeneity, those restrictions impose no relationship between AES_{XX} and the Allen elasticities that define σ_{TK} . Especially, the AES_{XX} term highlights that the compensated scale effect ε^H is all about substitution to the outside goods \hat{X} whereas treatment-control comparisons reveal only substitution between treatment and controls.

Proposition 4 establishes that the compensated scale effect ε^H summarizes everything about substitution that cannot be discovered with the treatment-control comparison $-\sigma_{TK}$. Moreover, none of the nine Allen substitution terms can be inferred from shares and σ_{TK} alone. Each of them also requires ε^H .

PROPOSITION 4 (Recovering the substitution matrix from DiD and scale). With PTPT, homogeneity, symmetry, and values for σ_{TK} , ε^H , and two expenditure shares, the entire AES matrix can be constructed as:

¹³ Note that the gap between ε^M and ε^H is small to the extent that a large majority of consumer spending is on the outside goods. The DiD is not characterized as either Hicksian or Marshallian because T and K have a common income effect that differences out.

¹⁴ The difference between the compensated ToT and the uncompensated ToT is $\omega_T \eta$.

$$\begin{pmatrix} AES_{TT} & AES_{TK} & AES_{TX} \\ AES_{KT} & AES_{KK} & AES_{KX} \\ AES_{XT} & AES_{XK} & AES_{XX} \end{pmatrix} = \begin{pmatrix} \frac{\varepsilon^H - \frac{\omega_K}{\omega_T} \sigma_{TK}}{\omega_T + \omega_K} & \frac{\varepsilon^H + \sigma_{TK}}{\omega_T + \omega_K} & -\frac{\varepsilon^H}{\omega_X} \\ \frac{\varepsilon^H + \sigma_{TK}}{\omega_T + \omega_K} & \frac{\varepsilon^H - \frac{\omega_T}{\omega_K} \sigma_{TK}}{\omega_T + \omega_K} & -\frac{\varepsilon^H}{\omega_X} \\ -\frac{\varepsilon^H}{\omega_X} & -\frac{\varepsilon^H}{\omega_X} & \frac{\omega_T + \omega_K}{\omega_X} \frac{\varepsilon^H}{\omega_X} \end{pmatrix} \quad (17)$$

Holding constant σ_{TK} and the shares, each of the nine entries varies with ε^H .

Proof. The AES_{XX} entry derives directly from (15). $\omega_X = 1 - \omega_T - \omega_K$. Homogeneity, symmetry, and PTPT require that $AES_{XT} = AES_{XK} = AES_{KX} = AES_{TX}$. Homogeneity of X demand together with the expression for AES_{XX} requires those four Allen elasticities to be each be $-\varepsilon^H/\omega_X$. The remaining four Allen elasticities then follow from symmetry, homogeneity, and the definition (14) of the shadow elasticity σ_{TK} . QED

Equation (17) shows than none of the nine substitution terms can be inferred from shares and treatment-control comparisons alone. Each of them also requires the scale effect ε^H . By proposition 3, DiD only told us about σ_{TK} .

Marshall's Laws of Derived demand were formulated by Marshall (1895) and Hicks (1936) under the more restrictive assumption that z_T and z_K are weakly separable from the X goods.¹⁵ In this case, the shadow elasticity σ_{TK} in u is itself the elasticity of substitution between T and K in the function that aggregates them in preferences.¹⁶ With this interpretation of DiD and σ_{TK} , equation (12) becomes Marshall's Laws expressed in elasticity form. Equations (12) and (14) show the more general case where $-DiD$ is a shadow elasticity of substitution in u .

IV. Bias correction

The straightforward case for interpreting a DiD estimate is when we are interested in ToT – the effect on T of $dt > 0 = dk$ – rather than the effect of a hypothetical aggregate treatment. Proposition 2 shows that $ToT = \lambda\varepsilon + (1-\lambda)DiD$, which we expect to be similar to the DiD estimate when the share treated is close to zero. For example, one could use a DiD estimate of Canada's policy experience as an estimate of what would happen to another small (compared to the world economy) and otherwise similar country that adopts the same policy because both effects would be $DiD dt$.

¹⁵ With randomly assigned treatments, z_T and z_K would, up to sampling error, have the same income elasticity and cross-price elasticities with each element of \tilde{X} . These are the conditions for weak separability with a homothetic aggregator.

¹⁶ The concept of shadow elasticity of substitution was not yet invented when Marshall, Hicks, and Allen were writing about these issues. Allen and shadow elasticities of substitution coincide when there are only two goods.

Even when the parameter of interest involves the effect ε of treating the entire market, price theory shows how DiDs can be used to obtain a reliable estimate of ε . Two instances follow, both where a scale effect can be inferred entirely from the results of a targeted treatment. The result is possible due to the presence of market frictions that create known heterogeneity in the spillover among the untreated. In IV.A., we describe how out-of-market controls in the cross section, combined with *DiD* estimates, can be used to recover the scale effect. Subsection IV.B. provides a setting, familiar from industrial organization and spatial economics, where no control is fully beyond the reach of the invisible hand. Nevertheless, with potential controls differing in terms of distance from the treatment, targeted treatments provide enough information to estimate a scale effect from two DiDs. Both examples have the scale effect outside the range spanned by the two DiDs.

IV.A. Outside- and within-market control groups

Suppose that a fraction n of the controls were beyond the reach of the invisible hand. That is, the spillover of a targeted treatment dt onto the in-market controls is $s_{Kt} dt$, as compared to zero for the out-of-market controls. The average spillover among the controls is therefore $(1-n)s_{Kt} dt$.

The scale effect is defined as before, $\varepsilon \equiv s_{Tt} + s_{Tk}$. For a targeted treatment of $dt = 1$, we denote the difference between the treated outcome and the average control outcome as $DiD(\lambda n)$:

$$DiD(\lambda n) \equiv s_{Tt} - (1 - n)s_{Kt} = \lambda n\varepsilon + (1 - \lambda n)DiD(0) \quad (18)$$

where the second equality follows from the element-by-element equations (7). $DiD(0)$ refers to the difference between ToT and the spillover effect s_{Kt} on the in-market controls. Unlike $DiD(0)$, $DiD(\lambda n)$ puts some weight on the scale effect. The scale effect coefficient is less than one both because only part of the market is treated ($\lambda < 1$) and because only some of the controls are outside the market ($n < 1$). Still, $DiD(\lambda n)$ over- or under-estimates the magnitude of the scale effect according to whether T and K are substitutes or complements, respectively. If none of the controls were in the market ($n = 1$), $DiD(\lambda n)$ would be the ToT but still differ from the scale effect because it does not include the effect on the treated of applying treatments to the untreated in their market.

Having at least some of the controls out of the market raises the possibility of recovering the scale effect from a meta-analysis. Specifically, assume that two DiDs are available from distinct markets with the same ε and $DiD(0)$ but different shares for the out-of-market controls (n) or different treatment shares as reflected in λ . Letting subscripts denote markets, the common scale effect can be written in the two-market case as a weighted average of DiD from each market:

$$\varepsilon = \delta DiD(\lambda_1 n_1) + (1 - \delta) DiD(\lambda_2 n_2) \quad (19)$$

$$\delta \equiv \frac{1 - \lambda_2 n_2}{\lambda_1 n_1 - \lambda_2 n_2} \quad (20)$$

One, but not both, of the markets could have $n = 0$. Note that the market-1 weight δ must be outside the unit interval and therefore put negative weight on one of the DiDs and a coefficient

greater than one on the other. According to (19), the scale effect is outside the range spanned by the two DiDs. The scale effect falls below (above) that range when treatments and controls are complements (substitutes), respectively.

Alternatively, a single targeted treatment permits recovery of the scale effect when the in-market controls can be distinguished from the out-of-market controls. Their outcomes each serve as a distinct comparison for the treated outcome. The DiD using in-market controls is $DiD(0)$, while the DiD using out-of-market controls is $DiD(\lambda) = ToT$. This is the special case of equation (19) that reduces to the ToT expression in equation (7). All of this is possible by having at least some controls that do not experience the full effect of the invisible hand.

IV.B. Market spillovers that diminish with distance

So far, we have considered controls as either in the same market as the treated, or out of the market entirely. Here we allow for a gray area, namely a model in which the magnitude of market spillovers diminishes with “distance” from the treatment. The model is sometimes known as the circle-city or Salop model (1979). A key result is a simple formula for inferring a scale effect on prices entirely from the results of a targeted treatment.

An integer number N of producers are evenly spaced around the circle, whose circumference is equal to one. Distance along the circle can be interpreted literally, as for an urban economics application. Alternatively, gaps between producers on the circle can represent other product-attribute differences.

We assume that each producer $i \in \{1, 2, \dots, N\}$ sets its price p_i as a weighted sum of its own cost c_i and the prices of its immediate neighbors:

$$p_i = \rho_0 + \rho_1 c_i + \rho_2 \frac{p_{i-1} + p_{i+1}}{2} \quad (21)$$

with the constants $\rho_1 > 0$ and $\rho_2 \in (-1, 1)$. Note that $p_0 = p_N$ and $p_{N+1} = p_1$, which appear in equation (21) when evaluated for producers $i \in \{1, N\}$, because of the model’s circular setup. Our Appendix I shows how the special case of equation (21) with $\rho_1 = \rho_2 = \frac{1}{2}$ can be derived from profit-maximizing and utility-maximizing behavior. Our purpose here is to show how it restricts the pattern of market spillovers.

Let $\{dp_1, \dots, dp_N\}$ denote the price effects of a targeted treatment that increases the costs only of producer 1, i.e. $dc_1 > 0 = dc_2 = \dots = dc_N$, and let $\{dP_1, \dots, dP_N\}$ denote the price effects of treating all producers (the entire market) with the same-magnitude cost shock as producer 1 experiences in the targeted treatment. Differencing the total derivative of equation (21) between the marketwide and targeted outcomes,

$$dP_1 - dp_1 = \rho_2 \frac{dP_N + dP_2}{2} - \rho_2 \frac{dp_N + dp_2}{2} \quad (22)$$

The symmetry of the model and the two treatments requires that $dp_2 = dp_N$ and $dP_1 = dP_2 = dP_N$.

$$dP_1 - dp_1 = \rho_2(dP_2 - dp_2) \quad (23)$$

By definition, the scale effect on prices is $dP_1 = dP_2 = \varepsilon$, the DiD for the targeted treatment is $dp_1 - dp_2$, and the ToT for the targeted treatment is dp_1 . Equation (23) can be rewritten as

$$\varepsilon = ToT + \frac{\rho_2}{1 - \rho_2} DiD \quad (24)$$

Note that the DiD in equation (24) refers to the treated-adjacent producers as the controls. To a close approximation, the ToT is the difference $dp_1 - dp_{N/2+1}$ between the treated firm's price and the price charged on the opposite side of the circle.¹⁷ Therefore, equation (24) allows the scale effect on prices to be inferred entirely from the results of a targeted treatment. With $\rho_2 = 1/2$, the scale effect is simply the sum of ToT and DiD. Appendix II further illustrates a broader principle that market frictions allow scale effects to be recovered from a targeted treatment's ToT and DiD.

An analysis related to (24) is possible for quantities. As with prices, the effect on quantities of a market wide cost change is outside the interval bracketed by the ToT and adjacent-producer DiD observed from a targeted cost change. However, unlike the scale effect on prices, the scale effect on quantities is closer to zero than either ToT or DiD.

V. Further examples of difference-in-differences in the marketplace

V.A. Models with time and region fixed effects

Without price theory as a guide, difference-in-differences estimates can easily be misinterpreted in geographical contexts. One case is an early set of studies attempting to detect imperfect competition in cigarette manufacturing in the form of “over-shifting” cigarette excise taxes (Sumner, 1981). Over-shifting means a \$1 per pack tax would increase the retail price of cigarettes by more than \$1 per pack, whereas “one-for-one passthrough” refers to a dollar-for-dollar correspondence between excise taxes and retail prices. These studies were executed with essentially a difference-in-differences framework by comparing states with large tax increases to states with little or no increase.

DiD pass-through studies found nearly one-for-one passthrough, but overlooked the possibility that retail prices in the control states were increased by the tax rates in the treatment

¹⁷ In this way, both (19) and (24) show the scale effect as a weighted sum of two different DiDs, except that the sum of (24)'s weights exceeds one.

states, which can occur through wholesale prices.¹⁸ If the control states were affected in this way, nationwide increases in excise taxes would be over-shifted even though the state DiD shows one-for-one pass through. If we interpret T as retail prices in the states with tax increases and K represents retail prices in the other states, that is the situation illustrated in Figures 2 and Figures 3. The blue arrow represents the retail-price effects of a nationwide tax increase. A national tax would increase prices more than a geographically-concentrated tax increase (green arrow), even in the states targeted by those taxes.

Another example is related to Jaffe, Minton, Mulligan, & Murphy (2019, p. Chapter 17), which concludes that business taxes reduce wages in the long run because the taxes reduce productivity. Nevertheless, an increase in business taxes in a particular locality may not reduce wages in that locality relative to the rest of the nation because workers have a choice of where to live and work. In effect, the wage in any locality is influenced by business taxes throughout the country, or even throughout the world. By failing to account for this, a DiD approach might not show any wage effect of business taxes.

If geographic differences in business taxes result in little or no geographic differences in wages, they might result in especially large geographic differences in employment. This is another case in which the geographic-specific effect is different from the aggregate effect, but this time with the former effect being greater.

Another policy question is the employment effect of public projects such as building a sports stadium or hosting a major event such as the Olympics. Early studies used something like a DiD approach and found a “multiplier”: that total employment in the vicinity of the stadium increased more than the number directly employed by the sports enterprise (Wanhill, 1983; Johnson, Obermiller, & Radtke, 1989). For example, complementary businesses such as restaurants, lodging, and parking were opened nearby. But later studies found that most, if not all, of the additional employment was pulled in from other localities (Dwyer & Forsyth, 2009).

Development economics includes experiments that encourage healthcare providers in treatment villages to supply more healthcare. Others incentivize more instructional effort by teachers in the treatment villages. Such experiments can be analogous to the sports-stadium studies. Namely, through factor markets the experiment reallocates resources from control villages to treatment villages. The per-capita effect of treating all villages would be different unless resources are moved with equal ease (or difficulty) between villages as from outside the village economy as a whole. In our notation, that condition is $\varepsilon = DiD$.

¹⁸ Suppose that, for example, cigarette manufacturers set one nationwide wholesale price because of concerns that regional wholesale price inequality would result in unauthorized wholesale orders and shipments in the low-price regions on behalf of the high-price regions. Such manufacturers would respond to an increase in one state’s excise rate by adjusting their nationwide wholesale price, and through that mechanism indirectly adjust retail prices throughout the nation. Later studies acknowledged this market mechanism’s effect on state differences (Keeler, et al. 1996, Evans, Ringel and Stech 1999, Adhikari 2004); see also Tennant (1950). Harris (1987) emphasizes the results of a federal tax change. Our Appendix II provides a model of such price setting, expressing its results in the format (1). The eigenvalues of S prove to be the national pass-through coefficient (NPTC) and a weighted harmonic mean of one and that same NPTC. The weights depend on the trans-shipment costs, with very little weight on NPTC. That is, due to trans-shipment, difference-in-differences tend to show a one-for-one “effect” of tax on retail price, regardless of the NPTC.

Regional DiD studies often include many regions that are treated at different times and different amounts. The average region has a small share; λ is close to zero in our notation. Neither the ToT nor the DiD reflects much of the scale effect. For instance, with cigarette excise taxes, the DiD indicates what would happen to retail prices in a state that increases its tax rate, while the national effect of a national increase would be different. More generally, the DiD in a many-region study closely approximates the effect on the treated of treating yet another region of similar size, even though it does not indicate the aggregate effect of treating all regions.

Newer studies have proposed to estimate the *ToT* by selecting control observations that are located at a given radius from treated observations. Butts' (2023) reexamination of the Tennessee Valley Authority analysis in Kline and Moretti (2014) is an example. The economic interpretation of such procedures depends on the outcome of interest and the structure of market spillovers. For example, for quantity-price relationships satisfying Hicksian symmetry, as in our Section III.B, a targeted treatment by itself does not provide enough information to infer scale effects. In contrast, the scale effect on prices of a cost treatment in the Salop model can be estimated as the sum of the ToT and adjacent-neighbor DiD observed from a geographically targeted treatment, as shown in our Section IV.B.

V.B. Welfare effects of random treatments

Let's examine treatment effects in a setting where the substitution effect exceeds the scale effect. A large pool of *ex ante* identical workers supplies hours on the intensive margin. Their population is normalized to one. From employers' perspective, any worker's hours are perfect substitutes for another's. In the baseline, each worker is paid the same hourly wage w and supplies the same hours. The aggregate demand for their hours is $D(w)$, with $D'(w) < 0$. The per capita supply of labor is $L(w)$, with $L'(w) > 0$.

An experiment selects a fraction τ of the workers for a wage subsidy $t \geq 0$.¹⁹ Their hours are denoted T per treated and τT in total. The untreated "controls" supply K per control and $(1-\tau)K$ in aggregate. To highlight the analogy with the model (1), our notation also includes k as a subsidy for the controls, although it is not emphasized here. Given values for τ , t and k , an equilibrium is a list $\{w, T, K\}$ of wage and hours satisfying:

$$K = L(w + k)$$

$$T = L(w + t)$$

$$(1 - \tau)K + \tau T = D(w)$$

¹⁹ This is a simplified version of Heckman, LaLonde, and Smith (1999) that focuses on incidence rather than employment effects. [also connect to the example at the beginning of the paper]

Unsurprisingly, $dK/dt < 0 < dT/dt$ and $dw/dt < 0$.²⁰ The subsidy benefits the treated and employers (even those who do not employ any treated) and harms the controls.²¹ Regardless of the share treated, the first magnitude can easily be less than the combined magnitudes of other two.

Take the case when labor demand is wage inelastic, $\tau \in (0,1)$, and the subsidy is small.²² The treated benefit from the subsidy, but employers benefit even more because they pay less for both treated workers and untreated workers. Shrinking the treated share does not change this result. If the treated are to be the primary beneficiary of the subsidy, demand needs to be wage elastic enough or supply inelastic enough.²³ Clearly, quantifying the scale effect is essential for understanding the relationship between the welfare effect on the treated and welfare effects more broadly.

This example also distinguishes the DiD estimator $dT/dt - dK/dt$ from the effect of treating all workers. The DiD estimate is $L'(w)$ because the subsidy moves treated and controls in opposite directions along the supply curve. The equilibrium quantity effect of subsidizing all workers, is a parameter of interest and results from shifting the supply curve downward by dt . As expected from the general substitutes case, this scale effect is closer to zero than the DiD estimate.

V.C. The missing intercept

Equilibrium spillovers have been acknowledged in a macroeconomics literature that uses cross-sectional comparisons to infer effects of national policies, such as cash-transfer programs intended to affect consumption. It recognizes that the consumption of households not receiving a transfer will be affected by the transfers going to other households. Wolf (2023) characterizes the effect as the “missing intercept” problem.

Our paper features a treatment that is uniform among a single treatment group. In that case, the intercept perspective can be seen by writing the treatment outcome dT from the first column of the effects matrix (7):

$$dT = dK + DiD * dt \tag{25}$$

the first term on the RHS of equation (25) is the spillover onto the controls. This “intercept” term would be differenced out if the version of (25) for controls were subtracted. Alternatively, the

²⁰ For this example, the matrix elements corresponding to (1) are $s_{Tt} = L'(w) \frac{(1-\tau)L'(w)-D'(w)}{L'(w)-D'(w)}$, $s_{Tk} = -(1-\tau)L'(w) \frac{L'(w)}{L'(w)-D'(w)}$, $s_{kT} = -\tau L'(w) \frac{L'(w)}{L'(w)-D'(w)}$, and $s_{Kk} = L'(w) \frac{\tau L'(w)-D'(w)}{L'(w)-D'(w)}$. These satisfy $\lambda = \tau$ and the parallel trends assumption (4).

²¹ The expressions for aggregate effects on surplus for treated, controls, and employers are $(dw+dt)\tau T$, $(1-\tau)Kdw$, and $-D(w)dw$, respectively.

²² A “small subsidy” refers to the comparative static $dt > 0$ in the neighborhood of $t = 0$, holding k constant at zero.

²³ For non-zero supply and demand elasticities, the aggregate benefit for the treated as a ratio to the aggregate employer benefit is $1 - \tau - D'(w)/L'(w)$.

intercept might be captured by other variables such as time effects added to the analysis for the purpose of holding constant determinants of the outcome that are distinct from the treatment.

Now consider the relation between outcomes and treatments when there are multiple treatment groups $g = 1, \dots, N$ each receiving a different treatment. Outcomes for controls and each treated group could be calculated sequentially by applying one treatment at a time. If DiD were a constant – not varying with the baseline – then the resulting relation would be linear as in (26):

$$dT_g = dK + DiD * dt_g \tag{26}$$

dK is the outcome change for the untreated, which reflects the market spillovers of all N of the treatments given. In Wolf (2023), the outcome is household consumption, with households differing in terms of the size of the treatment (stimulus check) they received. What we call DiD is estimated in his work (and others) with a panel regression at the household-by-quarter level including time fixed effects and estimated in first differences.

The market featured in our Figure 1 provides another example. DiD is the constant one for a log revenue outcome regardless of the functional form of market demand. Comparisons of log revenue to the productivity treatment in a cross-section of suppliers would show a linear relationship with treatment slope one.

If the missing intercept were identified, the econometrician would possess an estimate of the ToT , which, as we have discussed before, still differs from the scale effect of a uniform treatment on the entire market. At that point, however, estimates of both ToT and DiD can be entered into a variant of equation (12) to determine the scale effect ϵ . Therefore, Wolf (2023) can be interpreted as using a version of the solution described in our section IV.A. within a setting covered by our section III.

VI. Summary and conclusions

Markets are ubiquitous. Consumers and businesses do not live or work in isolation, even approximately so. Perhaps one reaction among those engaged in measurement is to actively attempt to isolate members of the treatment group. Clinical drug trials, for example, do try to prevent trial participants from trading with each other, that is, sharing or exchanging the treatments with others. Some clinical trials even discourage participants from communicating specifics about their trial experiences to prevent (what the investigators view as) potential bias or cross-contamination of results.

We take a different approach in this paper, which is to acknowledge trade and keep it at the center of the analysis. In our framework, parallel trends require the treatment and control outcomes to be weakly separable in utility, production, or cost from all other outcomes. Marshall's Laws of Derived Demand are thereby vehicles for several analytical results. One is that a DiD estimator measures the degree of substitution between treatments and controls, regardless of the fraction of the market that is treated and the magnitude of market spillovers (Proposition 3). In contrast, the

effect of treating the entire market is a “scale effect,” which is the degree of substitution with goods outside the market where treated and controls participate. The effect of the treatment on the treated (ToT) is a weighted average of the scale effect and the DiD, whereas the market spillovers are proportional to their difference (Proposition 2).

Proposition 2 also establishes that, assuming “parallel trends for parallel treatments” (PTPT), the eigenvalues of the treatment-effects matrix are the scale effect and the DiD. As a result, any arithmetic operations on treatment-effects matrices translate into the same operations on their respective scale effects and *DiDs*. This correspondence appeared in a few of our examples where treatment effects on demand were inverted, combined with an identity matrix, or combined with treatment effects on supply.

The presumption that market-equilibrium responses are typically dampened relative to experimental evidence (an example of which is provided in Banerjee and Duflo (2009)) may refer to the market-level feedback between supply and demand, which we have left only implicit in this paper. For the usual incidence reasons, for example, the market-level reduced form for the quantity elasticity of tax changes is $-\varepsilon^s \varepsilon^d / (\varepsilon^s - \varepsilon^d)$. This incidence coefficient reflects equilibrium dampening in the sense that it is less than both the demand elasticity magnitude $-\varepsilon^d$ and the supply elasticity ε^s .²⁴

However, there is more to market equilibrium than supply-demand feedback. In particular, the actions of market participants—even those on just one side of the market—are coordinated by prices. Either the controls are affected by the treatment, the treated would be further affected by treating the rest of the market, or some combination thereof. Market spillovers drive a wedge between *DiD* and the scale effect. The market-level demand elasticity ε^d , the market-level supply elasticity ε^s and the market-level incidence coefficient $\varepsilon^s \varepsilon^d / (\varepsilon^s - \varepsilon^d)$ are each examples of a scale effect. The *DiD* from an experiment with control and treatment in the same market recovers substitution effects instead of scale effects. Treatment-control comparisons by themselves do not even partially identify ε^d , ε^s , or $\varepsilon^s \varepsilon^d / (\varepsilon^s - \varepsilon^d)$.

Complementarity is the case when the treated are affected more by a full-market treatment than by receiving the same treatment while others in the market are untreated. Note that complementarity requires neither increasing returns nor externalities. It does not require that the treated and controls ever meet each other to trade. It does not require Leontief preferences or technology. Complementarity in this sense only means that the scale effect exceeds the substitution effect.

What econometricians sometimes call “spillover” effects are not well described as externalities – missing markets – because markets also transmit treatment effects to the untreated through prices. Analogizing spillover effects with externalities may give the wrong impression that such effects are rare or beyond basic economic training.

²⁴ The incidence coefficient can be derived in the usual way as the equilibrium quantity effect of a one price-unit wedge between market supply and market demand.

Per capita market spillover effects tend to decrease as the size of the treatment group goes to zero, but so does the aggregate treatment effect. Small-scale treatments thereby come with two disadvantages. One is that scale effects are especially obscured by substitution effects. Second, and surprisingly, the spillover effect is comparatively large in the aggregate.

Appendix I: Derivation of the circle-city pricing equation

A continuum of consumers is uniformly distributed around the same circle as the N producers cited in the main text. Consumers have a choice to consume only the outside good, or to travel to a producer and purchase one unit. If the purchase is from a producer charging p_i , the purchase results in a consumer surplus of $v - p_i$ minus a travel cost that is h times the distance traveled.

If the number of producers N or value v are too low, or the cost parameters h and $\{c_i\}$ too high, then the consumers furthest from producers will not make a purchase. Each producer would thereby operate in its own market, with no reaction to the pricing of other producers. We rule out this case.

We also assume that prices and costs are sufficiently similar across producers that no consumer would benefit by leapfrogging a producer to make a purchase. That is, the choice for the consumer is either to travel left or right, and then stop at the first producer encountered. Therefore, there are N distinct locations around the circle where the consumers are indifferent between purchasing from the two nearest producers. For the consumers indifferent between producer i and $i+1$, their location is at distance x from the former and $1/N - x$ from the latter:

$$x = \frac{1}{2N} + \frac{p_{i+1} - p_i}{2h} \quad (27)$$

The demand facing producer i is the sum of the consumers traveling from the producer's left and those traveling from the producer's right:

$$D\left(p_i; \frac{p_{i-1} + p_{i+1}}{2}\right) = \frac{1}{N} + \frac{\frac{p_{i-1} + p_{i+1}}{2} - p_i}{h} \quad (28)$$

This demand function is linear in the producer's own price and the average of the prices of adjacent competitors. A producer pricing exactly at that average receives $1/N$ of the city-wide demand.

The profit maximizing price for a producer i taking as given the prices of competitors is:

$$p_i = \frac{c_i + \frac{p_{i-1} + p_{i+1}}{2} + \frac{h}{N}}{2} \quad (29)$$

The condition (29) must hold for all integers i from 1 to N . It is a special case of equation (21) from the main text, with $\rho_0 = h/(2N)$ and $\rho_1 = \rho_2 = 1/2$.

Appendix II: National and local pass-through of excise taxes

The outcomes are retail prices T and K , expressed in levels (as in the literature). The treatments are excise tax rates t and k , respectively. Consumers are not mobile between the treatment and control areas, which have populations τ and $1-\tau$, respectively. Per-capita consumer demand in each area is a function $D()$ of the area's retail price. Supply prices are $T-t$ and $K-k$, respectively.

A national manufacturer sets retail prices in each area to maximize profits:

$$\begin{aligned} & [T - t - c((T - t) - (K - k))] \tau D(T) \\ & + [K - k - c((T - t) - (K - k))] (1 - \tau) D(K) \end{aligned} \quad (30)$$

where the average and marginal cost $c()$ depends on the gap in supply prices between areas. It is a convex function and minimized when the two areas have the same supply price. These assumptions reflect incentives for trans-shipping by area wholesalers, who reimburse the supply price to the manufacturer and handle excise tax payments. Especially, supply-price gaps incentivize wholesalers in the low-price areas to acquire more quantity than needed for their area and (before excise tax is determined) sell the excess to wholesalers in other areas.

The first order conditions for maximizing profit are:

$$\begin{aligned} & \tau \{D(T) + [T - t - c((T - t) - (K - k))] D'(T)\} \\ & = [\tau D(T) + (1 - \tau) D(K)] c'((T - t) - (K - k)) \\ & = (1 - \tau) \{D(K) + [K - k - c((T - t) - (K - k))] D'(K)\} \end{aligned} \quad (31)$$

By totally differentiating the two equations (31) in the neighborhood of $t = k$ while holding τ constant, yields expressions for dT and dK as functions of dt and dk . We write them in the matrix form (1) (not shown), and simplify them with two definitions:

$$\rho^d(T) \equiv \left[2 - \frac{D''(T)/D'(T)}{D'(T)/D(T)} \right]^{-1} > 0 \wedge \eta(T) \equiv \frac{T}{D(T)} D'(T) < 0 \quad (32)$$

The first is the pass-through coefficient ρ^d defined in the usual way as a transformation of the demand function. The second is the price elasticity of demand. Both depend on the retail price because they are not necessarily constant along the demand curve.

With the S matrix so derived, its eigenvalues are:

$$\varepsilon = \rho^d \wedge DiD = \left[\omega + (1 - \omega) \frac{1}{\rho^d} \right]^{-1} \quad (33)$$

DiD is a weighted harmonic mean of 1 and ρ^d , where the weight is $\omega \equiv \frac{\tau c''(0)}{\tau c''(0) - (1-\tau)\tau\eta} \in (0,1)$. None of the weight is on ρ^d as the trans-shipping cost function becomes more convex or as the treatment share τ approaches zero. The weight λ for calculating ToT from the two eigenvalues is, in this application, the population share τ .

References

- Adhikari, Deerga Raj. 2004. "Measuring market power of the US cigarette industry." *Applied Economics Letters* 11: 957–959.
- Angrist, Joshua D., and Jörn-Steffen Pischke. 2008. *Mostly harmless econometrics: An empiricist's companion*. Princeton, NJ: Princeton University Press.
- Athey, Susan, and Guido W. Imbens. 2017. "The state of applied econometrics: Causality and policy evaluation." *Journal of Economic Perspectives* 31: 3–32.
- Banerjee, Abhijit V., and Esther Duflo. 2009. "The experimental approach to development economics." *Annu. Rev. Econ.* 1: 151–178.
- Banzhaf, H. Spencer. 2021. "Difference-in-differences hedonics." *Journal of Political Economy* 129: 2385–2414.
- Borusyak, Kirill, Peter Hull, and Xavier Jaravel. 2022. "Quasi-experimental shift-share research designs." *The Review of Economic Studies* 89: 181–213.
- Butts, Kyle. 2023. "Difference-in-Differences Estimation with Spatial Spillovers." *arxiv.org*. <https://arxiv.org/abs/2105.03737>.
- Crépon, Bruno, Esther Duflo, Marc Gurgand, Roland Rathelot, and Philippe Zamora. 2013. "Do labor market policies have displacement effects? Evidence from a clustered randomized experiment." *The quarterly journal of economics* 128: 531–580.
- Cunha, Jesse M., Giacomo De Giorgi, and Seema Jayachandran. 2015. "The Impact of In-kind and Cash Transfers on Local Prices." *Economic Journal* 6: 195–230.
- De Chaisemartin, Clément, and Xavier d'Haultfoeuille. 2020. "Two-way fixed effects estimators with heterogeneous treatment effects." *American Economic Review* 110: 2964–2996.
- Deaton, Angus, and John Muellbauer. 1980. *Economics and consumer behavior*. Cambridge University Press.
- Dwyer, Larry, and Peter Forsyth. 2009. "Public sector support for special events." *Eastern Economic Journal* 35: 481–499.
- Egger, Dennis, Johannes Haushofer, Edward Miguel, Paul Niehaus, and Michael Walker. 2022. "General equilibrium effects of cash transfers: experimental evidence from Kenya." *Econometrica* 90: 2603–2643.
- Evans, William N., Jeanne S. Ringel, and Diana Stech. 1999. "Tobacco taxes and public policy to discourage smoking." *Tax policy and the economy* 13: 1–55.
- Glaeser, Edward L., and Joshua D. Gottlieb. 2009. "The wealth of cities: Agglomeration economies and spatial equilibrium in the United States." *Journal of economic literature* 47: 983–1028.
- Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift. 2020. "Bartik instruments: What, when, why, and how." *American Economic Review* 110: 2586–2624.
- Harris, Jeffrey E. 1987. "The 1983 increase in the federal cigarette excise tax." *Tax policy and the economy* 1: 87–111.
- Heckman, James J., and Edward Vytlacil. 2005. "Structural equations, treatment effects, and econometric policy evaluation 1." *Econometrica* 73: 669–738.
- Heckman, James J., Lance Lochner, and Christopher Taber. 1999. "Human capital formation and general equilibrium treatment effects: a study of tax and tuition policy." *Fiscal Studies* 20: 25–40.
- Heckman, James J., Robert J. LaLonde, and Jeffrey A. Smith. 1999. *The economics and econometrics of active labor market programs*. Vol. 3, in *Handbook of labor economics*, 1865–2097. Elsevier.

- Heckman, James, and Rodrigo Pinto. 2024. "Econometric causality: The central role of thought experiments." *Journal of Econometrics* 243 (1-2).
- Hicks, John. 1936. *The Theory of Wages*. London: Macmillan.
- . 1939/1975. *Value and Capital*. Oxford: Oxford University Press.
- Hong, Guanglei, and Stephen W. Raudenbush. 2006. "Evaluating kindergarten retention policy: A case study of causal inference for multilevel observational data." *Journal of the American Statistical Association* 101: 901–910.
- Huber, Martin. 2023. *Causal Analysis*. Cambridge: The MIT Press.
- Hudgens, Michael G., and M. Elizabeth Halloran. 2008. "Toward causal inference with interference." *Journal of the American Statistical Association* 103: 832–842.
- Jacobs, Jane. 1969. "Strategies for helping cities." *The American Economic Review* 59: 652–656.
- Jaffe, Sonia, Robert Minton, Casey B. Mulligan, and Kevin M. Murphy. 2019. *Chicago Price Theory*. Princeton University Press (ChicagoPriceTheory.com).
- Johnson, Rebecca L., Fred Obermiller, and Hans Radtke. 1989. "The economic impact of tourism sales." *Journal of Leisure Research* 21: 140–154.
- Keeler, Theodore E., Teh-wei Hu, Paul G. Barnett, Willard G. Manning, and Hai-Yen Sung. 1996. "Do cigarette producers price-discriminate by state? An empirical analysis of local cigarette pricing and taxation." *Journal of health economics* 15: 499–512.
- Kline, Patrick, and Enrico Moretti. 2014. "Local Economic Development, Agglomeration Economies, and the Big Push: 100 Years of Evidence from the Tennessee Valley Authority." *Quarterly Journal of Economics* 129 (1): 275-331.
- Manski, Charles F. 1993. "Identification of endogenous social effects: The reflection problem." *The review of economic studies* 60: 531–542.
- Marshall, Alfred. 1895. *Principles of Economics*. London: MacMillan and Co.
- McFadden, Daniel. 1963. "Constant elasticity of substitution production functions." *The Review of Economic Studies* 30: 73–83.
- Miguel, Edward, and Michael Kremer. 2004. "Worms: identifying impacts on education and health in the presence of treatment externalities." *Econometrica* 72: 159–217.
- Monte, Ferdinando, Stephen J. Redding, and Esteban Rossi-Hansberg. 2018. "Commuting, migration, and local employment elasticities." *American Economic Review* 108: 3855–3890.
- Munro, Evan, Stefan Wager, and Kuang Xu. 2021. "Treatment effects in market equilibrium." *arXiv preprint arXiv:2109.11647* (arXiv).
- Muralidharan, Karthik, and Paul Niehaus. 2017. "Experimentation at scale." *Journal of Economic Perspectives* 31: 103–124.
- Muralidharan, Karthik, Paul Niehaus, and Sandip Sukhtankar. 2017. "Direct Benefits Transfer in Food." *UC San Diego*.
- Salop, Steven C. 1979. "Monopolistic Competition with Outside Goods." *Bell Journal of Economics* 10 (1): 141-56. <http://www.jstor.org/stable/3003323>.
- Sobel, Michael E. 2006. "What do randomized studies of housing mobility demonstrate? Causal inference in the face of interference." *Journal of the American Statistical Association* 101: 1398–1407.
- Sraer, David, and David Thesmar. 2023. "How to use natural experiments to estimate misallocation." *American Economic Review* 113: 906–938.
- Sumner, Daniel A. 1981. "Measurement of monopoly behavior: an application to the cigarette industry." *Journal of Political Economy* 89: 1010–1019.

- Tennant, Richard B. 1950. "The American cigarette industry: a study in economic analysis and public policy." (*No Title*).
- Viviano, Davide. 2024. "Policy targeting under network interference." *Review of Economic Studies* rdae041.
- Wanhill, Stephen R. C. 1983. "Measuring the economic impact of tourism." *The Service Industries Journal* 3: 9–20.
- Wolf, Christian K. 2023. "The Missing Intercept: A Demand Equivalence Approach." *American Economic Review* 113 (8): 2232-69.