Separating Energy Substitution from Factor-Biased Productivity: Evidence from Norwegian Firms

Fuchen Zhang*

November 14, 2025

Latest version available here

Abstract

The clean–dirty substitution elasticity is central to climate policy, but available estimates are scarce and dispersed. Much of this dispersion reflects identification choices across two empirical routes. Price-based designs typically impose *Constant Elasticity (CE)* and identify a single slope linking log input ratios to log relative prices. Marginal-product designs estimate the production function and infer elasticities from relative marginal products, typically imposing *Hicks-Neutral (HN)* technology. I show that the two routes coincide only when both CE and HN hold. Otherwise, they diverge, and I make the source of divergence explicit. I then develop a bridging estimator that relaxes both CE and HN and jointly recovers firm-time elasticities and factor-biased productivities. It nests the two methods as special cases and allows testing CE and HN assumptions. Applied to Norwegian food-and-beverage firms (2010–2022), the estimates are near-Leontief: the clean–dirty elasticity clusters near zero, implying limited short-run within-firm fuel switching in response to relative prices. CE cannot be rejected, and CE-based estimates align closely with mine. HN is rejected, and HN-based approaches overstate substitutability, implying gross substitutes.

Keywords: Elasticity of substitution; energy transition; factor-biased productivity; production-function estimation

JEL codes: D24, Q41, Q58, Q55, C14, C23.

*Author: fuchang@uis.no. Affiliation: University of Stavanger. I would like to thank my PhD supervisors Torfinn Harding and Anders Åkerman. I would like to thank Jan De Loecker and seminar participants at NOITS 2025; the 47th Annual Meeting of the Norwegian Association of Economists; the University of Stavanger; and Kristiania University of Applied Sciences for constructive comments.

1 Introduction

Analyzing the energy transition requires credible estimates of the elasticity of substitution between clean and dirty energy. This parameter guides both policy and technology. On the policy side, it governs how price instruments such as carbon taxes and clean subsidies translate into aggregate emissions reductions, and how much comes from within-firm fuel switching versus output contraction and reallocation (Acemoglu *et al.*, 2012; Golosov *et al.*, 2014). When the elasticity is low, large price differentials induce little substitution, so emissions fall mainly through the scale channel. When it is high, small relative price changes reallocate energy use toward cleaner inputs, delivering more abatement at lower output losses. On the technology side, the same elasticity shapes the effects and incentives for clean-biased technical change: with greater substitutability, efficiency gains in clean energy displace dirty inputs more strongly, while price signals may provide weaker incentives for directed technical change and the diffusion of green technologies (Acemoglu, 2002; Acemoglu *et al.*, 2012; Beaudry *et al.*, 2010).

Despite its importance, empirical estimates are scarce and scattered. A nontrivial share of the dispersion reflects identification choices rather than pure context (country, sector, aggregation level, horizon).[1] Two empirical routes dominate. The price-based approach estimates how input mixes respond to relative prices. In practice, it often adopts a *constant-elasticity (CE)* specification and uses plausibly exogenous price variations for identification. Recent micro studies in manufacturing follow this route and typically find high substitutability (Jo, 2025; Leclair, 2024; Kaartinen and Prane, 2024). By contrast, the marginal-product–based route estimates the production function and then computes elasticities from relative marginal products. This route can relax CE and allow heterogeneous elasticities, but generally assumes *Hicks-Neutral technology (HN)*. Macro-data studies taking this route find elasticities at or below one (Papageorgiou *et al.*, 2017; Malikov *et al.*, 2018). I refer to the resulting elasticity estimates as CE-estimates and HN-estimates.

These contrasting assumptions create a "mutual exclusion": price-based designs impose CE but relax HN, while marginal-product designs impose HN but relax CE. As a result, the two routes need not estimate the same object. The first part of the paper makes explicit what each approach recovers when its own assumption does not strictly hold. The second part proposes an integrated method between them. It identifies heterogeneous elasticities separately from factor-biased productivities, and allows both assumptions to be tested on the same footing. I focus on micro data. The benchmark elasticity is the direct partial (Hicksian) elasticity of substitution. It captures how the input mix adjusts along an isoquant, *conditional on output and other inputs*. This is a transparent market-free object. I refer to it as the (conditional) elasticity.

[1] See Table 1 for the estimates.

I begin by characterizing the CE-estimate. Under CE (i.e., a two-input CES), a simple rule applies: regress the log input ratio on the log relative price, after addressing endogenous factor-biased productivity, and the slope gives the elasticity. The convenience comes from a strong restriction: a single common slope, regardless of the firm's input mix, technology, equipment, product scope, or adjustment frictions. In practice, slopes differ (e.g., it depends on capital supply: scarce clean-using machine supplies make inputs look more complementary; once the machine becomes sufficient, they become more substitutable). If CE is misspecified, even randomized price shocks deliver a weighted average of firm–time *unconditional* elasticities. Two issues follow. First, the benchmark conditional elasticity is a market-free technology object, and is typically non-negative to ensure a convex isoquant. The unconditional elasticity is an equilibrium object that bundles substitution with other margins shaped by market structure (e.g., scale under profit maximization, reallocation of other inputs).[2] It is context-specific and can even be negative. Second, a single-slope fit is vulnerable to standard negative-weight and outlier-sensitive concerns (Stoker, 1986; Powell *et al.*, 1989).

Subsequently, I characterize HN-estimates based on production-function methods, e.g., control approaches (Olley and Pakes, 1996; Levinsohn and Petrin, 2003; Ackerberg *et al.*, 2015; Gandhi *et al.*, 2020). HN assumption has two benefits. First, for estimating the production function, it simplifies endogeneity control. In control approaches, Hicks-neutrality supports the invertibility (monotonicity) of the proxy input (investment or materials), allowing unobserved productivity heterogeneity to be controlled. Second, when computing elasticities from the estimated production function, ratios of marginal products are robust because, under Hicks-neutrality, the unobserved productivity cancels. The trade-off is that HN can be restrictive. The literature has well established that input-specific efficiency differences are pervasive, arising from factor-biased technologies, managerial practices, and infrastructure (e.g., grid access). Accordingly, I clarify what the marginal-product-based approach recovers. Under misspecified HN, HN-estimate approximates the true elasticity plus a typically positive productivity-selection term. Intuitively, observations using more of X_1 often also have higher X_1-biased productivity.[3] If this productivity-selection channel is not controlled, the local relation between relative marginal products and the input mix becomes too flat, inflating the elasticity. Formally, in a general profit-maximization framework with a joint distribution of productivities and prices, the HN estimate decomposes into the true elasticity, a productivity-selection component, and several small terms. If input-biased hetero-

[2]Similar to the unconditional elasticity, the Allen–Uzawa and Morishima elasticities are also equilibrium objects. They are defined on the cost function under cost minimization with all inputs freely adjustable. However, CE-estimates are specific to the empirical context and need not coincide with those cost-based objects. I thus label it *unconditional* to avoid conflating.

[3]By "often" I mean the typical case where the demand response to prices and productivities are not strongly negatively correlated, ensuring that higher X_1 use doesn't systematically come with lower X_1-biased productivity.

geneity is large, those small terms can combine to push the HN-estimate below zero.[4] Imposing CE and HN together in the marginal-product approach is particularly problematic when input bias is substantial. The joint restrictions can force a high-elasticity CES fit (gross substitutes) even when the true technology features gross complements. I use Monte Carlo simulations to corroborate these characterizations of CE-estimates and HN-estimates.

Rather than treating price-based and marginal-product-based approaches as alternatives, I propose a bridging estimator that relaxes both CE and HN, nesting them as special cases. The procedure jointly recovers firm–time substitution elasticities and factor-biased productivities. It combines the price-based approach with the technical-control method of Ackerberg *et al.* (2015): using observed input quantities, expenditures, and a proxy (e.g., materials) to control for unobserved productivities and relying on multidimensional timing restrictions for identification. Plausibly exogenous price shifters can be incorporated to sharpen identification and make the source of variation transparent. The estimates allow CE and HN to be tested. When either assumption fails, the framework delivers a full, data-based decomposition of the divergence in estimated elasticities: for the CE-estimate, the gap splits into weighting and unconditionality components; for HN, it splits into productivity-selection plus small remainder terms. Each component is separately quantifiable.

I apply the method to linked Norwegian administrative micro data (2010–2022), merging: (i) annual plant–level energy reports with physical quantities and expenditures of seven energy groups; (ii) annual firm–level financial statements; and (iii) monthly employer–employee matches. I estimate a five–input translog value–added function with Hicks-neutral revenue productivity and clean–, dirty–, and labor–augmenting productivities.[5]

The estimates imply a near-Leontief relation between clean and dirty energy in the food and beverages industry: firm–time substitution elasticities cluster near zero.[6] I benchmark CE- and HN-based elasticity estimates against my estimates. The CE assumption cannot be rejected, and CE-based estimates align closely with my estimates. The HN assumption is rejected, and HN-based estimates overstate substitutability, making clean and dirty substitutes, globally or at least

[4] A "negative" elasticity means substitutability has gone beyond the perfect-substitutes limit. It implies a *concave* isoquant, and thus it is a misspecification flag rather than a meaningful technology. Intuitively, as substitutability passes the perfect-substitutes bound, the substitution elasticity flips from very large positive values to very large negative values.

[5] (i). All fossil fuels are classified as dirty energy, and electricity, steam/district heating, and solid biomass/waste are considered clean. (ii). Following the technical control approach (Ackerberg *et al.*, 2015), identification requires that intermediate materials enter gross output in a Leontief relationship with value added. The Norwegian administrative financial statements provide a transparent breakdown of "other operating costs," allowing us to separate relatively non-substitutable physical intermediate goods from more substitutable inputs, e.g., purchased services. Accordingly, I use intermediate goods as the proxy variable and treat the rest as a separate input in the value-added production function. (iii). I cannot allocate the capital stock to clean- versus dirty-using components. Accordingly, I prefer proxying firms' ability to utilize each energy type through clean- and dirty-augmenting productivity terms.

[6] More industries will be investigated in future work.

locally for a substantial set of observed input mixes.

The near-Leontief relation points to limited scope for within-firm fuel switching in response to relative-price instruments. Adjustment instead occurs through lower output, reallocation toward cleaner producers, and energy-biased technical change/equipment turnover. For example, in dairies that pasteurize with steam from LPG/diesel boilers, a carbon tax mainly reduces throughput and can erode competitiveness. Switching is minimal without capital spending on high-temperature heat pumps or electric boilers. Because substitution is weak, subsidizing clean energy prices could, in principle, raise emissions via scale effects if output expands. Moreover, estimated clean-relative-to-dirty productivity varies more widely than hicks-neutral and labor-augmenting productivities, underscoring potentially strong distributional impacts of climate policy across firms.

Using the estimated productivity dynamics, I test for path dependence: whether a temporary clean- or dirty-biased productivity shock generates a persistent and widening advantage. If path dependence holds, a temporary green R&D subsidy can permanently redirect technical change and achieve a long-term energy transition. Under strong complementarity between clean and dirty, Acemoglu *et al.* (2012) predict no path dependence and directed R&D always tilts toward the less efficient input until energy productivities converge. The evidence is consistent with these predictions. I find no path dependence: temporary advantages dissipate rather than widen. Two components matter for the advantage: own-shock responses and cross-energy responses. Own-shock responses of clean- and dirty-augmenting productivity decay at similar rates. Cross-effects are asymmetric: a dirty-biased (clean-augmenting) shock depresses the subsequent path of clean-biased (dirty-augmenting) productivity, but not enough to overturn overall reversion.[7] Conversely, a clean-biased (dirty-augmenting) shock lifts dirty-biased (clean-augmenting) productivity. This asymmetry is in line with Acemoglu *et al.* (2012)'s prediction that the relatively weaker input (clean) is targeted first. Finally, I document a robust association between labor-augmenting and clean-biased productivity, consistent with evidence on green technology and labor bias (Hedne, 2025).

This paper's first contribution is a transparent framework for measuring substitution elasticities. It places CE- and HN-based estimates on comparable footing and shows they agree only when both assumptions hold. It also provides implementable decompositions so reported numbers are comparable. For illustration, the CE estimate can be written as the target elasticity plus two pieces: a weighting component and an unconditionality component. This reduces misinterpretation and guides calibration. CE-based market-equilibrium elasticities should not be used as technology elasticities. Either adjust them by stripping those margins using the decomposition, or treat them as reduced-form targets and calibrate a model that excludes the market channels

[7]Because clean and dirty are strong complements in my estimates, factor-augmenting innovations display the "opposite-bias" property (Acemoglu, 2002): a clean-augmenting innovation is dirty-biased in the sense that it raises the equilibrium intensity of the other input at given relative prices, and vice versa.

already embedded in the estimate to avoid double counting. HN-based elasticities should be used as substitution parameters only when Hicks-neutrality is credible.

The paper is related to production–function estimation. It is close to recent nonparametric control-function work that allows factor-biased productivities (Demirer, 2022; Ackerberg *et al.*, 2023). I accommodate richer input sets and multiple factor-biased productivities to study substitutability across inputs, trading off against a mild functional-form assumption (arbitrary nested CES or translog) that delivers separability of productivity parameters. The framework complements the factor-biased productivity approaches (Krusell *et al.*, 2000b; León-Ledesma *et al.*, 2010; Doraszelski and Jaumandreu, 2018; Zhang, 2019) by allowing more flexible functional forms and non-constant substitution elasticities. It broadens classic control-function approaches (Olley and Pakes, 1996; Levinsohn and Petrin, 2003; Ackerberg *et al.*, 2015; Gandhi *et al.*, 2020) by extending identification beyond hicks-neutrality.

Finally, the framework is portable beyond clean–dirty energy to contexts where input–biased productivity and substitutability are first–order, such as skill–biased technical change and wage inequality (Krusell *et al.*, 2000b; Acemoglu, 2002; Beaudry *et al.*, 2010), labor–capital substitution and the labor share (Karabarbounis and Neiman, 2013), and multi–origin input substitution in global supply chains (Antràs and Chor, 2013).

The paper proceeds as follows. Section 2 reviews substitution concepts and Section 3 motivates a better understanding of identification assumptions. Section 4 and Section 5 formalize CE and HN-estimates, and discuss the biases. Section 6 presents the integrated method. Section 7 and Section 8 report the empirical application to Norwegian micro data, and robustness. Section 9 concludes.

2 Clean–Dirty Elasticity of Substitution

This section defines the clean–dirty elasticity of substitution and explains its relevance for policy design.

2.1 Elasticity of Substitution

The (direct partial) elasticity of substitution quantifies how easily one input can be traded for another while holding output and the remaining inputs fixed.[8] For two inputs X_1 and X_2 with

[8]See Hicks (1932) for the two-input definition and McFadden (1963) for multi-input generalizations. See Blackorby and Russell (1989) for its relationship with two alternative definitions. I emphasize direct partial elasticity because it is defined at the production-function level and does not require market structure assumptions, whereas the Allen and Morishima elasticities are cost-based measures derived under cost minimization with all inputs fully adjustable.

marginal products F_{X_i} and prices P_i, the elasticity is

$$\sigma_{X_1X_2} \equiv \left.\frac{dln(X_1/X_2)}{dln(F_{X_2}/F_{X_1})}\right|_{Y,\boldsymbol{X}_{-\{1,2\}}} = \left.\frac{dln(X_1/X_2)}{dln(P_2/P_1)}\right|_{Y,\boldsymbol{X}_{-\{1,2\}}} \geqslant 0. \tag{1}$$

When $0 < \sigma_{X_1X_2} < 1$, the inputs are *gross complements*. For instance, an increase in the relative price of X_1 leads to a relative increase in the expenditure on it. When $\sigma_{X_1X_2} > 1$, they are *gross substitutes*.

Throughout, I refer to $\sigma_{X_1X_2}$ (**conditional**) elasticity of substitution, emphasizing that it holds Y and $X_{-\{1,2\}}$ fixed.

2.2 Clean–Dirty Elasticity

The elasticity of substitution between clean (C) and dirty (D) energy, σ_{CD}, governs the curvature of energy isoquants and hence the ease with which firms can replace D by C while holding output fixed.[9] As a result, σ_{CD} is pivotal for the cost of decarbonization and the design of climate policy.

Carbon tax. For a marginal increase in the price of dirty energy P_D, the total Marshallian response of D can be decomposed into a Hicksian substitution effect, holding output Y constant, and a scale effect:

$$\frac{dlnD}{dlnP_D} = \underbrace{\overbrace{-\left.\frac{\partial lnD}{\partial lnP_C/P_D}\right|_{lnY}}^{\leqslant 0}}_{\text{Substitution effect}} + \underbrace{\overbrace{\frac{\partial lnD}{\partial lnY}\frac{\partial lnY}{\partial lnP_D}}^{\leqslant 0}}_{\text{Scale effect}} \tag{2}$$

The magnitude of substitution term is increasing in σ_{CD}: higher substitutability amplifies the reallocation away from D toward C. The scale term captures output contraction from higher energy costs and depends on market structure and the general-equilibrium closure.

When σ_{CD} is low, abatement relies relatively more on output contraction, making a given emissions target costlier and typically requiring a higher and more persistent carbon tax (Acemoglu *et al.*, 2012; Golosov *et al.*, 2014). When σ_{CD} is high, the substitution channel does most of the work, so a lower, shorter-lived tax can achieve the same abatement.[10]

Clean subsidy. For a marginal subsidy to clean energy, the decomposition is analogous:

$$\frac{dlnD}{d(-lnP_C)} = \underbrace{\overbrace{-\left.\frac{\partial lnD}{\partial lnP_C/P_D}\right|_{lnY}}^{\leqslant 0}}_{\text{Substitution effect}} + \underbrace{\overbrace{\frac{\partial lnD}{\partial lnY}\frac{\partial lnY}{\partial d(-lnP_C)}}^{\geqslant 0}}_{\text{Scale effect}} \tag{3}$$

[9] See Figure A.1 for an illustration of isoquants under different elasticities.
[10] See Figure A.2 for an illustration.

By lowering the overall energy price, a clean subsidy generally raises output, resulting in a positive scale effect. Emissions fall only if the substitution response dominates. If $\sigma_{CD} < 1$, a standalone clean subsidy can perversely increase D unless paired with a carbon tax.[11]

Clean-augmenting technical change An increase in clean-augmenting efficiency Ω_C is isomorphic to a fall in P_C:

$$\frac{dlnD}{dln\Omega_C} = \underbrace{\overbrace{-\left.\frac{\partial lnD}{\partial ln\Omega_D/\Omega_C}\right|_{lnY}}^{\leqslant 0}}_{\text{Substitution effect}} + \underbrace{\overbrace{\frac{\partial lnD}{\partial lnY}\frac{\partial lnY}{\partial ln\Omega_C}}^{\geqslant 0}}_{\text{Scale effect}} \tag{4}$$

Hence clean-augmenting innovation mirrors a clean subsidy: with high σ_{CD}, innovation reduces emissions; with low σ_{CD}, rebound via the scale channel can offset substitution unless complemented by carbon pricing.[12]

Under gross substitutes assumption, the first-best policy combines a carbon tax with a clean R&D subsidy (Acemoglu *et al.*, 2012). However, technological and infrastructural constraints can imply regional or short-run complementarity, so leaning on R&D or clean subsidies risks temporary emission increases. Because atmospheric carbon accumulates, even short-run rebounds are consequential, motivating the need for a better understanding of the elasticity.

3 Literature and motivation

Despite its salience for welfare analysis and policy design, empirical estimates of the clean–dirty elasticity remain scarce. Much of the dispersion appears to come from identification choices, rather than pure context heterogeneity (country, sector, aggregation-level, horizon).

3.1 Identification choice

Table 1 summarizes representative estimates and their design assumptions. Most papers adopt one of two simplifications: (i) *Constant Elasticity (CE)*, which treats the elasticity as invariant to prices, technology, and scale; and (ii) *Hicks-Neutral (HN)*, which rules out input-biased productivity heterogeneity.

Recent micro studies in manufacturing impose CE and use plausibly exogenous price variations to address endogeneity arising from factor-biased productivity. They often recover high substitutability (e.g. Jo, 2025; Leclair, 2024; Kaartinen and Prane, 2024). By contrast, macro-data

[11] See Figure A.3 for an illustration.
[12] See Figure A.4 for an illustration.

Table 1: Selected estimates of clean-–dirty substitution elasticity and identification assumptions

Study	Aggregation	Estimates	Identification assumption	
			CE	**HN**
Papageorgiou *et al.* (2017)	Macro: Country–industry	**1.38 ∼ 2.87**	**Yes**	**Yes**
Malikov *et al.* (2018)	Macro: Country–industry	**> 1 only in 1/3 of observations**	**No**	**Yes**
Jo (2025)	Micro: French manufacturing firms	**3.32**	**Yes**	**No**
Leclair (2024)	Micro: Indian steel establishments	**2.224**	**Yes**	**No**
Kaartinen and Prane (2024)	Micro: Swedish manufacturing firms	**1.39 ∼ 4.4**	**Yes**	**No**

studies relax the CE restriction and scrutinize functional-form flexibility, finding elasticities at or below unity (Papageorgiou *et al.*, 2017; Malikov *et al.*, 2018). They adopt HN to avoid concerns regarding endogenous factor-biased productivity.

This "mutual exclusion" in identification strategy, where one strand fixes elasticity and chases exogenous prices while the other allows flexibility but assumes factor neutrality, means that the two strands do not recover the same object. This motivates a framework that makes objects explicit for each approach and places CE and HN on equal footing, so differences across studies can be interpreted as differences in identification rather than being mistaken for context. Accordingly, I will review CE- and HN-based designs, discuss their restrictions, and show what each truly recovers when these assumptions do not hold in a macro-level short-to-medium-run context.[13]

4 Constant Elasticity (CE) Assumption

This section formalizes the *Constant Elasticity (CE)* assumption and the associated reduced-form estimation. I then discuss the restriction of this assumption and the potential bias when the assumption does not hold.

4.1 Setup and review

Assumption CE. *The elasticity of substitution between X_1 and X_2 is a constant, $\sigma_{X_1X_2} \equiv \sigma \geq 0$.*

Under CE, X_1 and X_2 must enter the production technology through a two-input CES aggregator, possibly embedded in a larger production function with other inputs $\boldsymbol{Z}$ (Arrow *et al.*, 1961;

[13]Estimates hinge on aggregation and horizon: macro (country–industry) mix within-firm substitution with between-firm reallocation and map to aggregates; micro (firm/plant) isolate within-firm mechanisms and heterogeneity. Short-run elasticities (with quasi-fixed capital/infrastructure) are typically smaller; long-run ones allow adoption/technical change, are often larger, and harder to identify.

Berndt and Christensen, 1973):

$$\textbf{CE} \iff Y = \Omega_R F\big(\Omega_1 X_1,\ \Omega_2 X_2,\ \boldsymbol{Z}\big) = \Omega_R F'\Big(CES(\Omega_1 X_1, \Omega_2 X_2),\ \boldsymbol{Z}\Big), \tag{5}$$

with

$$CES(\Omega_1 X_1, \Omega_2 X_2) = \Big[\alpha_1(\Omega_1 X_1)^{\frac{\sigma-1}{\sigma}} + \alpha_2(\Omega_2 X_2)^{\frac{\sigma-1}{\sigma}}\Big]^{\frac{\sigma}{\sigma-1}}, \qquad \alpha_i > 0,\ \alpha_1 + \alpha_2 = 1. \tag{6}$$

where Ω_R and Ω_i denote the hicks-neutral revenue productivity and factor-augmenting productivity of input X_i, and σ, the X_1-X_2 substitution elasticity, is independent of output Y and the levels of other inputs $\boldsymbol{Z}$. Let $\gamma \equiv (\sigma-1)/\sigma \leq 1$ denote the substitutability parameter. The CES family nests well known Leontief (perfect complements:$\sigma \to 0$), Cobb–Douglas ($\sigma \to 1$), and linear (perfect substitutes: $\sigma \to \infty$) as special cases.

Imagine a panel involving firm j and time t. Equation (7) implies a log-linear expression of the marginal rate of substitution MRS, motivating the reduced form regression:

$$\begin{aligned} MRS_{jt} &= \frac{P_{1,jt}}{P_{2,jt}} = \frac{\alpha_1}{\alpha_2}\Big(\frac{\Omega_{1,jt}}{\Omega_{2,jt}}\Big)^{\frac{\sigma-1}{\sigma}}\Big(\frac{X_{1,jt}}{X_{2,jt}}\Big)^{-\frac{1}{\sigma}} \\ \implies ln\frac{X_{1,jt}}{X_{2,jt}} &= \sigma \cdot ln\frac{P_{2,jt}}{P_{1,jt}} + \overbrace{\underbrace{(\sigma-1)ln\frac{\Omega_{1,jt}}{\Omega_{2,jt}}}_{\text{Factor-biased productivity}} + \underbrace{\varepsilon_{jt}}_{\text{Noise}}}^{\text{Unobserved}} \end{aligned} \tag{7}$$

Estimation is complicated by endogeneity, as factor-biased productivity is unobserved. Early work employs time trends or fixed effects, while recent studies seek exogenous shifters, such as Bartik instruments built on policy shocks, to identify σ.

4.2 Why the elasticity of substitution may not be constant

The CE assumption requires the elasticity of substitution between two inputs to be invariant to the economic environment. This may be untenable once one allows for richer technologies, factor and product heterogeneity (supply curves, quality), and variation in factor availability and infrastructure (e.g., access to renewables). Even within the CES class, nesting induces state dependence: cross–nest elasticities vary with prices, quantities, and productivities through output shares. For more flexible functional forms, elasticities can be explicitly local functions of the data.

A nested-CES example. Consider a two-tier production technology in which each energy type $i \in \{C, D\}$ combines with a specific machine K_i, and the two composites Q_i are then aggregated

into final output Y:

$$\begin{aligned}
&\text{Clean composite:} && Q_{C,jt} = \left[\alpha_C\, C_{jt}^{\frac{\sigma_c-1}{\sigma_c}} + (1-\alpha_C)\, K_{C,jt}^{\frac{\sigma_c-1}{\sigma_c}}\right]^{\frac{\sigma_c}{\sigma_c-1}},\\
&\text{Dirty composite:} && Q_{D,jt} = \left[\alpha_D\, D_{jt}^{\frac{\sigma_d-1}{\sigma_d}} + (1-\alpha_D)\, K_{D,jt}^{\frac{\sigma_d-1}{\sigma_d}}\right]^{\frac{\sigma_d}{\sigma_d-1}},\\
&\text{Outer aggregation:} && Y_{jt} = \left[\alpha_Q\, Q_{C,jt}^{\frac{\sigma_o-1}{\sigma_o}} + (1-\alpha_Q)\, Q_{D,jt}^{\frac{\sigma_o-1}{\sigma_o}}\right]^{\frac{\sigma_o}{\sigma_o-1}}.
\end{aligned} \tag{8}$$

Suppose the within-composite elasticities satisfy $\sigma_c, \sigma_d < 1$ (energy and the related machine are complements) while the outer elasticity satisfies $\sigma_o > 1$ (the two composites are substitutes).[14]

What is the elasticity of substitution between clean and dirty energy, σ_{CD}, in this setting? By Sato (1967), the cross-nest elasticity between two inputs equals a share-weighted average of the elasticities of the nests along the path linking them. In (8), σ_{CD} is a weighted average of $(\sigma_c, \sigma_d, \sigma_o)$ with weights given by endogenous output shares. Therefore, σ_{CD} varies with input uses, prices, and productivities.

The intuition is that substitution between C and D operates through two margins: (i) reallocation across composites at rate σ_o, and (ii) adjustment within composites at rates σ_c and σ_d. Suppose dirty-specific machines are ample so that energy prices on the dirty side are not binding, the effective bottleneck is the clean composite. When clean machines K_C are scarce/expensive, adjustment within the clean composite is difficult and $\sigma_{CD} \to \sigma_c < 1$ (gross complements). As K_C becomes abundant, adjustment shifts to the outer margin and $\sigma_{CD} \to \sigma_o > 1$ (gross substitutes).

4.3 Bias

I now consider cases where the CE assumption fails in the data and characterize the resulting bias of CE-type estimators. I work with the following general representation:

$$x_{jt} \;=\; \int_{p_0}^{p_{jt}} \sigma(p, \boldsymbol{Z}_{jt})\; dp \;+\; \underbrace{\varepsilon_{jt}}_{\text{noise}}, \qquad x_{jt} \equiv ln\frac{X_{1,jt}}{X_{2,jt}}, \quad p_{jt} \equiv ln\frac{P_{2,jt}}{P_{1,jt}}, \tag{9}$$

where $\sigma(\cdot)$ denotes the point elasticity of substitution, possibly varying with relative prices p_{jt} and firm characteristics $\boldsymbol{Z}_{jt}$ (e.g., the price level, output Y_{jt}, other inputs $X_{-\{1,2\},jt}$, and factor-biased productivities $\boldsymbol{\Omega}_{jt}$).[15] For clarity, I treat $\boldsymbol{Z}_{jt}$ as the collection of all channels that may generate endogeneity, and ε_{jt} is a separable error.[16]

[14] Productivities parameters are omitted to streamline the exposition.

[15] The lower limit p_0 is an arbitrary normalization.

[16] See, e.g., Matzkin (2003); Chesher (2003); Imbens and Newey (2009) for identification with non-separable noise.

I refer to the OLS slope from regressing relative demand on relative prices as the *CE-estimate*:

$$\hat{\sigma}^{\text{CE}} \equiv \frac{\text{Cov}(x_{jt},\, p_{jt})}{\text{Var}(p_{jt})}, \tag{10}$$

I will proceed in steps, starting from the most transparent benchmark and generalizing the conditions sequentially.[17]

Case A. *$\sigma(p_{jt}, \boldsymbol{Z}_{jt})$ effectively reduces to $\sigma(p_{jt})$.*[18] In this benchmark, characteristics-driven endogeneity is absent.

Proposition 1. *Given $\sigma(p_{jt})$, CE-estimate is a convex combination of all elasticities on the compact interval $\mathcal{P} = [\min\{p_{jt}\}, \max\{p_{jt}\}]$, where the weights only depend on the distribution of p_{jt}:*

$$\hat{\sigma}^{CE} = \int_{\mathcal{P}} w(p)\, \sigma(p)\, dp, \qquad w(p) > 0, \quad \int_{\mathcal{P}} w(p)\, dp = 1,$$

Proposition 1 echoes the classic average derivative result: When independent variables have continuous support with strictly positive density, the OLS estimate is a convex combination of point derivatives.[19]

By contrast, if the empirical support of p_{jt} is discrete or gapped, convexity need not hold, and some weights on individual elasticities can be *negative*.

Corollary 1. *Given $\sigma(p_{jt})$, CE-estimate is an affine combination of individual elasticities $\sigma(p_{jt})$. The weights reflect the distribution of p_{jt} and the monotonicity and curvature of σ. Some weights can be negative if σ is not monotonic on the discrete or gapped support of p_{jt}.*

Consequently, negative weights may push $\hat{\sigma}^{\text{CE}}$ outside the range of individual point elasticity.[20]

I illustrate the emergence of negative weights in the nested–CES setting of (8). Consider a short-run environment with quasi-fixed capital stocks K_C, K_D. Firms cost–minimize over (C, D) taking prices as given. Variations in output and capitals are eliminated so that the clean–dirty elasticity σ_{CD} depends only on the relative price $p \equiv ln(P_C/P_D)$. There are three price regions: (i) Low p: The clean-specific machine binds and substitution loads on the clean inner nest. Thus, more weight is on σ_c, and σ_{CD} is low (complements). (ii) High p: Similarly, the dirty machine is binding, σ_d dominates, and σ_{CD} is low (complements). (iii) Moderate p: Neither machine binds, σ_o dominates, and σ_{CD} is high (substitutes). Suppose the panel is composed mainly of firms

[17] See Appendix A for more detailed discussion.

[18] E.g., σ does not depend on $\boldsymbol{Z}_{jt}$, or $\boldsymbol{Z}_{jt}$ has no within-sample variation.

[19] See Stoker (1986); Powell *et al.* (1989) for the identification of the average elasticity weighted by the density of p_{jt} when σ is continuously differentiable and p_{jt} is draw from a continuous, differentiable, and bounded distribution.

[20] See Figure A.5 for a simple illustration.

in regions (i) and (ii) (e.g., due to persistent heterogeneity in resource access or infrastructure). By Proposition 1, the CE estimator is a convex combination of point elasticities over the entire observed price range, hence it implicitly averages elasticities from the middle region (iii) even if that region is thinly represented in the data. Consequently, the CE-estimate can exceed all individual elasticities, resulting in the negative weight concern in Corollary 1.[21]

In the left panel of Figure 1, I corroborate this mechanism with a Monte Carlo that progressively thins support in the moderate-price region (iii): starting from a baseline uniform price distribution, I carve out a symmetric gap within region (iii) and increase its length (x-axis) while holding all other primitives fixed. In each simulation, I plot the maximum and minimum point elasticities across firms, and the CE-estimate. As the gap widens, the CE slope remains largely unchanged, even as observations from (iii) vanish, and eventually exceeds the maximum realized point elasticity, i.e., negative weight cncern.

CE estimators are also *outlier-sensitive*. Observations at the tails of the relative–price distribution exert high leverage on $\hat{\sigma}^{CE}$. When $\sigma(p)$ varies with p, elasticities evaluated at extreme prices can be far from those relevant for the bulk of firms.[22]

I run a targeted simulation where most firms face similar relative prices (and hence similar point elasticities ≈ 0.5), while a small, randomly chosen tail faces extreme relative prices with much higher point elasticities ≈ 1.5. I vary only the share of tail firms (x-axis) and the size of the price gap (mild vs. massive). The right panel of Figure 1 shows the pattern: even a small fraction of the tail exerts high leverage on the CE-estimate, pulling it toward the tail elasticity. The distortion is significantly larger when the price gap is larger.

Case B: $p \perp \boldsymbol{Z}$. Even if σ varies with firm characteristics $\boldsymbol{Z}$, full exogeneity of relative prices implies a similar result.

Corollary 2. *Under $p \perp \boldsymbol{Z}$, CE-estimate remains an affine combination of individual elasticities, with higher weight on more prevalent characteristics (e.g., larger probability mass $\mathbb{P}(\boldsymbol{Z})$).*

However, $p \perp \boldsymbol{Z}$ can be a demanding condition in practice.

Case C: *Relative prices p_{jt} are randomly assigned, but $p \not\perp \boldsymbol{Z}$.* Random price shocks justify $p \perp \boldsymbol{Z}$ under cost minimization with fixed other inputs and exogenous productivities. Outside this ideal, firms may optimally *adjust characteristics* $\boldsymbol{Z}_{jt}$ (e.g., output, other inputs, technology adoption), so $p \not\perp \boldsymbol{Z}$.

[21] Because re-expressing the estimator using only the observed individual elasticities then forces algebraic weights of opposite signs.

[22] See Figure A.6 for a simple illustration.

Proposition 2. *Given randomly assigned relative prices, the CE estimator $\hat{\sigma}^{CE}$ is an affine combination of individual "unconditional" elasticities of substitution, defined as follows:*

$$\underbrace{\sigma^{UN}}_{\textbf{\textit{Unconditional}}} = \underbrace{\sigma(p, \boldsymbol{Z})}_{\textbf{\textit{(Conditional)}}} + \underbrace{\int \overbrace{\frac{\partial \sigma(p, \boldsymbol{Z})}{\partial \boldsymbol{Z}}}^{\neq 0 \text{ without CE}} \frac{d\boldsymbol{Z}}{dp}\, dp}_{\textit{Characteristics adjustment}} \tag{11}$$

As in (1), the conditional elasticity holds $\boldsymbol{Z}_{jt}$ fixed and thus captures the pure substitution response. When CE fails, optimal adjustments in $\boldsymbol{Z}_{jt}$ alter $\sigma(p, \boldsymbol{Z}_{jt})$, so CE-estimates add these adjustment effects on top of substitution. Intuitively, the conditional elasticity is a technology object, describing how the input ratio adjusts along a fixed-output isoquant and does not depend on market structures. It is typically non-negative to ensure convex isoquants. By contrast, the

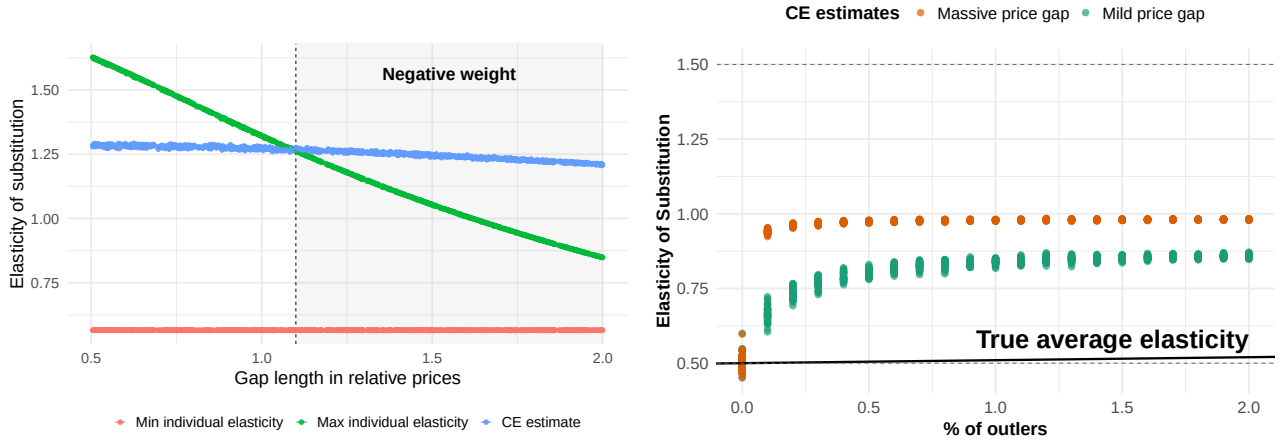


Figure 1: Monte Carlo simulation of negative weight (left) and outlier (right) in CE estimation

Notes: Left panel: Each Monte Carlo replication simulates $n = 1000$ firms under the nested-CES technology in (8). The inner nests are complements aggregations ($\sigma_c = \sigma_d = 0.2$), and the outer nest is a substitutes aggregation ($\sigma_o = 2$). Firms share similar input K_C, K_D and output Y and choose (C, D) to minimize cost given prices (Case A condition holds). Three regions based on log relative price p are: regions (i) when $p \in (-\infty, -1.6)$, regions (ii) when $p \in (1.6, +\infty)$, and regions (iii) when $p \in (-1.6, 1.6)$. For each firm, log relative price p is randomly drawn from the interval with gap: $[-3, -a] \cup [a, 3]$. $a \in [0.5, 2]$ is randomly drawn in each dataset and shown on the x-axis. CE-estimate, and maximal and minimal individual elasticity are plotted based on y-axis. Results are based on 1,000 replications. CE weights must be negative when CE-estimate exceeds the maximum individual elasticity (gap > 1.1). Right panel: Each Monte Carlo replication simulates $n = 1000$ firms producing with a nested CES: $Y = CES(CES(X_1, Z_1), X_2)$ so that elasticities of substitution between X_1, X_2 is not a constant. The inner nest is a complements aggregation ($\sigma = 0.2$), and the outer nest is a substitutes aggregation ($\sigma = 2$). Firms share similar input Z_1 and output Y and choose (X_1, X_2) to minimize cost given prices (Case A condition holds). For most firms, relative prices are similar, yielding point elasticites $\sigma_{X_1X_2} \approx 0.5$. A subset of "outlier" firms faces much higher relative prices with $\sigma_{X_1X_2} \approx 1.5$ in equilibrium. The share of outliers in each dataset is randomly drawn and shown on the x-axis. CE-estimates and the true average elasticities are plotted based on y-axis. I consider two scenarios for the outlier price gap (massive vs. mild). Results are based on 1,000 replications per scenario.

unconditional elasticity is an equilibrium object, bundling pure substitution with other margins depending on market structures. Its value is context-specific (e.g., production function, market structure) and can even be negative.[23]

I illustrate this bias in the nested–CES example of (8). Consider a short-run setting in which the capital stocks K_C, K_D are quasi-fixed, while firms choose (C, D) to maximize profits, so output Y is endogenous to price shocks. An exogenous increase in the dirty price P_D induces a negative scale effect. With K_C, K_D fixed, this scale adjustment operates through changes in energy uses C, D, which reallocate output shares across the inner and outer nests. Consequently, the clean–dirty elasticity σ_{CD}, as a share-weighted average, shifts due to the scale effect (e.g., toward the clean inner-nest elasticity as the clean machine is binding). CE-estimate therefore recovers the unconditional elasticity that embeds both substitution and scale-driven reallocation. Using such a CE elasticity as an input parameter in counterfactuals that already allow for profit-driven output adjustment double-counts the scale channel and biases the predicted effects.

I run a targeted simulation with a nested-CES production. Firms face similar relative prices (x-axis) but maximize profits over inputs, so scale adjustments operate alongside substitution. For each simulation, I report the CE-estimate, i.e., the unconditional elasticity, and the true

[23]Similar to the unconditional elasticity, the Allen–Uzawa and Morishima elasticities are also equilibrium objects. They are defined on the cost function under cost minimization with all inputs freely adjustable. However, CE-estimates are specific to the empirical context and need not coincide with those cost-based objects. I thus label it *unconditional* to avoid conflating.

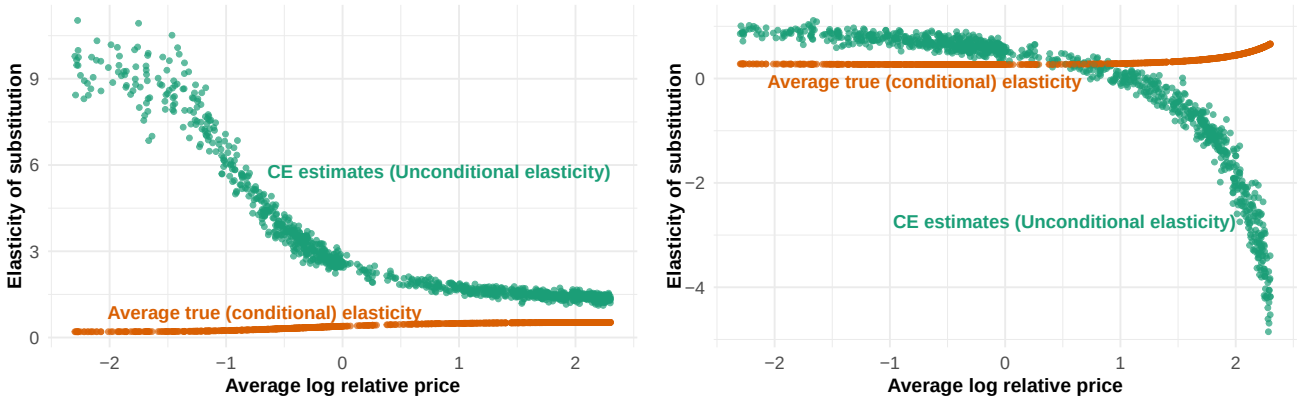


Figure 2: Monte Carlo simulation of unconditional elasticity bias in CE estimation

Notes: Each Monte Carlo replication simulates $n = 1000$ firms producing with a nested CES: $CES(CES(X_1, Z_1), X_2)$ so that elasticities of substitution between X_1, X_2 is not a constant. Firms draw Z_1 i.i.d. and choose (X_1, X_2) to maximize profit given prices. All firms face similar relative prices with the average price randomly drawn and shown on the x-axis. CE-estimates and the true average elasticities are plotted based on y-axis. In the left panel, the inner nest is a complements aggregation ($\sigma = 0.2$), and the outer nest is a substitutes aggregation ($\sigma = 2$). I exchange the inner and outer nests' elasticities in the right panel. Results are based on 1,000 replications per scenario.

conditional elasticity. Figure 2 shows two cases that swap the inner- and outer-nest elasticities. Three messages emerge: (i) the unconditional–conditional gap can be quantitatively large; (ii) its sign is context-dependent; and (iii) the unconditional elasticity can even be negative in this Monte Carlo with a standard market setup.

The unconditional elasticity can be decomposed into conditional elasticity and adjustment effects if the production function and market structure are estimated. For instance, consider price-taking firms that maximize profit in X_1, X_2 with all other inputs fixed. Let $f_i \equiv \partial lnF/\partial lnX_i$ and $f_{ij} \equiv \partial^2 lnF/\partial lnX_i \partial lnX_j$. Then [24]

$$\sigma = \frac{f_1 + f_2}{f_1 + f_2 + 2f_{12} - f_{11}\frac{f_2}{f_1} - f_{22}\frac{f_1}{f_2}}, \qquad \sigma^{UN} = \frac{\sum_i f_{1i} + \sum_i f_{2i}}{\left[1 + \frac{f_{12}}{f_1} - \frac{f_{22}}{f_2}\right]\sum_i f_{1i} + \left[1 + \frac{f_{12}}{f_2} - \frac{f_{11}}{f_1}\right]\sum_i f_{2i}}, \tag{12}$$

Case D: *Imperfectly assigned p_{jt}.* When relative prices are not exogenously shocked, the CE estimand continues to be a weighted average of unconditional elasticities, and imperfect assignment can further generate non-convex (possibly negative) weights.

A common design is the shift–share (Bartik) instruments that distribute aggregate shocks by predetermined shares. Credible identification requires orthogonality in shock space and reliability of shares. Otherwise, endogenous or mismeasured shares, heterogeneous shocks, and spillovers can distort, or flip the sign of weights (Bartik, 1991; Goldsmith-Pinkham *et al.*, 2020; Adão *et al.*, 2019; Borusyak *et al.*, 2022). Furthermore, elasticity estimated based on a DiD with heterogeneous price shock timing may combine "cohort-by-event-time" unconditional substitution responses using non-convex weights (Goodman-Bacon, 2021; Sun and Abraham, 2021; Callaway and Sant'Anna, 2021; de Chaisemartin and D'Haultfœuille, 2020, 2022). I do not pursue the further details, which are outside the scope of this paper.

5 Hicks-Neutral (HN) assumption

This section formalizes the *Hicks-Neutral (HN)* approach that underlies much of the production–function literature, presents estimable equations, states identification conditions, and clarifies the bias when factor–biased productivity is misspecified.

5.1 Setup and review

Assumption HN. *There is no variation in factor–biased productivities:* $\Omega_{i,jt} \equiv 1$.

[24] Derivations are provided in Appendix A.3.

Under HN, firms differ only through a Hicks-neutral productivity term Ω_R. Let

$$Y_{jt} \;=\; \Omega_{R,jt}\, F\big(X_{1,jt},\, X_{2,jt},\, \boldsymbol{Z}_{jt}\big) \quad \Longrightarrow \quad y_{jt} \;=\; f\big(x_{1,jt},\, x_{2,jt},\, \boldsymbol{z}_{jt}\big) \;+\; \overbrace{\underbrace{\omega_{R,jt}}_{\substack{\text{Hicks-neutral}\\ \text{productivity}}} \;+\; \underbrace{\varepsilon_{jt}}_{\text{noise}}}^{\text{Unobserved}} \tag{13}$$

where lower case denotes logs. Because fixed–effects may not fully absorb $\omega_{R,jt}$ and valid instruments for inputs are difficult to find (Griliches and Mairesse, 1995; Ackerberg *et al.*, 2007), recent work relies on control approaches to identify $f(\cdot)$. The procedure involves two stages. The first stage constructs a control for unobserved productivity $\omega_{R,jt}$ using observables. This can be done either (i) via a proxy (e.g., investment/materials), as in Olley and Pakes (1996); Levinsohn and Petrin (2003); Ackerberg *et al.* (2015), or (ii) via an expenditure share equation, as in Gandhi *et al.* (2020). The second stage imposes a timing restriction based on Markov $\omega_{R,jt}$ and pins down the production function and the productivity law of motion by GMM, using quasi–fixed inputs and lagged behavior as instruments.

First stage. The proxy variable approach is built on the following assumption.

Assumption 1. *There exists an observed proxy v_{jt} (e.g., investment or materials) such that its demand is $v_{jt} = V\big(\omega_{R,jt},\, X_{1,jt},\, X_{2,jt},\, \boldsymbol{Z}_{jt}\big)$ and V is strictly monotone in $\omega_{R,jt}$.*

Assumption 1 implies an inversion $\omega_{R,jt} = V^{-1}\big(v_{jt}, X_{1,jt}, X_{2,jt}, \boldsymbol{Z}_{jt}\big)$. Substituting in (13) yields a first–stage regression that absorbs $\omega_{R,jt}$ into a flexible control:

$$\begin{aligned} y_{jt} &= f\big(\, x_{1,jt},\; x_{2,jt},\; \boldsymbol{z}_{jt}\,\big) + \underbrace{V^{-1}(v_{jt}, X_1,\; X_2,\; \boldsymbol{Z})}_{\omega_{R,jt}} + \varepsilon_{jt} \\ &= \Phi\big(\, v_{jt}, x_{1,jt},\; x_{2,jt},\; \boldsymbol{z}_{jt}\,\big) + \varepsilon_{jt} \end{aligned}$$

where Φ can be parameterized with rich time/industry/region interactions to reflect market variations. Given estimate $\widehat{\Phi}$, the fitted productivity expression is

$$\widehat{\omega}_{R,jt} \;=\; \widehat{\Phi}\big(v_{jt}, x_{1,jt}, x_{2,jt}, \boldsymbol{z}_{jt}\big) \;-\; f\big(x_{1,jt}, x_{2,jt}, \boldsymbol{z}_{jt}\big). \tag{14}$$

Alternatively, the expenditure share approach is based on:

Assumption 2. *For flexible input(s) $v_{jt} \in \{x_{1,jt}, x_{2,jt}, \boldsymbol{z}_{jt}\}$, the firm is price–taking in v_{jt} and chooses v_{jt} to maximize profit.*

Under Assumption 2, the estimation is built on log expenditure share of v_{jt}:

$$lnS_{v,jt} \;=\; \Phi\big(x_{1,jt}, x_{2,jt}, \boldsymbol{z}_{jt}\big) \;-\; \varepsilon_{jt}, \qquad \Phi(\cdot) \equiv ln\Big(\frac{\partial y_{jt}}{\partial v_{jt}}\Big), \tag{15}$$

where the Hicks-neutral term cancels.[25] Estimating (time/industry/region-specific) $\widehat{\Phi}$ identifies the marginal product in v. by the Fundamental Theorem of Calculus, fitted productivity is

$$\widehat{\omega}_{R,jt} \;=\; y_{jt} \;-\; \widehat{\varepsilon}_{jt} \;-\; \int \exp\big[\widehat{\Phi}(v, \boldsymbol{x}_{jt})\big]\, dv \;-\; g(\boldsymbol{x}_{jt}), \tag{16}$$

for some function $g(\cdot)$ of the remaining inputs $\boldsymbol{x}_{jt} = \{x_{1,jt},\ x_{2,jt},\ \boldsymbol{z}_{jt}\} \setminus \{v_{jt}\}$.

Second stage. The production function is identified by the timing assumption:

Assumption 3. *$\omega_{R,jt}$ follows an ith–order Markov process, so that $\omega_{R,jt} - \mathbb{E}[\omega_{R,jt} \mid \{\omega_{R,j,t'}\}_{t'=t-i}^{t-1}]$ is unpredictable, and thus independent of the information set $\mathcal{I}_{jt}$ (behaviors decided before the realization of $\omega_{R,jt}$).*

In practice, information set $\mathcal{I}_{jt}$ includes predetermined inputs and lagged choices. Using (14) or (16), identification proceeds via GMM with moment conditions

$$\mathbb{E}\Big[\Big(\widehat{\omega}_{R,jt} - \mathbb{E}[\widehat{\omega}_{R,jt} \mid \{\widehat{\omega}_{R,j,t'}\}_{t'=t-i}^{t-1}]\Big) \cdot Z_{jt}\Big] = 0, \qquad Z_{jt} \in \mathcal{I}_{jt}, \tag{17}$$

which pin down $f(\cdot)$ and the productivity law of motion. Once $f(\cdot)$ is estimated, individual input elasticities follow from expression (12).

CE versus HN. Relative to CE regressions, production function methods allow rich functional forms and heterogeneous elasticities, but they require HN to avoid unobservables inside the production function. CE approaches, by contrast, can accommodate factor–biased productivity if relative prices are exogenous. This complementarity motivates the integrated strategy developed in Section 6, which relaxes both restrictions simultaneously.

5.2 Why factor-biased productivities exist

Heterogeneity in factor-augmenting productivity arises along at least three well-documented dimensions. (i) *Technology*: Directed technical change and the adoption of factor-biased technologies differentially raises the effective productivity across inputs (Acemoglu, 2002; Acemoglu *et al.*, 2012; Krusell *et al.*, 2000a; Autor *et al.*, 2003; Beaudry *et al.*, 2010). (ii) *Management*: Cross-firm differences in managerial practices can elevate effective labor and materials productivity, and complement the impact of factor-biased technical change (Bresnahan *et al.*, 2002; Bloom *et al.*, 2012; Syverson, 2011). (iii) *Input heterogeneity and infrastructure*: Unobserved variation in factor quality (e.g., worker skill, experience) and in local infrastructure (e.g., grid capacity, renewables

[25] To streamline the exposition, I assume that ε_{jt} is measurement error or unpredictable shock realize after the decision of flexible input v_{jt} so that $\mathbb{E}[\exp(v_{jt})]$ does not show up in the first-order condition.

access, storage and transmission constraints) produces persistent differences in efficiency of particular inputs (Katz and Murphy, 1992; Syverson, 2011; Joskow, 2011; Borenstein, 2012; Gonzales *et al.*, 2023).

5.3 Intuition for bias

I now consider settings in which the HN assumption is misspecified and characterize the resulting distortion in elasticity estimates. I begin with an intuition result under an innocuous CE benchmark.

Intuition for overestimation. Assume the true technology aggregates (X_1, X_2) via a CES, so that the relative input demand decomposes into a price component and a factor–biased productivity component:

$$\underbrace{ln\frac{X_{2,jt}}{X_{1,jt}}}_{\text{Demand variations}} = \alpha + \underbrace{\sigma\, ln\frac{P_{1,jt}}{P_{2,jt}}}_{\text{Price variations}} + \underbrace{(\sigma-1)\, ln\frac{\Omega_{2,jt}}{\Omega_{1,jt}}}_{\text{Productivity variations}} \tag{18}$$

Let $x_{jt} \equiv ln(X_{2,jt}/X_{1,jt})$, $p_{jt} \equiv ln(P_{1,jt}/P_{2,jt})$, and $\omega_{jt} \equiv ln(\Omega_{2,jt}/\Omega_{1,jt})$. Variance decomposition implies that

$$\text{Var}(x_{jt}) = \sigma^2\text{Var}(p_{jt}) + \underbrace{(\sigma-1)^2\text{Var}(\omega_{jt}) + \sigma(\sigma-1)\text{Cov}(p_{jt},\omega_{jt})}_{\text{Shut down by HN}} \tag{19}$$

Estimation under HN impose $\omega_{jt} \equiv 0$ in (19) and attributes variation in x_{jt} entirely to p_{jt}. Unless the covariance term in (19) is sufficiently negative to offset $(\sigma-1)^2\text{Var}(\omega_{jt})$, the productivity contribution in brace is strictly positive. Hence, the elasticity of substitution must be *overestimated* to let price-driven variation match demand variance.

Example. To gauge the sign of "HN–shut–down" term in (19), consider a parsimonious system with both endogenous factor supplies and endogenous (e.g., directed) factor–biased productivity.[26] Relative prices respond to an exogenous supply shifter ρ_{jt}^{exo}, and input demand with supply elasticity $\rho > 0$: $p_{jt} = \rho_{jt}^{\text{exo}} - \rho\, x_{jt}$. Relative factor–biased productivity responds to an exogenous productivity shifter θ_{jt}^{exo}, and relative profitability of developing productivity with elasticity $\theta > 0$, as in Acemoglu (2002): $\omega_{jt} = \theta_{jt}^{\text{exo}} + \theta\big(-\frac{\omega_{jt}}{\sigma} + \frac{\sigma-1}{\sigma}x_{jt}\big)$. Suppose $\rho > \theta$.

The "HN–shut–down" term is negative only when (i) Two inputs are complements but very close to unit elasticity. (ii) Strong productivity-driven shocks relative to supply shocks. For instance, when $\rho = 0.1$, $\theta = 0.05$, the requirements are (i) $0.922 < \sigma < 1$. (ii) $\text{Var}(\rho_{jt}^{\text{exo}})/\,\text{Var}(\theta_{jt}^{\text{exo}}) <$

[26] See Section B.1 for more details.

0.171 as $\sigma = 0.93$, or $V_\rho/V_\theta < 2.264$ as $\sigma = 0.97$. Overall, the "HN–shut–down" term should be typically positive.

5.4 Bias

This subsection formally investigates the bias when HN assumption is violated.

Setup. Let log output be generated by a well-behaved production function[27]

$$y_{jt} \;=\; f\big(\,\omega_{1,jt} + x_{1,jt},\; \omega_{2,jt} + x_{2,jt},\; z_{jt}\,\big)$$

with (log) flexible inputs x_i and their factor-augmenting productivities ω_i, and (log) quasi-fixed inputs z. For expositional clarity, I omit the hicks-neutral productivity (think of it as perfectly controlled in an ideal HN estimation) and normalize $\mathbb{E}[\omega_i] = 0$. Firms choose x_1, x_2 to maximize profit given $(\omega_1, \omega_2, p_1, p_2, z)$, following the continuous and differentiable joint density $\mathbb{Q}_p(\omega_1, \omega_2, p_1, p_2, z)$ with compact support where p_i denotes log prices.

Under HN, one estimates f by a function g that projects y onto observed inputs, i.e.

$$\min_{g(\cdot)} \mathbb{E}\Big[\big(y_{jt} - g(x_{1,jt}, x_{2,jt}, z_{jt})\big)^2\Big] \quad \Rightarrow \quad g(x_1, x_2, z) = \mathbb{E}[f(\omega_1 + x_1, \omega_2 + x_2, z) \mid x_1, x_2, z]$$

I call the elasticity and substitutability computed from g via the usual formula (12) the *HN estimator:*

$$\hat{\sigma}^{\text{HN}}(x_1, x_2, z) = \frac{g_{x_1} + g_{x_2}}{g_{x_1} + g_{x_2} + 2g_{x_1x_2} - g_{x_1x_1}\frac{g_{x_2}}{g_{x_1}} - g_{x_2x_2}\frac{g_{x_1}}{g_{x_2}}}, \quad \hat{\lambda}^{\text{HN}} = \frac{\hat{\sigma}^{\text{HN}} - 1}{\hat{\sigma}^{\text{HN}}}. \tag{20}$$

A tractable two-input-CES benchmark. To isolate the mechanism, suppose X_1 and X_2 enter through a CES aggregator (I suppress jt and z for brevity):

$$F\big(\Omega_1 X_1, \Omega_2 X_2\big) \;=\; \Big[\alpha_1(\Omega_1 X_1)^{\frac{\sigma-1}{\sigma}} + \alpha_2(\Omega_2 X_2)^{\frac{\sigma-1}{\sigma}}\Big]^{\frac{s_R\sigma}{\sigma-1}}, \qquad \Omega_i = e^{\omega_i},\; \alpha_i > 0,\; \alpha_1 + \alpha_2 = 1,$$

so the true elasticity between X_1 and X_2 is the constant σ, and s_R denotes the scale elasticity. Let S_i denote the output share of X_i, with $S_1 + S_2 = s_R < 1$. By first-order conditions of profit maximization, $p_i = \ln F_{X_i}$, and thus one can map densities from (ω, p) to (ω, x). Bayes' rule thus implies the conditional density of productivities is $\mathbb{P}_\omega(\omega_1, \omega_2 \mid x_1, x_2)$. By Fisher's identity, derivatives of g can be computed, yielding the HN-estimate.[28]

[27] f is four times continuously differentiable and satisfies "well-behaved" regularity in (McFadden, 1963)

[28] See Appendix B.2 for the derivation.

In what follows, all expectations and covariances are understood to be conditional on (x_1, x_2). HN substitutability estimate can be decomposed as:

$$\hat{\lambda}^{\mathrm{HN}} = \underbrace{\frac{\mathbb{E}[S_1 S_2]}{g_{x_1} g_{x_2}} \cdot \frac{g_{x_1} + g_{x_2}}{s_R} \cdot \lambda}_{\text{Averaging curvature}} + \underbrace{2\left(\frac{\mathrm{Cov}\big(S_1,\, \partial_{x_1} \ln \mathbb{P}_\omega\big)}{g_{x_1}} + \frac{\mathrm{Cov}\big(S_2,\, \partial_{x_2} \ln \mathbb{P}_\omega\big)}{g_{x_2}}\right)}_{\text{Substitution-selection bias}} + \underbrace{\frac{1}{g_{x_1} + g_{x_2}}\left(R_{11}\frac{g_{x_2}}{g_{x_1}} + R_{22}\frac{g_{x_1}}{g_{x_2}} + R_{12}\right)}_{\text{Quadratic/score curvature}}. \tag{21}$$

The first bracket is the bias of curvature attenuation. The second bracket is the *productivity selection bias*. Intuitively, firms that use X_1 more intensively typically do so because their production technology is more X_1-biased, not only because X_1 is exceptionally cheap. A marginal-product-only specification treats this technological selection as if it were heightened price sensitivity. The final bracket collects second–order Fisher curvature and score–variance terms, and does not overturn the bias implications under mild conditions.

I first state the main proposition and then examine the sources of the bias.

Assumption 4 (Scale elasticity bias)**.** *Let* $\varsigma \equiv (g_{x_1} + g_{x_2}) - s_R > 0$. *For some small* $\nu > 0$,

$$\big|\, g_{x_1} - g_{x_2} - \mathbb{E}[S_1 - S_2]\big(1 + \varsigma/s_R\big)\,\big| < \nu\,\varsigma.$$

Assumption 5 (Substitution–selection)**.** $\sum_i \mathrm{Cov}\big(S_i,\ \partial_{x_i} \ln \mathbb{P}_\omega\big) > 0$ *holds on a non-negligible set of* (x_1, x_2) .

Assumption 6 (Quadratic/score curvature)**.** $R_{ii} \geq 0$ *for* $i = 1, 2$, *and* $R_{12} \leq \sqrt{R_{11} R_{22}}$.

Proposition 3. *For* (x_1, x_2) *satisfying Assumptions 4–6, the HN estimator exhibits two biases:*

(i) Averaging curvature *that pushes* $\hat{\sigma}^{\mathrm{HN}}$ *toward one.*

(ii) Substitution–selection *that inflates measured substitutability. Consequently, locally,* $\hat{\sigma}^{\mathrm{HN}} \geq \sigma$, *and for sufficiently strong selection,* $\hat{\lambda}^{\mathrm{HN}}$ *can exceed one, implying a negative* $\hat{\sigma}^{\mathrm{HN}}$.

No selection. Even when there is no selection through the conditional density $\mathbb{P}_\omega$ (e.g., when $\mathbb{Q}_x(\omega_1, \omega_2, x_1, x_2)$ is independent of (x_1, x_2)), HN-estimate is biased due to the nonlinearity of the production function (averaging curvature). In particular, the HN substitutability parameter $\lambda \equiv (\sigma - 1)/\sigma$ is attenuated toward zero:

$$\hat{\lambda}^{\mathrm{HN}} \;=\; \xi(x_1, x_2) \cdot \frac{\sigma - 1}{\sigma}, \qquad \xi(x_1, x_2) \;=\; 1 - \frac{\mathrm{Var}(S_1 \mid x_1, x_2)}{\mathbb{E}[S_1 \mid x_1, x_2]\,\mathbb{E}[S_2 \mid x_1, x_2]} \in (0, 1], \tag{22}$$

Hence, every HN elasticity estimate is biased toward one relative to the true σ. The bias converges to zero only when the true production function is log-linear (i.e., Cobb-Douglas $\sigma \to 1$) or when there is no variation in relative factor-biased productivity conditional on observed inputs.

With selection bias. When $\partial_{x_i} \ln \mathbb{P}_\omega \neq 0$, the selection bias further biases the estimate:

Bias in scale elasticity. A useful intermediate object is the estimated scale elasticity:

$$g_{x_1} + g_{x_2} = s_R + \underbrace{(s_R - 1)\ \mathrm{Cov}\Big(f,\ \sum_{i=1}^{2} \partial_{p_i} \ln \mathbb{Q}_p\ \Big)}_{\text{Scale–selection bias}},$$

Because higher prices (holding x fixed) typically occur in the right tail of $\mathbb{Q}_p$, the score $\sum_i \partial_{p_i} \ln \mathbb{Q}_p$ is generally negative and more so at higher f. This yields a positive covariance term in standard environments, in line with the "transmission bias" intuition of Olley and Pakes (1996): more productive firms choose higher inputs, so $g_{x_1} + g_{x_2}$ is upward–biased.

Averaging curvature. The magnitude of curvature attenuation involving selection depends on how the upward bias in the estimated scale elasticity, $g_{x_1} + g_{x_2}$, is allocated across the two inputs. By Assumption 4, I focus on regular cases without highly asymmetric allocation, implying $g_{x_1}, g_{x_2} > 0$, as low or negative input elasticity naturally discounts the credibility of estimates. Algebraically, averaging curvature bias is

$$\frac{\mathbb{E}\big[S_1 S_2\big]}{g_{x_1} g_{x_2}}\ \frac{g_{x_1} + g_{x_2}}{s_R} \cdot \lambda\ =\ \underbrace{\frac{\mathbb{E}[S_1]\mathbb{E}[S_2]}{g_{x_1} g_{x_2}}\ \frac{g_{x_1} + g_{x_2}}{s_R}}_{\text{Allocation of scale–selection bias}}\ \cdot \xi\ \lambda,$$

where $\xi \in (0, 1]$ is the curvature factor defined in (22). Under Assumption 4, the allocation term is strictly below $\frac{1}{\xi}$, so the implied $\hat{\sigma}^{\mathrm{HN}}$ is still pulled toward one.

Substitution–selection bias. This is the key force inflating $\hat{\lambda}^{\mathrm{HN}}$. It decomposes into a relative–price (substitution) component and a level–price (scale) component:

$$\begin{aligned}
&\frac{\mathrm{Cov}\big(S_1,\ \partial_{x_1} \ln \mathbb{P}_\omega\big)}{g_{x_1}}\ +\ \frac{\mathrm{Cov}\big(S_2,\ \partial_{x_2} \ln \mathbb{P}_\omega\big)}{g_{x_2}} \\
&\quad =\ \underbrace{(1 - \lambda)\,\mathrm{Cov}\left(\Delta S,\ \frac{\partial_{p_1} \ln \mathbb{Q}_p}{g_{x_1}} - \frac{\partial_{p_2} \ln \mathbb{Q}_p}{g_{x_2}}\right)}_{\text{Relative–price (substitution) selection}} \\
&\quad +\ \underbrace{\frac{1}{2}\left(1 - \frac{\lambda}{s_R}\right)\Big[s_R\big(g_{x_1}^{-1} - g_{x_2}^{-1}\big)\,\mathrm{Cov}\Big(\Delta S,\ \textstyle\sum_i \partial_{p_i} \ln \mathbb{Q}_p\Big) + \big(g_{x_1}^{-1} + g_{x_2}^{-1}\big)\,\mathrm{Cov}\Big(\Delta S,\ \big[\textstyle\sum_i \partial_{p_i} \ln \mathbb{Q}_p\big]\Delta S\Big)\Big]}_{\text{Level–price (scale) selection}},
\end{aligned}$$

where $\Delta S \equiv S_1 - S_2$. Conditional on (x_1, x_2) and productivities, a larger expenditure tilt toward

X_1 (higher ΔS) is typically sustained by a lower relative price $p_1 - p_2$, with at most a modest move in the overall price level $p_1 + p_2$. Hence the substitution term is typically positive and dominant because higher ΔS co–occurs with lower $p_1 - p_2$, tilting the posterior $\mathbb{P}_\omega$ toward X_1–biased technical change consistent with the observed mix. Assumption 5 rules out joint densities that would make the scale component negative and dominant.[29] Furthermore, as the dispersion of factor–biased productivities rises, both $\mathrm{Var}(\Delta S)$ and the dispersion of price scores increase, amplifying the substitution component and pushing $\hat{\lambda}^{\mathrm{HN}}$ further upward.

5.5 Bias under joint assumptions

The previous subsection considered flexible estimation under HN and showed that $\hat{\sigma}^{\mathrm{HN}}$ is typically *locally* upward biased. In much applied work, however, researchers may impose *both* HN and CE by specifying a CES aggregator for (X_1, X_2) and setting $\Omega_i \equiv 1$ in estimation. Under these joint restrictions the reported elasticity is a *single, global* number. This section shows why the joint specification tends to *globally overstate* substitutability.

A tractable two-input CES benchmark. To isolate the mechanisms, suppose the true technology (suppressing jt and z) is

$$F(\Omega_1 X_1, \Omega_2 X_2) = \Big[\alpha_1(\Omega_1 X_1)^{\frac{\sigma-1}{\sigma}} + \alpha_2(\Omega_2 X_2)^{\frac{\sigma-1}{\sigma}}\Big]^{\frac{s_R \sigma}{\sigma-1}}, \qquad \Omega_i = e^{\omega_i},\ \alpha_i > 0,\ \alpha_1 + \alpha_2 = 1.$$

Under HN and CE, a total derivative on the marginal rate of substitution delivers the compact approximation:

$$\hat{\sigma}^{\mathrm{HN+CE}} \approx \left(\mathbb{E}\left[\frac{d\ln(F_{X_1}/F_{X_2})}{d\ln(X_2/X_1)}\right]\right)^{-1} \approx \sigma\Big[\underbrace{1}_{\text{Price effect}} + \underbrace{(\sigma-1)\,\xi}_{\text{Averaging curvature}} \underbrace{-\,\mathbb{E}\Big(\frac{(1-\sigma)\,d\ln(\Omega_1/\Omega_2)}{d\ln(X_2/X_1)}\Big)}_{\text{Productivity selection}}\Big]^{-1}, \tag{23}$$

where $\xi \in (0,1)$ is a curvature term stemming from the non-log-linearity of the CES aggregator, which pulls estimates toward one. The last term is the share of productivity-driven variation in relative inputs: Based on Equation (19),

$$\mathbb{E}\left[\frac{(1-\sigma)\,d\ln(\Omega_1/\Omega_2)}{d\ln(X_2/X_1)}\right] \approx \frac{(1-\sigma)^2\,\mathrm{Var}(\omega) \;+\; \sigma(1-\sigma)\,\mathrm{Cov}(\omega, p)}{\mathrm{Var}(x)} \in [0,1), \tag{24}$$

Absent an implausibly large negative $\mathrm{Cov}(\omega, p)$, the share is non-negative.

[29] Locally, the scale term can matter when (i) one input's scale elasticity g_{x_i} is much larger than the other's (ruled out by Assumption 4), and/or (ii) $\mathbb{Q}_p$ makes the overall price level co–move strongly with expenditure composition.

Equation (23) delivers several testable predictions. First, when factor-biased productivity dispersion is substantial, the productivity share is largely positive, implying systematic overstatement of elasticity. Second, when dispersion in factor-biased productivity is limited, the curvature term dominates and pulls estimates toward one. Consequently, if $\sigma > 1$ (substitutes), modest underestimation can occur. Third, because the productivity share is bounded by one, negative elasticity estimates typically arise only when the true $\sigma < 1$ (averaging curvature is negative).[30][31]

I run a targeted simulation with CES production and input-biased productivity. The true elasticity is either 0.5 (complements) or 2 (substitutes). Holding other primitives fixed, I vary only the dispersion of the factor-biased productivity (x-axis). For each simulation, I plots the HN-estimate alongside the true elasticity. I repeat this under two price-variation regimes: baseline and a lower-variance case that increases the relative role of productivity. The results in Figure 3 confirm all testable predictions.

Furthermore, I show that the joint restrictions CE+HN can *partially pin down* the estimated elasticity when factor–biased productivity is highly dispersed.

An example with large factor-biased heterogeneity. Suppose clean and dirty are truly complements, and there are two firm types: (i) *dirty–biased* (very high Ω_C/Ω_D), and (ii) *clean–biased* (very low Ω_C/Ω_D), so that energy mixes are extremely dispersed even though price dispersion is modest. Under CE, the two inputs enter a CES aggregation $CES(\Omega_C X_C, \Omega_D X_D; \sigma)$ (possibly inside a larger aggregator with other inputs). If one further imposes HN, the model must rationalize observations with $X_C \approx 0$ without attributing them to a very high Ω_C/Ω_D.

Hence, for a candidate CES with $\sigma \leqslant 1$ (gross complements), the aggregator penalizes asymmetry at the boundary:

$$\lim_{X_C \to 0^+} CES(X_C, X_D; \sigma) = 0 \quad \text{and} \quad \lim_{X_C \to 0^+} \frac{\partial CES}{\partial X_D}(X_C, X_D; \sigma) \equiv 0,$$

so the marginal contribution of X_D collapses when X_C is scarce. Economically, at finite prices, a firm would not purchase any X_D when $X_C \approx 0$. The only remaining fit is to raise substitutability: choosing $\sigma > 1$ keeps $\partial CES/\partial X_D$ bounded away from zero and lets X_D "carry" the aggregation. Hence, in data with substantial variations in factor-biased productivities, the joint CE+HN restrictions push the estimate toward $\hat{\sigma} > 1$.

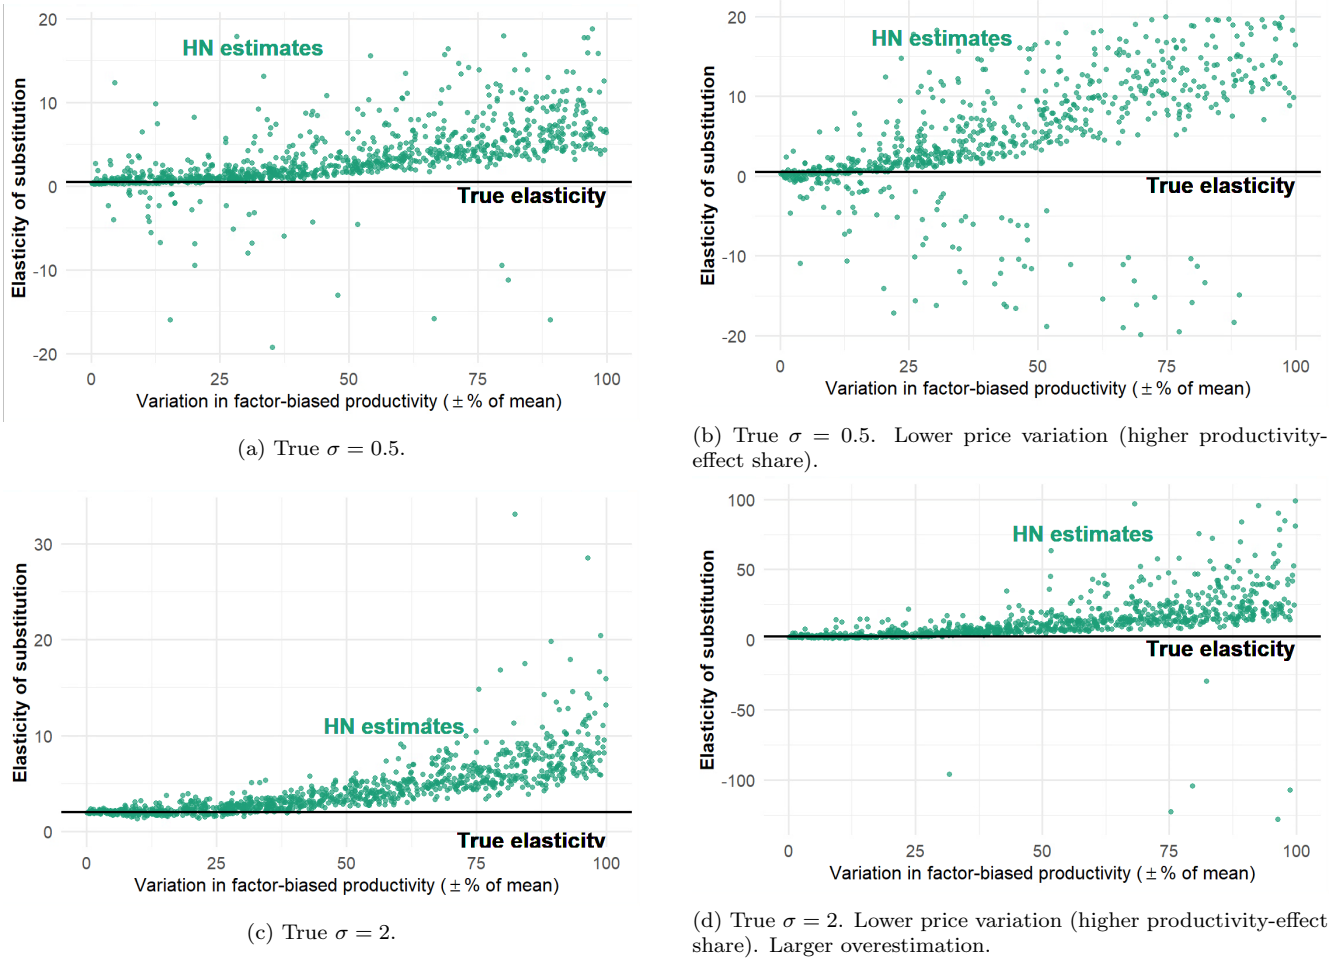
[30] A "negative" elasticity σ means substitutability $\lambda = (\sigma - 1)/\sigma$ has gone beyond one, the perfect-substitutes limit. It implies a *concave* isoquant, and thus it is a misspecification flag rather than a meaningful technology.

[31] Furthermore, when $\sigma < 1$ and the global estimate is negative, as the productivity share rises, the marginal impact on the global estimate is ambiguous: locally negative elasticities move toward zero (intensive margin), but more evaluation points can cross into the negative-elasticity region (extensive margin).

5.6 Limitations.

I focus on the X_1-X_2 elasticity while holding other inputs fixed and exploiting a profit-maximization mapping from (ω, p) to (ω, x). With additional flexible inputs and cost minimization, equilibrium shares are co-determined across all dimensions, and characterizing HN bias requires joint restrictions on the full production structure. Also, the analysis is silent on the spillover effect into $\hat{\sigma}$ by omitting factor-biased productivities of the inputs other than (X_1, X_2).[32]

[32]These issues may not alter the core message, but they complicate closed-form characterizations. I therefore leave a general multi-input treatment to future work.



(a) True $\sigma = 0.5$.

(b) True $\sigma = 0.5$. Lower price variation (higher productivity-effect share).

(c) True $\sigma = 2$.

(d) True $\sigma = 2$. Lower price variation (higher productivity-effect share). Larger overestimation.

Figure 3: Monte Carlo simulation of bias in HN estimation

Notes: Each Monte Carlo replication simulates $n = 1000$ firms producing with a CES: $CES(X_1, X_2)$ without hicks-neutral productivity. Firms draw prices and factor-biased productivity i.i.d. and choose (X_1, X_2) to maximize profit. The variation in factor-biased productivity is randomly draw and plotted on x-axis. HN-estimates under joint assumptions, and the true average elasticities are plotted based on y-axis. In the right two panels, variations in prices are reduced to 50% so that productivity-effect share is higher. Results are based on 1,000 replications per scenario.

I study bias under a direct least–squares fit of output that minimizes residuals from a flexible $g(x_1, x_2, \boldsymbol{z})$, while omitting Hicks-neutral productivity. Extending the analysis to the control–function approach is substantially more involved because both stages can be biased. If the first stage were correctly specified, the identification of the second stage in control approaches primarily exploits within-firm time variation (Gandhi *et al.*, 2020), which tempers inflated NFBP bias caused by cross-sectional technological heterogeneity in simple projections. However, first-stage controls are themselves biased when factor-biased productivities vary: proxies can be weak or non-monotone in the relevant productivity components, and input and capital mismeasurement further contaminate the constructed controls, propagating bias into the second stage (de la Cal Medina, 2025; Collard-Wexler and De Loecker, 2025). Accordingly, NFBP-based estimates may remain upward-biased (and occasionally produce implausible negatives) even in control-function implementations.

6 General method

The complementarity between CE and HN estimators motivates an integrated approach that relaxes both restrictions simultaneously. This section presents the method.

6.1 Setup

I consider an estimation environment close to recent work on factor–biased productivity (Doraszelski and Jaumandreu, 2018; Zhang, 2019; Demirer, 2022).

Assumption 7 (Setup)**.** *Firms add value to intermediates using predetermined capital K and a vector of flexible inputs $\boldsymbol{X}^f$ according to the value–added function*

$$VA_{jt} \equiv \Omega_{Rjt} \left[F\big(K_{jt}, (\boldsymbol{\Omega^f_{jt}})' \boldsymbol{X^f_{jt}}\big) \right]^{s_R}, \tag{25}$$

where $F(\cdot)$ is constant–returns–to–scale (CRS) and s_R is the scale elasticity. Ω_{Rjt} is Hicks-neutral (revenue) productivity and $\boldsymbol{\Omega}^f_{jt}$ collects factor–biased productivities on flexible inputs. Capital's productivity is absorbed into Ω_{Rjt}, so factor–biased productivities are interpreted relative to K. Within period t, firms take K_{jt} and productivity vector as given and choose $\boldsymbol{X}^f_{jt}$ to maximize profits as price–takers in the flexible–input markets.

Assumption 8 (Estimation)**.** *The researcher observes K_{jt}, $\boldsymbol{X}^f_{jt}$, flexible–input expenditures $\boldsymbol{E}^f_{jt}$, intermediates M_{jt}, and value added with noise $VA^O_{jt} = VA_{jt} \exp(\varepsilon_{jt})$, where ε_{jt} is a measurement error or unpredictable shock that does not enter the firm's choice problem. All productivity parameters are unobserved.*

In practice, observing both prices and quantities for all flexible inputs can be difficult. The framework can accommodate data limitations by allowing parsimony in the set of factor–biased productivities.

A central challenge is endogeneity from unobserved productivity. Under Assumption 7, gross output is Leontief in intermediates and a production aggregator, which implies that materials demand is informative about Hicks-neutral productivity. I therefore adopt a proxy–variable assumption:

Assumption 9 (Proxy variable)**.** *Materials demand satisfies* $M_{jt} = V\big(\Omega_{R,jt},\, \boldsymbol{\Omega}^{f}_{jt},\, K_{jt},\, \boldsymbol{X}^{f}_{jt},\, \boldsymbol{E}^{f}_{jt}\big)$*, with* V *strictly monotone in* $\Omega_{R,jt}$*.*

Assumption 9 is a strengthened proxy–variable condition, made feasible by richer observables (flexible–input prices) and by explicitly allowing factor–biased productivities. Intuitively, two firms that are identical in all-dimensional productivities and input uses, and face similar input supply curves should have the same materials demand. Strict monotonicity then permits inversion of the materials demand to control for Hicks-neutral productivity.

Finally, I impose a standard timing restriction to separate predictable and unpredictable productivity components.

Assumption 10 (Timing assumption)**.** *Let* $\boldsymbol{\omega}_{jt} \equiv \ln \boldsymbol{\Omega}_{jt}$*. The vector* $\boldsymbol{\omega}_{jt}$ *follows an* i*th–order joint Markov process such that* $\boldsymbol{\omega}_{jt} - \mathbb{E}[\boldsymbol{\omega}_{jt} \mid \{\boldsymbol{\omega}_{j,t'}\}_{t'=t-i}^{t-1}]$ *is unpredictable and independent of the information set* $\mathcal{I}_{jt}$ *(the choices decided prior to the realization of* $\boldsymbol{\omega}_{jt}$*).*

Assumption 10 states that among firms with the same realized history of productivities, they are expected to be equally productive now. Thus, their current productivity differences are orthogonal to predetermined controls, i.e., information set including current K_{jt}, lagged inputs, etc.

I focus on estimating the substitution between input among demand and remain agnostic about the decomposition of revenue into production and output demand, a difficult task in practice (Klette and Griliches, 1996; van Dijcke, 2024).[33]

6.2 First-stage

The first stage controls for productivity. For any flexible input $X_i \in \boldsymbol{X}^f$, profit maximization implies that its expenditure share of value added equals its contribution share:

$$S_{X_i}\big(K_{jt}, (\boldsymbol{\Omega}^{\boldsymbol{f}}_{\boldsymbol{jt}})' \boldsymbol{X}^{\boldsymbol{f}}_{\boldsymbol{jt}}\big) \;\equiv\; \frac{d \ln VA_{jt}}{d \ln X_{i,jt}} \;=\; \frac{E_{X_i,jt}}{VA_{jt}}, \tag{26}$$

[33] When separating production from demand is the goal, additional structure is required, e.g., cost minimization (Hall, 1988, 1989; De Loecker *et al.*, 2020), CRS production (Grieco *et al.*, 2016), or isoelastic demand with gross output CRS in materials (Kroft *et al.*, 2024).

and the capital share follows from Euler's theorem for homogeneous functions

$$\begin{aligned} s_R\, VA_{jt} \;&=\; \frac{\partial VA_{jt}}{\partial K_{jt}}\, K_{jt} + \sum_{X_{i,jt}\in \boldsymbol{X}^f_{jt}} \frac{\partial VA_{jt}}{\partial X_{i,jt}}\, X_{i,jt} \\ \implies\; S_K\big(K_{jt}, (\boldsymbol{\Omega^f_{jt}})'\boldsymbol{X^f_{jt}}\big) \;&=\; s_R - \sum_{X_{i,jt}\in \boldsymbol{X}^f_{jt}} S_{X_i,jt}. \end{aligned} \tag{27}$$

The intuition is that a 1% increase in all inputs raises VA by s_R%.

To isolate productivity parameters, I look for functional forms under which factor-biased productivities are separable in (26)–(27).

Assumption 11 (Separable functional from)**.** *$F(\cdot)$ is either an arbitrary nested CES or a translog aggregator.*

Assumption 11 ensures that factor-biased productivities can be controlled by observables $(K_{jt}, \boldsymbol{X}^f_{jt}, \boldsymbol{S}^f_{jt})$ implied by (26)–(27). I will emphasize translog in the micro–empirical application. See Zhang (2025) for more discussion on nested CES.

To avoid control–approach nonidentification (Gandhi *et al.*, 2020), I adopt the "technical control" approach of Ackerberg *et al.* (2015). Inverting the materials demand in Assumption 9 gives $V^{-1}\big(M_{jt}, \boldsymbol{\Omega}^f_{jt}\, K_{jt}\, \boldsymbol{X}^f_{jt}, \boldsymbol{E}^f_{jt}\big)$ as a control for $\Omega_{R,jt}$. Since factor-biased productivities are themselves controlled by $(K_{jt}, \boldsymbol{X}^f_{jt}, \boldsymbol{S}^f_{jt})$, Hicks-neutral productivity can be eventually controlled by $(M_{jt}, K_{jt}, \boldsymbol{X}^f_{jt}, \boldsymbol{S}^f_{jt}, \boldsymbol{E}^f_{jt})$. This leads to a flexible first–stage regression:

$$va^O_{jt} \;=\; \phi\big(m_{jt}, k_{jt}, \boldsymbol{x}^f_{jt}, \boldsymbol{e}^f_{jt}\big) + \varepsilon_{jt}, \tag{28}$$

where lower–case variables denote logs. Applying time/industry/region interactions in $\phi(\cdot)$ can absorb market heterogeneity.

A practical complication is that $\boldsymbol{S}^f$ depends on expected value added, while only VA^O is observed.[34] Two workarounds are available. First, under Leontief gross output and competitive materials markets, M_{jt} scales expected VA (Kroft *et al.*, 2024). One can then replace contribution shares by expenditures and material use, as in (28). Second, observe that the pre–log relationship has the implicit form $VA_{jt} = \Phi\big(VA_{jt}, M_{jt}, K_{jt}, \boldsymbol{X}^f_{jt}, \boldsymbol{E}^f_{jt}\big)$. Under conditions ensuring a unique fixed point, this yields $VA_{jt} = \Phi'(M_{jt}, K_{jt}, \boldsymbol{X}^f_{jt}, \boldsymbol{E}^f_{jt})$, i.e., the same first–stage after logging. I favor the technical–control route and do not pursue primitive conditions for the fixed–point approach.

The first stage returns predicted value added $\hat{va}_{jt} = \hat{\phi}(\cdot)$ and the measurement error $\hat{\varepsilon}_{jt} = va^O_{jt} -$

[34] Zhang (2019) highlighted this concern, and proves that Φ can be strictly increasing in the noise under mild conditions, and identification follows.

$\hat{\phi}(\cdot)$. Substituting into (25)–(27) gives a mapping from observables to productivity components:

$$\hat{\boldsymbol{\omega}}_{jt} \;=\; \mathbb{W}\big(k_{jt}, \boldsymbol{x}^f_{jt}, \boldsymbol{e}^f_{jt}, \hat{va}_{jt}\big), \tag{29}$$

where $\mathbb{W}(\cdot)$ is parametric in all unknown technology parameters and thus can pin down the value–added function.

6.3 Second-stage

The second stage is a GMM step typical of control approaches, now in multiple productivity dimensions. By Assumption 10,

$$\mathbb{E}\big[\big(\hat{\boldsymbol{\omega}}_{jt} - \mathbb{E}[\hat{\boldsymbol{\omega}}_{jt} \mid \{\hat{\boldsymbol{\omega}}_{j,t'}\}_{t'=t-i}^{t-1}]\big)\, Z_{jt}\big] = 0, \qquad Z_{jt} \in \mathcal{I}_{jt}. \tag{30}$$

To parameterize productivity dynamics, let

$$\boldsymbol{\omega}_{jt} \;=\; \Psi\big(\{\boldsymbol{\omega}_{j,t'}\}_{t'=t-i}^{t-1}\big) + \boldsymbol{\upsilon}_{jt},$$

with Ψ the productivity evolution process, and $\boldsymbol{\upsilon}_{jt}$ the unpredictable productivity shock. In stylized applications, Ψ can be a VAR(1) (or diagonal AR(1) if interactions are deemed weak). The GMM moments then take the form

$$\mathbb{E}\big[\big(\hat{\boldsymbol{\omega}}_{jt} - \Psi(\{\hat{\boldsymbol{\omega}}_{j,t'}\}_{t'=t-i}^{t-1})\big)\, Z_{jt}\big] = 0, \qquad Z_{jt} \in \mathcal{I}_{jt}, \tag{31}$$

identifying the technology and the productivity evolution. Productivity parameters and elasticities of substitution are recovered from the estimated parameters. Thus, one can test CE and HN assumptions after the estimation.

To sharpen identification and make the source of variation transparent, plausibly exogenous relative-price shifters from the CE reduced form can be incorporated as additional instruments in the GMM. Their exogeneity was argued under CE. In the broader setting here, I state a mild benchmark: the shifters $Z_{p,jt}$ are uncorrelated with relative factor-biased productivities between X_1 and X_2. This yields the additional exclusion moment conditions

$$\mathbb{E}[(\hat{\omega}_{1,jt} - \hat{\omega}_{2,jt})\, Z_{p,jt}] = 0. \tag{32}$$

Implementation can be tailored to instrument credibility and strength. For instance, if the price–shifter exclusion is more credible, one may give those moments higher weight or drop the

weakest timing moments to avoid over-identification.[35]

6.4 Identification

Gandhi *et al.* (2020) highlight nonidentification risks for control approaches in gross–output settings, while their first stage implicitly imposes HN. To release HN, I adopt the technical–control strategy of Ackerberg *et al.* (2015) to avoid this problem. Furthermore, note that imperfect output competition can introduce additional variation helpful for identification (Ackerberg and De Loecker, 2024). My method targets input–demand substitution and does not require strong output-demand–side structure. Demand shocks, when present, can assist identification. Moreover, while Gandhi *et al.* (2020) emphasize that identification hinges on unpredictable shocks (so perfectly predictable input prices are problematic), my framework allows time-firm-varying flexible input supply curve variations. The unpredictable components, i.e., orthogonal to the information set $\mathcal{I}_{jt}$, may serve as useful sources of identification. Finally, because I integrate CE and HN ideas, plausibly exogenous instruments commonly used in CE–style reduced forms can be included in the information set to enhance transparency and robustness.

6.5 Relation to the literature

The method nests CE–style reduced forms within a production–function control approach (Olley and Pakes, 1996; Levinsohn and Petrin, 2003; Ackerberg *et al.*, 2015), and is closely related to recent nonparametric work with factor-biased productivities (Demirer, 2022; Ackerberg *et al.*, 2023). I accommodate richer input and multiple factor-biased productivities to study substitutability, trading off against mild structure that yields separability of productivities parameters (arbitrary nested CES or translog). Relative to factor-biased productivity estimation approaches, Krusell *et al.* (2000b); León-Ledesma *et al.* (2010); Doraszelski and Jaumandreu (2018); Zhang (2019), I permit more flexible functional forms. Relative to Grieco *et al.* (2016); Gandhi *et al.* (2020), I extend identification to settings with non–neutral productivities.

7 Empirical application

7.1 Data and specification.

I apply the method to Norwegian administrative micro data, 2010–2022. The firm–year panel merges (i) annual firm financial statements, (ii) monthly employer–employee records, and (iii) an-

[35]Instrument credibility may be tested by sensitivity to the chosen moment set and standard diagnostics (e.g., Hansen's J test, first-stage relevance, and weak-identification checks).

nual plant-level energy reports for manufacturing.[36] The energy survey records quantities (GWh) and expenditures for: electricity; petroleum products; gas; coal; steam and district heating; solid biomass and waste; and transport fuels. Norway's production mix is relatively clean and gas-intensive among fossil fuels. To limit mechanical zeros in "dirty" use, I classify all fossil-energy inputs (including gas) as dirty, and group electricity, steam/district heating, and solid biomass/waste as clean. Results are robust when clarifying transport fuels as intermediate material. I will focus on the food & beverage industry, and expand the analysis to other sectors in the future.[37]

The panel includes observed value added VA^O_{jt}, capital K_{jt}, flexible input uses: labor L_{jt}, clean energy C_{jt}, and dirty energy D_{jt}, and input expenditures: labor $E_{L,jt}$, clean energy $E_{C,jt}$, dirty energy $E_{D,jt}$, intermediate goods M_{jt}, and an "other inputs" aggregate O_{jt}, including advertising, service, outsourcing, office, and transportation renting.

To recover the clean-–dirty elasticity of substitution, I estimate a five-input translog value-added function with factor-biased productivities for labor and energies[38]:

$$VA^O_{jt} = \Omega_{Rjt}\Big[F\big(K_{jt},\ \Omega^L_{jt}L_{jt},\ \Omega^C_{jt}C_{jt},\ \Omega^D_{jt}D_{jt},\ O_{jt}\big)\Big]^{s_R}\exp(\varepsilon_{jt}).$$

I treat O_{jt} separately from M_{jt} because those "other inputs" seem substitutable, and are thus unlikely to enter the Leontief gross-output function required by the technical-control approach (Ackerberg *et al.*, 2015). Furthermore, since the physical quantity for O_{jt} is unavailable, its productivity can not be separately identified. I drop its factor-biased productivity and follow Grieco *et al.* (2016) to avoid using its expenditure as quantity in estimation. I cannot allocate the capital stock to clean- versus dirty-using components. Accordingly, I prefer proxying firms' ability to utilize each energy type through clean- and dirty-augmenting productivity terms.

7.2 My method: clean and dirty are near-Leontief

Panel a of Figure 4 plots the firm–year distribution of the estimated clean-–dirty elasticities from my production–function approach.[39] The mass concentrates near zero with essentially no right tail, indicating that clean and dirty energy behave as *strong complements* (a near–Leontief relationship). While Leontief is often motivated by engineering constraints (e.g., fixed energy mixes within

[36]The firm–plant merge is validated by comparing firm-reported energy expenditures in financial statements with aggregated plant-reported energy costs. The two series closely align on average.

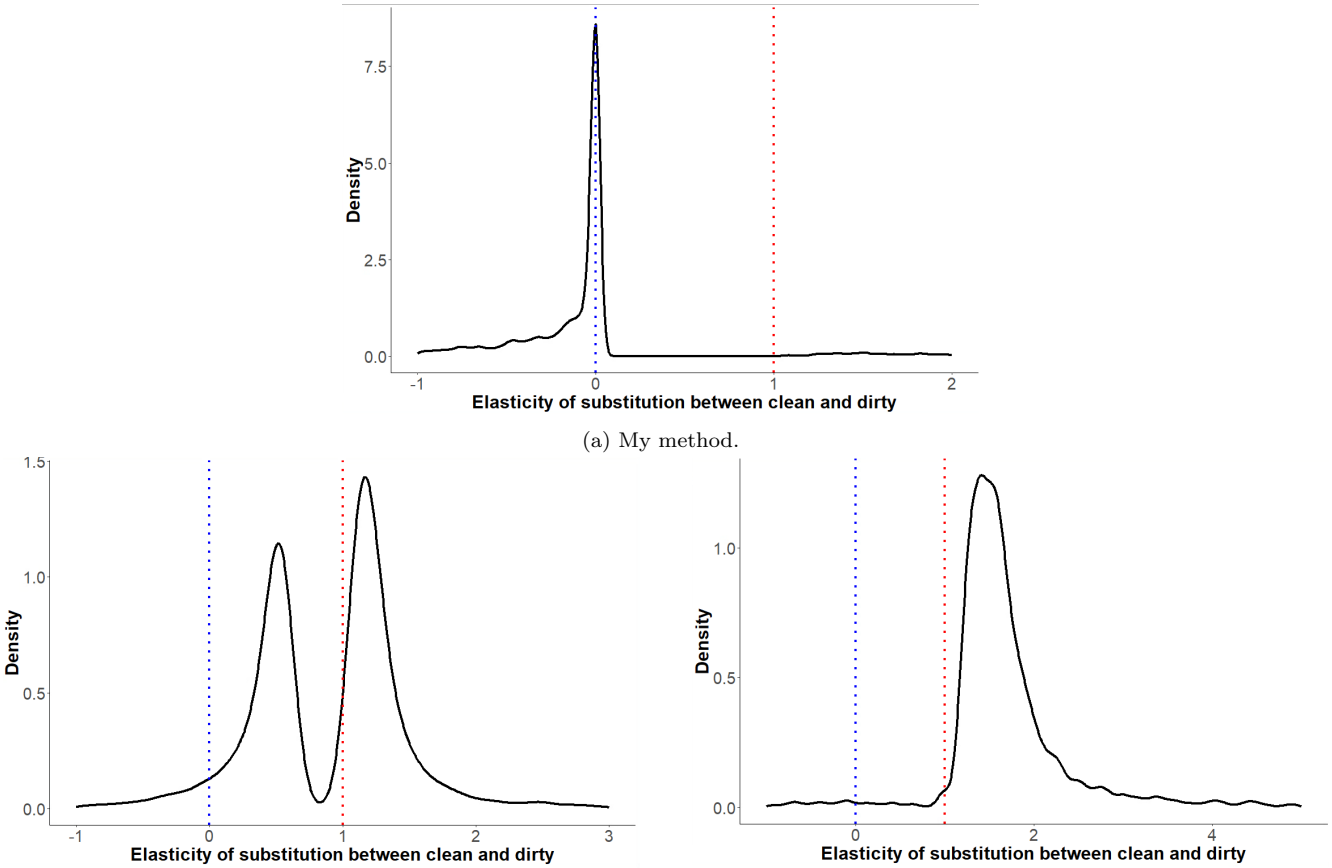
[37]In the food & beverage industry sample, clean energy accounts on average for 60.4% of total energy use in GWh. Within the clean basket, electricity is 86.7%, steam/district heating 11.0%, and solid biomass/waste 2.4%. Within the dirty basket, coal is negligible (0.0%); gas accounts for 63.1%, petroleum products 25.2%, and transport fuels 11.7%.

[38]See Appendix C for the translog estimation details and the two-stage control procedure with multiple productivity components.

[39]See Table A.1 for the translog coefficients of my method and other methods.

processes), these firms are not necessarily single–product. The estimates therefore suggest that, even across product lines within the firm, the scope for fuel substitution is limited.

With near–Leontief substitution, green transition operates primarily through output contraction, cross–firm reallocation, and energy–biased technical change/equipment turnover, rather than within–firm fuel switching induced by relative price incentives. Consequently, carbon pricing primarily reduces emissions via the negative scale channel: lower output and reallocation away from dirty–intensive producers, rather than within–firm fuel switching. Dirty–intensive firms may face higher effective abatement costs and greater transition risk (exit/stranded assets), and market concentration may rise as surviving firms are cleaner on average. Conversely, policies that expand clean supply, e.g., subsidize clean use, may generate a positive scale effect that raises emissions when substitution is negligible. As future work, combining these micro responses with cross-firm



(a) My method.

(b) ACF.

(c) TWFE.

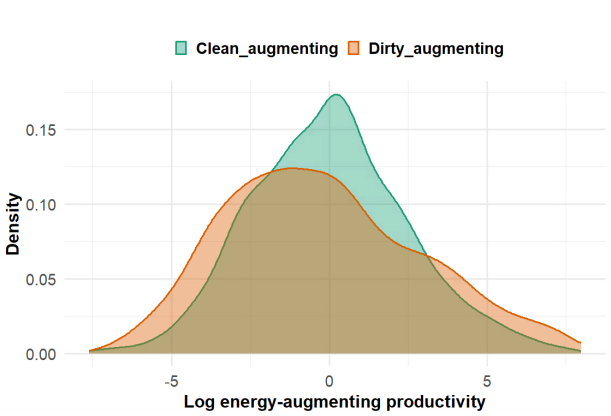
Figure 4: Clean–dirty elasticity of substitution estimates based on my method and two HN-based methods

Notes: Distribution plot of the clean–dirty elasticity of substitution of each observation in food & beverage industry based on the production function estimates of my method (Panel a), ACF method (Panel b), and TWFE (Panel c). "ACF" denotes the control-function approach of Ackerberg *et al.* (2015). "TWFE" denotes two-way fixed effects (firm-time) used to absorb hicks-neutral productivity heterogeneity.

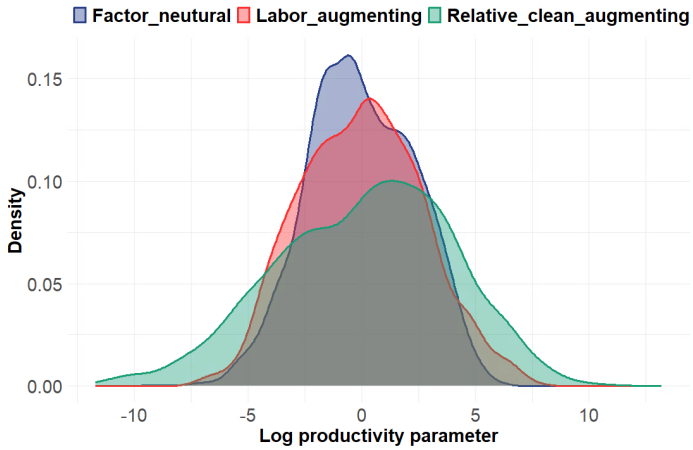
demand substitutes may shed light on industry-level substitutability and aggregated-level policy outcome (Oberfield and Raval, 2021).

The right panel of Figure 5 contrasts the distribution of clean–relative–to–dirty factor–augmenting productivity with the distributions of labor–augmenting and Hicks-neutral (revenue) productivities. Dispersion in the energy bias is markedly larger than in the other two dimensions, indicating substantial heterogeneity in firms' ability to use clean versus dirty inputs. This heterogeneity suggests sizeable distributional consequences of climate policy across firms. The pattern remains, but dispersion falls substantially once we control for firm fixed effects (panel d), indicating that most heterogeneity reflects cross-firm differences rather than within-firm variation.[40]

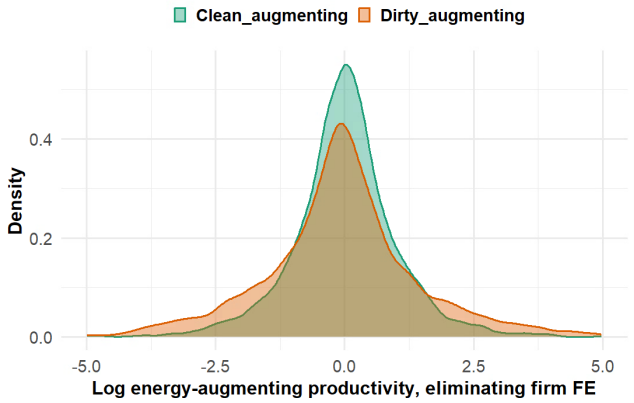
[40]Table A.3 reports pairwise correlations: the clean–relative–to–dirty productivity is weakly negatively correlated



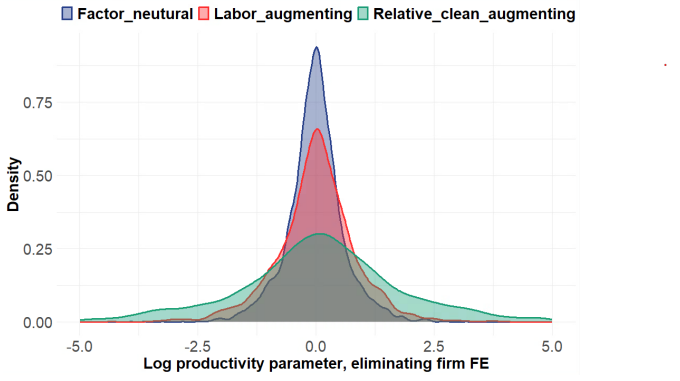
(a) Clean & dirty.



(b) Energy & labor & hicks-neutral.



(c) Clean & dirty, eliminating firm fixed effects.



(d) Energy & labor & hicks-neutral, eliminating firm fixed effects.

Figure 5: Hicks-neutral, and labor-, clean-, and dirty-augmenting productivity estimates based on my method

Notes: Distribution plot of estimated productivity parameters of each observation in food & beverage industry based on the production function estimates of my method. The left panel shows clean- and dirty-augmenting productivities. The right panel presents the productivity of clean relative to dirty and the other two productivities. Firm fixed effects are eliminated in the two panels below.

Figure 6 plots impulse responses from the estimated law of motion for productivity, tracing the passthrough of innovations to clean- and dirty-augmenting productivities over time. Because clean and dirty are strong complements in my estimates, factor-augmenting innovations display the "opposite-bias" property (Acemoglu, 2002): a clean-augmenting innovation is dirty-biased in the sense that it raises the equilibrium intensity of the other input at given relative prices, and vice versa. I test for path dependence/technology specialization: whether a temporary clean- or dirty-biased productivity shock generates a persistent and widening advantage (a central assumption in directed technical change models, e.g., Acemoglu *et al.* (2012, 2023)). The top panels show no evidence for such dynamics: temporary advantages recede rather than snowball. This tempers expectations that temporary subsidies alone can trigger a permanent energy transition. The result may not be surprising, since path dependence typically requires the two inputs to be substitutes. When inputs are strong complements, a temporary lead tends to raise the marginal value of the lagging input and erode divergence. Mechanically, our estimates show that the own-shock

with Hicks-neutral productivity and weakly positively correlated with labor–augmenting productivity. However, these associations vanish once the two measures enter jointly in a regression with fixed effects.

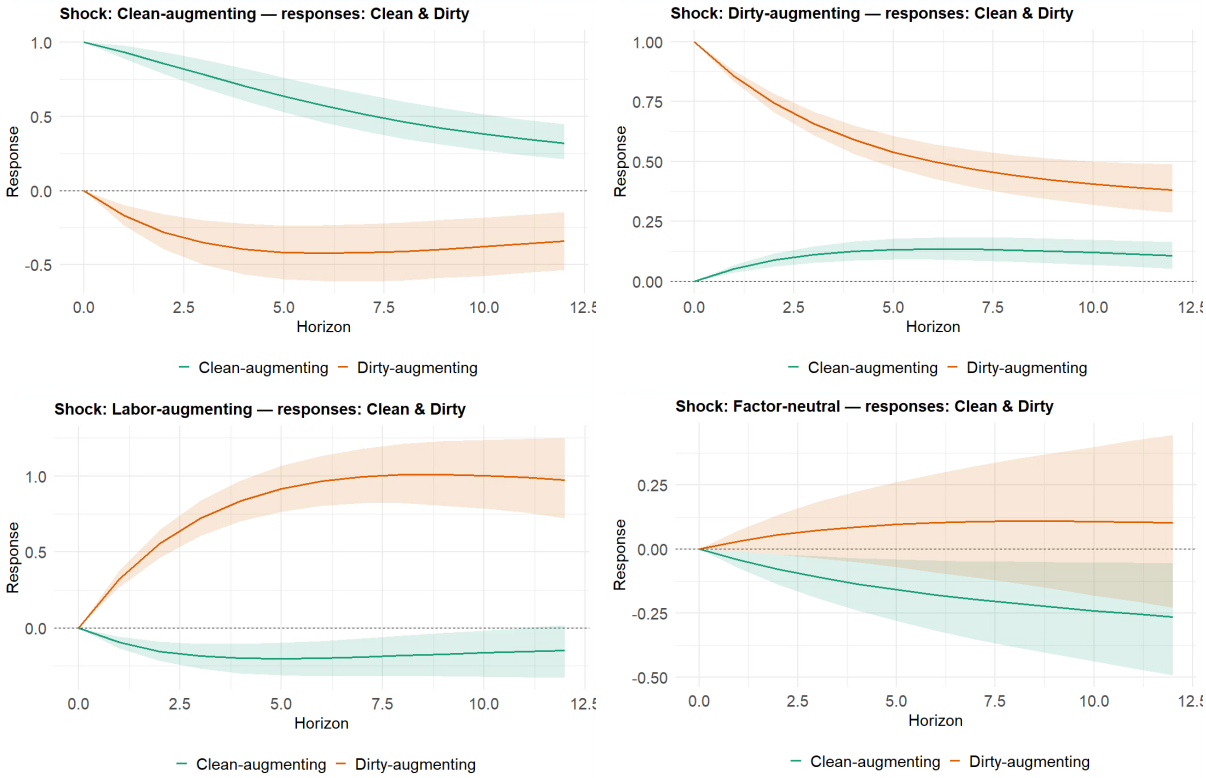


Figure 6: Passthrough of productivity shocks on clean- and dirty-augmenting productivities

Notes: Based on the estimated productivity evolution process, I plot the passthrough of productivity shocks on clean- and dirty-augmenting productivities. 95% bootstrap confidence intervals are included.

Table 2: Summarizing clean–dirty elasticity of substitution estimates of all production function estimation methods based on different identification assumptions.

				HN		
				No	Yes	Yes
				Factor-biased productivity	ACF	TWFE
				(1)	(2)	(3)
CE	No	Translog	(a)	$\sigma_{CD} \approx 0$	Mixed, locally > 1	
	Yes	Nested CES $((C, D), K, L, O)$	(b)	0.045 (0.014)	6.004 (17.070)	2.213 (1.445)
	Yes	CES nested in CD $((C, D), K, L, O)$	(c)	-0.184 (0.010)	4.247 (6.873)	-9.63 (61.001)
	Yes	Single-level CES (C, D, K, L, O)	(d)	0.367 (0.010)	4.383 (2.837)	3.763 (2.464)

Notes: The table reports estimates of the clean–dirty elasticity of substitution in food & beverage industry from all estimators. "ACF" denotes the control-function approach of Ackerberg *et al.* (2015). "TWFE" denotes two-way fixed effects (firm-time) used to absorb hicks-neutral productivity heterogeneity. "Factor-biased productivity" denotes the methods that control input-augmenting productivities. Within functional forms, each pair of parentheses denotes a CES aggregation (nest). In the CES-in-Cobb–Douglas specification, only clean-augmenting productivity is included, as other factor-biased components are not separately identified and are absorbed into the hicks-neutral term. Cell (1a) reports estimates from my proposed method; cells (1b–1d) are CE-based; cells (2a–3a) are HN-based; and cells (2b–3d) jointly impose both CE and HN.

responses of clean and dirty productivity fade at similar speeds. Cross-effects are asymmetric: a dirty-biased (clean-augmenting) shock depresses the subsequent path of clean-biased (dirty-augmenting) productivity, but not strongly enough to overturn the overall reversion. Conversely, a clean-biased (dirty-augmenting) shock lifts dirty-biased (clean-augmenting) productivity. Taken together, dirty-biased technical change yields a relatively more persistent advantage. This asymmetric pattern is consistent with Acemoglu *et al.* (2012): when inputs are strong complements, directed R&D initially concentrates in the less efficient margin (clean-augmenting) until productivities converge and R&D spreads to both, though my evidence is less pronounced than their mechanism. The bottom panels examine the other shocks. Increases in labor-augmenting or hicks-neutral productivity tend to induce clean-biased technical change alongside a dirty-biased recession. The link between energy-biased and hicks-neutral productivity is noisy, while the association with labor-augmenting productivity is highly robust, consistent with evidence on green technology and labor bias (Hedne, 2025).

7.3 CE-based and HN-based methods.

Two diagnostics follow from the estimates. First, given the tight mass near zero in Figure 4, I *cannot reject the CE assumption* in this application. I expect CE–based methods to deliver elasticities close to mine. Second, the left panel of Figure 5 shows substantial dispersion in clean– and dirty–augmenting productivities across firms, leading to a *rejection of HN*. Consequently, approaches that impose HN would tend to misattribute productivity heterogeneity to price–driven substitution and *overstate the elasticity*, sometimes even making clean and dirty gross substituites.

Table 2 compares clean–dirty elasticities across estimators.

Column (1) controls for labor-, clean-, and dirty-augmenting productivities. Cell (1a) repeats my baseline estimates. Cells (1b)–(1d) impose versions of the CE restriction. As expected, (1b) is close to (1a) since CE assumption cannot be rejected. In contrast, (1c) (CES nested in Cobb–Douglas) differs because only clean-augmenting productivity is controlled, as other input biases are absorbed by the hicks-neutral productivity. In (1d) (single-level CES), a single substitution parameter is imposed across all inputs, so the reported coefficient need not coincide with the clean–dirty elasticity of interest. Together, (1c)–(1d) underscore the importance of flexible functional forms for grasping the cross-energy elasticity.

Column (2) implements the control-function estimator of Ackerberg *et al.* (2015), while Column (3) applies two-way fixed effects to absorb hicks-neutral heterogeneity.[41] Panel b and c of Figure 4 and cells (2a)–(3a) report translog-based estimates. Consistent with the bias analysis, omitting energy-augmenting productivity that significantly vary across firms yields *substantially larger elasticities, even local substituites*. Cells (2b)–(3d) impose CE in addition to HN. The resulting single elasticities imply *global substitutes*.

8 Robustness

Transport fuels. On average, transport fuels account for 11.7% of dirty-energy use. Because transport activity is tied to logistics rather than the production energy composite, value added may combine transport services with production in a Leontief manner, biasing the elasticity estimates. Excluding transport fuels from dirty energy, and reclassifying it as intermediate inputs leaves the main result intact: the estimated clean–dirty elasticity moves even closer to zero (see Figure A.7). The result implies that transport fuels are, if anything, more substitutable than the energy used directly in production.

[41] In implementing ACF, I use the same first stage as in my method, controlling for prices. So, ACF's differences from my method arise primarily from the second-stage specification. Comparing ACF with TWFE is consistent with Section 5.6: with a well–specified first stage, the control–function estimator leans on within–firm time variation and is therefore less prone to the upward bias induced by cross-firm dispersion in factor-biased productivities.

9 Conclusion

This paper clarifies what leading empirical strategies for estimating substitution elasticities recover under two common, but potentially misspecified, simplifications: *constant elasticity* (CE) and *Hicks-Neutral* (HN). I then propose an integrated estimator that jointly relaxes both assumptions and recovers firm–time elasticities together with factor-biased productivities. Applying the estimator to Norwegian administrative microdata for 2010–2022, I find a near-Leontief relationship between clean and dirty energy in food and beverage industry. In these data, CE cannot be rejected, whereas HN is rejected. Consequently, CE-based estimates are close to mine, while HN-based estimates substantially overstate substitutability.

The estimates I recover can be inherently local, specific to country, sector, and time horizon. The goal is therefore not a universal "true" elasticity, but a transparent identification toolkit with diagnostics for weighting, selection, and curvature that make assumptions visible and testable. The framework is portable to other context where sustainability and factor-biased productivity are first-order, including skill-biased technical change and wage inequality (Krusell *et al.*, 2000b; Acemoglu, 2002; Beaudry *et al.*, 2010), labor–capital substitution and the labor share (Karabarbounis and Neiman, 2013; De Loecker *et al.*, 2020), and multi-origin input substitution in global supply chains (Antràs and Chor, 2013).

References

Acemoglu, D. (2002). Directed technical change. *The review of economic studies*, **69** (4), 781–809.

—, Aghion, P., Barrage, L. and Hémous, D. (2023). *Climate Change, Directed Innovation, and Energy Transition: The Long-Run Consequences of the Shale Gas Revolution*. Tech. Rep. DP18454, CEPR Discussion Paper, London.

—, —, Bursztyn, L. and Hemous, D. (2012). The environment and directed technical change. *American Economic Review*, **102** (1), 131–66.

Ackerberg, D., Hahn, J. and Pan, Q. (2023). *Nonparametric Identification Using Timing and Information Set Assumptions with an Application to Non-Hicks Neutral Productivity Shocks*. Working papers.

Ackerberg, D. A., Benkard, C. L., Berry, S. and Pakes, A. (2007). Econometric tools for analyzing market outcomes. In J. J. Heckman and E. E. Leamer (eds.), *Handbook of Econometrics*, vol. 6A, Elsevier, pp. 4171–4276.

—, Caves, K. and Frazer, G. (2015). Identification properties of recent production function estimators. *Econometrica*, **83** (6), 2411–2451.

— and De Loecker, J. (2024). *Production Function IdentiÖcation Under Imperfect Competition*. Working papers.

Adão, R., Kolesár, M. and Morales, E. (2019). Shift-share designs: Theory and inference. *Quarterly Journal of Economics*, **134** (4), 1949–2010.

Antràs, P. and Chor, D. (2013). Organizing the global value chain. *Econometrica*, **81** (6), 2127–2204.

Arrow, K. J., Chenery, H. B., Minhas, B. S. and Solow, R. M. (1961). Capital-labor substitution and economic efficiency. *The Review of Economics and Statistics*, **43** (3), 225–250.

Autor, D. H., Levy, F. and Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. *Quarterly Journal of Economics*, **118** (4), 1279–1333.

Bartik, T. J. (1991). *Who Benefits from State and Local Economic Development Policies?* Kalamazoo, MI: W.E. Upjohn Institute for Employment Research.

Beaudry, P., Mark, D. and Ethan, L. (2010). Should the personal computer be considered a technological revolution? evidence from u.s. metropolitan areas. *Journal of Political Economy*, **118** (5), 988–1036.

Berndt, E. R. and Christensen, L. R. (1973). The internal structure of functional relationships: Separability, substitution, and aggregation. *The Review of Economic Studies*, **40** (3), 403–410.

Blackorby, C. and Russell, R. R. (1989). Will the real elasticity of substitution please stand up? (a comparison of the allen/uzawa and morishima elasticities). *The American Economic Review*, **79** (4), 882–888.

Bloom, N., Sadun, R. and Reenen, J. V. (2012). Americans do IT better: US multinationals and the productivity miracle. *American Economic Review*, **102** (1), 167–201.

Borenstein, S. (2012). The private and public economics of renewable electricity generation. *Journal of Economic Perspectives*, **26** (1), 67–92.

Borusyak, K., Hull, P. and Jaravel, X. (2022). Quasi-experimental shift-share research designs. *Review of Economic Studies*, **89** (1).

Bresnahan, T. F., Brynjolfsson, E. and Hitt, L. M. (2002). Information technology, workplace organization, and the demand for skilled labor: Firm-level evidence. *Quarterly Journal of Economics*, **117** (1), 339–376.

Callaway, B. and Sant'Anna, P. H. C. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, **225** (2), 200–230.

Chesher, A. (2003). Identification in nonseparable models. *Econometrica*, **71** (5), 1405–1441.

Collard-Wexler, A. and De Loecker, J. (2025). Productivity and capital measurement error. Manuscript. Earlier version circulated as NBER Working Paper No. 22437 (2016).

de Chaisemartin, C. and D'Haultfœuille, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review*, **110** (9), 2964–2996.

— and D'Haultfœuille, X. (2022). *Difference-in-Differences Estimators of Intertemporal Treatment Effects*. NBER Working Paper 29873, National Bureau of Economic Research.

de la Cal Medina, J. (2025). Weak identification robust methods for production function estimation. Manuscript.

De Loecker, J., Eeckhout, J. and Unger, G. (2020). The Rise of Market Power and the Macroeconomic Implications*. *The Quarterly Journal of Economics*, **135** (2), 561–644.

Demirer, M. (2022). Production function estimation with factor-augmenting technology: An application to markups.

Doraszelski, U. and Jaumandreu, J. (2018). Measuring the bias of technological change. *Journal of Political Economy*, **126** (3), 1027–1084.

Gandhi, A., Navarro, S. and Rivers, D. A. (2020). On the identification of gross output production functions. *Journal of Political Economy*, **128** (8), 2973–3016.

Goldsmith-Pinkham, P., Sorkin, I. and Swift, H. (2020). Bartik instruments: What, when, why, and how. *American Economic Review*, **110** (8), 2586–2624.

Golosov, M., Hassler, J., Krusell, P. and Tsyvinski, A. (2014). Optimal taxes on fossil fuel in general equilibrium. *Econometrica*, **82** (1), 41–88.

Gonzales, L. E., Ito, K. and Reguant, M. (2023). The investment effects of market integration: Evidence from renewable energy expansion in chile. *Econometrica*, **91** (5), 1659–1693.

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, **225** (2), 254–277.

Grieco, P. L. E., Li, S. and Zhang, H. (2016). Production function estimation with unobserved input price dispersion. *International Economic Review*, **57** (2), 665–690.

Griliches, Z. and Mairesse, J. (1995). *Production Functions: The Search for Identification*. NBER Working Paper 5067, National Bureau of Economic Research.

Hall, R. E. (1988). The relation between price and marginal cost in u.s. industry. *Journal of Political Economy*, **96** (5), 921–947.

— (1989). *Invariance Properties of Solow's Productivity Residual*. Nber working papers, National Bureau of Economic Research, Inc.

Hedne, M. H. (2025). Is green technological change skill biased? , working paper; most recent version available online.

Hicks, J. (1932). *Macmillan, London.*, pp. 460–472.

Imbens, G. W. and Newey, W. K. (2009). Identification and estimation of triangular simultaneous equations models without additivity. *Econometrica*, **77** (5), 1481–1512.

Jo, A. (2025). Substitution between clean and dirty energy with biased technical change. *International Economic Review*, **66** (2), 883–902.

Joskow, P. L. (2011). Comparing the costs of intermittent and dispatchable electricity generating technologies. *American Economic Review: Papers & Proceedings*, **101** (3), 238–241.

Kaartinen, M. and Prane, O. (2024). Carbon pricing and fuel switching by firms: Theory and evidence.

Karabarbounis, L. and Neiman, B. (2013). The global decline of the labor share*. *The Quarterly Journal of Economics*, **129** (1), 61–103.

Katz, L. F. and Murphy, K. M. (1992). Changes in relative wages, 1963–1987: Supply and demand factors. *Quarterly Journal of Economics*, **107** (1), 35–78.

Klette, T. J. and Griliches, Z. (1996). The inconsistency of common scale estimators when output prices are unobserved and endogenous. *Journal of Applied Econometrics*, **11** (4), 343–361.

Kroft, K., Luo, Y., Mogstad, M. and Setzler, B. (2024). *Imperfect Competition and Rents in Labor and Product Markets: The Case of the Construction Industry*. Working Paper 27325, National Bureau of Economic Research.

Krusell, P., Ohanian, L. E., Ríos-Rull, J.-V. and Violante, G. L. (2000a). Capital-skill complementarity and inequality: A macroeconomic analysis. *Quarterly Journal of Economics*, **115** (2), 583–622.

—, —, Ríos-Rull, J.-V. and Violante, G. L. (2000b). Capital-skill complementarity and inequality: A macroeconomic analysis. *Econometrica*, **68** (5), 1029–1053.

Leclair, E. M. (2024). Balancing production and carbon emissions with fuel substitution. *Available at SSRN: https://ssrn.com/abstract=4855214 or http://dx.doi.org/10.2139/ssrn.4855214*.

Levinsohn, J. and Petrin, A. (2003). Estimating production functions using inputs to control for unobservables. *The Review of Economic Studies*, **70** (2), 317–341.

León-Ledesma, M. A., McAdam, P. and Willman, A. (2010). Identifying the elasticity of substitution with biased technical change. *American Economic Review*, **100** (4), 1330–57.

Malikov, E., Sun, K. and Kumbhakar, S. C. (2018). Nonparametric estimates of the clean and dirty energy substitutability. *Economics Letters*, **168** (C), 118–122.

Matzkin, R. L. (2003). Nonparametric estimation of nonadditive random functions. *Econometrica*, **71** (5), 1339–1375.

McFadden, D. (1963). Constant elasticity of substitution production functions. *The Review of Economic Studies*, **30** (2), 73–83.

— (1978). Cost, revenue, and profit functions. In M. Fuss and D. McFadden (eds.), *Production Economics: A Dual Approach to Theory and Applications*, vol. 1, *1*, McMaster University Archive for the History of Economic Thought.

Oberfield, E. and Raval, D. (2021). Micro data and macro technology. *Econometrica*, **89** (2), 703–732.

Olley, G. S. and Pakes, A. (1996). The dynamics of productivity in the telecommunications equipment industry. *Econometrica*, **64** (6), 1263–1297.

Papageorgiou, C., Saam, M. and Schulte, P. (2017). Substitution between Clean and Dirty Energy Inputs: A Macroeconomic Perspective. *The Review of Economics and Statistics*, **99** (2), 281–290.

Powell, J. L., Stock, J. H. and Stoker, T. M. (1989). Semiparametric estimation of index coefficients. *Econometrica*, **57** (6), 1403–1430.

Sato, K. (1967). A two-level constant-elasticity-of-substitution production function. *The Review of Economic Studies*, **34** (2), 201–218.

Shephard, R. (1953). *Cost and Production Functions*. Princeton University Press.

Stoker, T. M. (1986). Consistent estimation of scaled coefficients. *Econometrica*, **54** (6), 1461–1481.

Sun, L. and Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, **225** (2), 175–199.

Syverson, C. (2011). What determines productivity? *Journal of Economic Literature*, **49** (2), 326–365.

van Dijcke, D. (2024). *On the non-identification of revenue production functions*. Ssrn working papers.

Zhang, F. (2025). Misplaced confidence: Estimating substitution structures in production functions.

Zhang, H. (2019). Non-neutral technology, firm heterogeneity, and labor demand. *Journal of Development Economics*, **140**, 145–168.

ONLINE APPENDIX

A CE Bias

Bias in Constant-Elasticity (CE) Estimates When True Elasticity Varies

Consider firms that use two inputs X_1 and X_2 (and possibly others) with input prices P_1 and P_2. Without assuming a constant elasticity of substitution (CE), the relative input demand can be expressed as:

$$y_{jt} = ln\frac{X_{1,jt}}{X_{2,jt}} = \int \sigma(p, Z_{jt})\, dp + \varepsilon_{jt},$$

where $\sigma(\cdot)$ denotes the *point elasticity of substitution*, defined as a function of the relative price $p_{jt} = ln\frac{P_{2,jt}}{P_{1,jt}}$ and other firm characteristics Z_{jt} (e.g., the price level $P_{1,jt}$, output Y_{jt}, other inputs $X_{-\{1,2\},jt}$, or factor-biased productivity ω_{jt}). I assume $\sigma(\cdot)$ is non-negative, continuous, and differentiable, i.e., properties implied by a "well-behaved" production function.[42] ε_i is mean-zero noise (omitted for clarity in the derivation below). To simplify exposition, I first consider the case in which σ effectively depends only on relative prices. Time subscript t may be omitted in the following.

A.1 Case 1: $\sigma(p_{jt}, Z_{jt}) = \sigma(p_{jt})$

Suppose σ does not depend on Z_{jt}, or there is no variation in Z_{jt}. I order observed relative prices such that $p_1 \leq p_2 \leq \cdots \leq p_N$.

(1) Adjacent Pairs: Slope as an Affine Combination of Point Elasticities. For any adjacent firms $i, i+1$, the slope of their relative demand in prices equals the average elasticity between p_i and p_{i+1}:

$$\beta_i = \frac{\int_{p_i}^{p_{i+1}} \sigma(p)\, dp}{p_{i+1} - p_i}.$$

Assuming that adjacent individual elasticities $\sigma(p_i), \sigma(p_{i+1})$ are all different. I will return to this assumption later. This slope can be written as a weighted average of the two point elasticities:

$$\beta_i = s_i^i\, \sigma(p_i) + s_i^{i+1}\, \sigma(p_{i+1}), \qquad s_i^i = \frac{\sigma(p_{i+1}) - \beta_i}{\sigma(p_{i+1}) - \sigma(p_i)}, \qquad s_i^i + s_i^{i+1} = 1. \tag{A.1}$$

[42] A well-behaved production function satisfies the properties defined in McFadden (1963). Here, I further require it to be three times continuously differentiable, which holds for most functional forms commonly used in applied work. See Shephard (1953); McFadden (1978) for formal derivations.

If $\sigma(p)$ is monotonic over (p_i, p_{i+1}), then β_i lies between $\sigma(p_i)$ and $\sigma(p_{i+1})$, so the weights form a convex combination. Otherwise, weights may be negative.

The magnitude of s_i^i reflect the local monotonicity and curvature of σ. If σ is increasing and convex (concave), then $s_i^i > 0.5$ (< 0.5); if decreasing, the opposite holds.

(2) The CE-estimate as a Convex Combination of Adjacent Slopes. Since p_i is exogenous by setup, the CE-estimate is the OLS coefficient from regressing factor demand on prices:

$$\hat{\sigma}^{CE} = \frac{\sum_i (p_i - \bar{p})(y_i - \bar{y})}{\sum_i (p_i - \bar{p})^2} = \frac{\sum_{i<j} (p_i - p_j)(y_i - y_j)}{\sum_{i<j} (p_i - p_j)^2} = \frac{\sum_{i<j} (p_i - p_j)^2 \frac{y_i - y_j}{p_i - p_j}}{\sum_{i<j} (p_i - p_j)^2}. \tag{A.2}$$

For any nonadjacent pair $j > i + 1$, their slope can be decomposed into adjacent slopes:

$$\frac{y_i - y_j}{p_i - p_j} = \sum_{k=i}^{j-1} \frac{p_k - p_{k+1}}{p_i - p_j} \cdot \frac{y_k - y_{k+1}}{p_k - p_{k+1}}. \tag{A.3}$$

Thus, the CE-estimate is a convex combination of all adjacent slopes:

$$\hat{\sigma}^{CE} = \sum_{i=1}^{N-1} w_i \, \beta_i, \qquad w_i = \frac{\sum_{k \le i < j} (p_k - p_j)(p_i - p_{i+1})}{\sum_{k<j} (p_k - p_j)^2}. \tag{A.4}$$

(3) The CE-estimate as an Affine Combination of Individual Elasticities. Combining (A.1) and (A.4) gives:

$$\hat{\sigma}^{CE} = \sum_{i=1}^{N} \lambda_i \, \sigma(p_i), \qquad \lambda_i = w_{i-1} s_{i-1}^i + w_i s_i^i, \tag{A.5}$$

where $\lambda_1 = w_1 s_1^1, \lambda_N = w_{N-1} s_{N-1}^N$.

Negative weights The weights w_i are observable from data, but s_i^i and s_i^{i+1} depend on the unknown function $\sigma(p)$. If $\sigma(p)$ is monotonic over the observed range, then $\lambda_i \ge 0$, so $\hat{\sigma}^{CE}$ is a *convex combination* of true point elasticities. If $\sigma(p)$ is nonmonotonic, some weights may be negative, and $\hat{\sigma}^{CE}$ can lie outside the range of true elasticities.

Sensitivity to Outlier Prices. To see how extreme prices distort weights, consider the marginal effect of increasing the largest price p_N:

$$\frac{dw_i}{dp_N} = \frac{iN^2(p_{i+1} - p_i)\left\{\text{Var}(p) - 2\left(\bar{p} - \frac{\sum_{j\leq i} p_j}{i}\right)(p_N - \bar{p})\right\}}{\left[\sum_{k<j}(p_k - p_j)^2\right]^2}. \tag{A.6}$$

As i increases, $\left(\bar{p} - \frac{\sum_{j\leq i} p_j}{i}\right)$ decreases. Hence, raising p_N shifts weights from low-price firms to high-price firms. Because the term in braces is quadratic in p_N with negative coefficient $\frac{1-N}{N^2}$ and axis of symmetry $N\left[\frac{\sum_{j\leq i} p_j}{i} - \frac{N+1}{N}\frac{\sum_{j\leq N} p_j}{N-1}\right]$, large outlier prices assign disproportionately large weight to the last pair of firms.

Finally, within the top pair $(N-1, N)$, the marginal impact on relative weight share depends on the local monotonicity and curvature of $\sigma(p)$:

$$\frac{ds^N_{N-1}}{dp_N} = \frac{s^N_{N-1}}{p_N - p_{N-1}}\left[\frac{1}{s^N_{N-1}} - 1 - \frac{\sigma'(p_N)}{\frac{\int_{p_{N-1}}^{p_N}\sigma'(p)\,dp}{p_N - p_{N-1}}}\right]. \tag{A.7}$$

When $\sigma(p)$ is increasing and convex, or decreasing and concave, over (p_{N-1}, p_N), the derivative is negative, as $\frac{1}{s^N_{N-1}} - 1 < 1 < \frac{\sigma'(p_N)}{\frac{\int_{p_{N-1}}^{p_N}\sigma(p)dp}{p_N - p_{N-1}}}$, implying that a higher p_N *intensifies the inequality* between the two firms, i.e., whichever already has the larger share receives a disproportionately larger weight. The conclusion holds when $\sigma(p)$ is decreasing and convex, or increasing and concave

I now discuss the assumption that adjacent elasticities differ, i.e. $\sigma(p_i) \neq \sigma(p_{i+1})$. When $\sigma(p_i) = \sigma(p_{i+1})$, equation (A.1) no longer provides a unique affine decomposition because $s^i_i + s^{i+1}_i = \frac{\beta_i}{\sigma(p_i)}$. Two potential concerns arise. First, the overall CE-estimate remains an affine combination of point elasticities even in this case. To see this, unless all elasticities are identical, there must exist another observation j such that $\sigma(p_j) \neq \sigma(p_i)$. Using equation (A.3), the slope β_i can then be decomposed into slopes between (i, j) and $(i+1, j)$, each of which is an affine combination of $\sigma(p_i), \sigma(p_{i+1}), \sigma(p_j)$. Substituting these decompositions back into equation (A.4) ensures that the CE estimator remains an affine combination of individual elasticities. Second, when considering the marginal effect of an increase in p_N, $\sigma(p_i) \neq \sigma(p_{i+1})$ affects only the local share allocation between the two adjacent firms $N-1$ and N. The aggregate reweighting pattern across firms, reflected in $\frac{dw_i}{dp_N}$, remains robust.

To summarize, I consider outlier prices in the context without negative weight concern. Outlier prices reallocate weight toward the pair of outlier firms and amplify weight imbalances within that pair. Consequently, CE-estimates can be heavily biased toward the local elasticity of firms facing

extreme prices.

A.2 Case 2: $p \perp Z$

When σ depends on firm-specific characteristics Z, but relative prices p are exogenous, i.e., statistically independent of Z, CE estimator remains an affine combination of individual elasticities. The intuition is straightforward. Conditional on a given $Z = Z$, the estimated elasticity within that subgroup is an affine combination of $\sigma(p \mid Z = z)$, as established earlier. The overall CE-estimate can thus be viewed as a weighted average across these subgroups:

$$\hat{\sigma}^{CE} = \int \hat{\sigma}^{CE}(z)\, dF_Z(z),$$

where $F_Z(z)$ denotes the distribution of Z. Hence, the CE estimator remains an affine combination of all individual elasticities, where weights additionally reflect the empirical prevalence of each characteristic Z.

A.3 Derivation of substitution elasticity for translog function

For convenience, I omit the firm and time subscripts, and productivity parameters.

By definition of elasticity of substitution, except for two inputs X_1 and X_2, output and the other inputs remain unchanged. The total differential of revenue function yields

$$\frac{dlnX_1}{dlnX_2} = -\frac{\frac{\partial lnR}{\partial lnX_2}}{\frac{\partial lnR}{\partial lnX_1}}$$

Accordingly, the direct partial elasticity of substitution can be solved as:

$$\begin{aligned}
\sigma_{X_1X_2} &= \frac{dln\frac{X_1}{X_2}}{dln\frac{MRX_2}{MRX_1}} = \frac{dlnX_1 - dlnX_2}{dln\frac{\partial lnR}{\partial lnX_2} - dln\frac{\partial lnR}{\partial lnX_1} + dlnX_1 - dlnX_2} \\
&= \frac{dlnX_1 - dlnX_2}{\frac{\frac{\partial^2 lnR}{\partial (lnX_2)^2} dlnX_2 + \frac{\partial^2 lnR}{\partial lnX_1 \partial lnX_2} dlnX_1}{\frac{\partial lnR}{\partial lnX_2}} - \frac{\frac{\partial^2 lnR}{\partial (lnX_1)^2} dlnX_1 + \frac{\partial^2 lnR}{\partial lnX_1 \partial lnX_2} dlnX_2}{\frac{\partial lnR}{\partial lnX_1}} + dlnX_1 - dlnX_2} \\
&= \frac{\frac{\partial lnR}{\partial lnX_2} + \frac{\partial lnR}{\partial lnX_1}}{\frac{\partial lnR}{\partial lnX_2} + \frac{\partial lnR}{\partial lnX_1} + 2\frac{\partial^2 lnR}{\partial lnX_1 \partial lnX_2} - \frac{\partial^2 lnR}{\partial (lnX_2)^2} \frac{\frac{\partial lnR}{\partial lnX_1}}{\frac{\partial lnR}{\partial lnX_2}} - \frac{\partial^2 lnR}{\partial (lnX_1)^2} \frac{\frac{\partial lnR}{\partial lnX_2}}{\frac{\partial lnR}{\partial lnX_1}}}
\end{aligned}$$

I have verified that this expression is consistent with the formula derived in McFadden (1963).

Based on translog production function and profit maximization problem:

$$\frac{\partial lnR}{\partial lnX_i} = \frac{E_{X_i}}{R} \qquad \frac{\partial^2 lnR}{\partial lnX_i \partial lnX_j} = s_R \alpha_{X_i X_j}$$

Thus, the clean–dirty elasticity is

$$\sigma_{CD} = \frac{\frac{E_D}{s_R R} + \frac{E_C}{s_R R}}{\frac{E_D}{s_R R} + \frac{E_C}{s_R R} - \frac{E_D}{s_R R}\frac{s_R R}{E_C}\alpha_{CC} - \frac{E_C}{s_R R}\frac{s_R R}{E_D}\alpha_{DD} + 2\alpha_{CD}}$$

Furthermore, I investigate the direct partial elasticity unconditional on output.[43] Since revenue function is homogeneous of degree s_R, total derivative of Euler's theorem implies that

$$\frac{dlnX_1}{dlnX_2} = -\frac{\sum_i \frac{\partial^2 lnR}{\partial lnX_2 \partial lnX_i}}{\sum_i \frac{\partial^2 lnR}{\partial lnX_1 \partial lnX_i}}$$

Accordingly, the unconditional elasticity is

$$\sigma^{UN}_{X_1 X_2} = \frac{\sum_i \frac{\partial^2 lnR}{\partial lnX_1 \partial lnX_i} + \sum_i \frac{\partial^2 lnR}{\partial lnX_2 \partial lnX_i}}{\left[1 + \frac{\frac{\partial^2 lnR}{\partial lnX_1 \partial lnX_2}}{\frac{\partial lnR}{\partial lnX_1}} - \frac{\frac{\partial^2 lnR}{\partial (lnX_2)^2}}{\frac{\partial lnR}{\partial lnX_2}}\right] \sum_i \frac{\partial^2 lnR}{\partial lnX_1 \partial lnX_i} + \left[1 + \frac{\frac{\partial^2 lnR}{\partial lnX_1 \partial lnX_2}}{\frac{\partial lnR}{\partial lnX_2}} - \frac{\frac{\partial^2 lnR}{\partial (lnX_1)^2}}{\frac{\partial lnR}{\partial lnX_1}}\right] \sum_i \frac{\partial^2 lnR}{\partial lnX_2 \partial lnX_i}}$$

[43] Conditional on other inputs if exist.

B HN bias

B.1 Example with endogenous productivity and supplies

Following the intuition based on CES production, I develop a simple system in which both relative factor supplies and relative factor–biased productivity are endogenous, and derive the variance components in (19). Let

$$x_{jt} \equiv \ln \frac{X_{1,jt}}{X_{2,jt}}, \quad p_{jt} \equiv \ln \frac{P_{1,jt}}{P_{2,jt}}, \quad \omega_{jt} \equiv \ln \frac{\Omega_{1,jt}}{\Omega_{2,jt}},$$

CES production implies $x_{jt} = \sigma\,(-p_{jt}) + (\sigma-1)\omega_{jt}$. I close the system with two equilibrium mappings:

(i) Endogenous supply. Relative prices respond to an exogenous supply shifter ρ_{jt}^{exo}, and input demand with supply elasticity $\rho > 0$: $p_{jt} = \rho_{jt}^{\text{exo}} - \rho\, x_{jt}$

(ii) Directed technical change. Relative factor–biased productivity responds to an exogenous productivity shifter θ_{jt}^{exo}, and relative profitability of developing productivity with elasticity $\theta > 0$, as in Acemoglu (2002): $\omega_{jt} = \theta_{jt}^{\text{exo}} + \theta\big(-\frac{\omega_{jt}}{\sigma} + \frac{\sigma-1}{\sigma}x_{jt}\big)$. Suppose $\rho > \theta$

The solution of the system is

$$\Delta \equiv 1 + \sigma\rho + \theta(\rho - \sigma + 2), \qquad \begin{aligned} x_{jt} &= \frac{(\sigma+\theta)\,\rho_{jt}^{\text{exo}} + (\sigma-1)\,\theta_{jt}^{\text{exo}}}{\Delta}, \\ p_{jt} &= \frac{\big[1+\theta(2-\sigma)\big]\rho_{jt}^{\text{exo}} - \rho(\sigma-1)\,\theta_{jt}^{\text{exo}}}{\Delta}, \\ \omega_{jt} &= \frac{\theta(\sigma-1)\,\rho_{jt}^{\text{exo}} + \big(1+\sigma\rho\big)\theta_{jt}^{\text{exo}}}{\Delta}. \end{aligned}$$

Let $\text{Cov}(\rho_{jt}^{\text{exo}}, \theta_{jt}^{\text{exo}}) = 0$, $V_\rho \equiv \text{Var}(\rho_{jt}^{\text{exo}})$, $V_\theta \equiv \text{Var}(\theta_{jt}^{\text{exo}})$. Then

$$\begin{aligned} V_p \equiv \sigma^2\,\text{Var}(p_{jt}) &= \frac{\sigma^2\big[1+\theta(2-\sigma)\big]^2 V_\rho + \rho^2\sigma^4 V_\theta}{\Delta^2}, \\ V_\omega \equiv (\sigma-1)^2\,\text{Var}(\omega_{jt}) &= \frac{\theta^2(\sigma-1)^4 V_\rho + (\sigma-1)^2\big(1+\sigma\rho\big)^2 V_\theta}{\Delta^2}, \\ C_{p\omega} \equiv \sigma(\sigma-1)\text{Cov}(p_{jt},\omega_{jt}) &= \frac{\theta\sigma(\sigma-1)^2\big[1+\theta(2-\sigma)\big]V_\rho + \rho\sigma^2(\sigma-1)\big(1+\sigma\rho\big)V_\theta}{\Delta^2}. \end{aligned}$$

The "HN–shut–down" term $V_\omega + C_{p\omega}$ in (19) is negative only when (i) two inputs are complements but *very close* to unit elasticity and (ii) *strong* productivity-driven movements relative to supply

shocks:

$$\frac{V_\rho}{V_\theta} < \frac{(1+\sigma\rho)[\frac{\sigma^2\rho}{1-\sigma} - 1 - \sigma\rho]}{\sigma(\sigma+\theta)} \qquad \frac{\rho - 1 + \sqrt{\rho^2+6\rho+1}}{4\rho} < \sigma < 1$$

These can be substantially restrictive. For instance, when $\rho = 0.1$, $\theta = 0.05$, the admissible range for σ is $0.922 < \sigma < 1$, and price shocks should be limited, e.g., $V_\rho/V_\theta < 0.171$ as $\sigma = 0.93$, and $V_\rho/V_\theta < 2.264$ as $\sigma = 0.97$. Overall, the example reinforces the main intuitive result that HN tends to **overstate** substitutability.

The simple system differs from Acemoglu (2002) in its input market setup, but it can share common features. From the closed form above,

$$\frac{dp_{jt}}{d\rho_{jt}^{\text{exo}}} = \frac{1+\theta(2-\sigma)}{\Delta}, \qquad \frac{d\omega_{jt}}{d\theta_{jt}^{\text{exo}}} = \frac{1+\sigma\rho}{\Delta}.$$

Hence, a supply shock generally raise p_{jt} less than one-for-one, and can even reduce it if $\sigma > 2 + \frac{1}{\theta}$ (Price shocks are fully offset by directed innovation, i.e., an upper-sloping demand curve). Directed innovation can exhibit amplification $\partial\omega_{jt}/\partial\theta_{jt}^{\text{exo}} > 1$ when $\sigma > 2 + \rho$ (Supply restrictions make directed innovation more profitable). Furthermore, the cross shock derivatives are

$$\frac{dp_{jt}}{d\theta_{jt}^{\text{exo}}} = \frac{-\rho(\sigma-1)}{\Delta}, \qquad \frac{d\omega_{jt}}{d\rho_{jt}^{\text{exo}}} = \frac{\theta(\sigma-1)}{\Delta}.$$

which is guided by the sign of $\sigma - 1$ since it determines whether a X_1-augmenting technical change is X_1-biased (higher output share).

B.2 CES special case

I derive the HN substitutability estimate. By first-order conditions of profit maximization, $p_i = \ln F_{X_i}$, and thus one can map densities from (ω, p) to (ω, x).

$$\mathbb{Q}_x(\omega_1, \omega_2, x_1, x_2) = \mathbb{Q}_p\big(\omega_1, \omega_2, \ln F_{X_1}, \ln F_{X_2}\big)\; s_R^2 \left|\frac{\sigma-1}{\sigma}\right| \tag{B.1}$$

By Bayes' rule, the conditional density of productivities is

$$\mathbb{P}_\omega(\omega_1, \omega_2 \mid x_1, x_2) = \frac{\mathbb{Q}_p\big(\omega_1, \omega_2,\ lnF_{X_1},\ lnF_{X_2}\big)}{\iint_{\mathbb{R}^2} \mathbb{Q}_p\big(w_1, w_2,\ lnF_{X_1},\ lnF_{X_2}, z\big)\, dw_1 dw_2}. \tag{B.2}$$

By Fisher's identity, the score of the conditional density satisfies

$$\partial_{x_i} \ln \mathbb{P}_\omega = \partial_{x_i} \ln \mathbb{Q}_x - \mathbb{E}[\partial_{x_i} \ln \mathbb{Q}_x \mid x_1, x_2]\,, \tag{B.3}$$

$$\partial^2_{x_i} \ln \mathbb{P}_\omega = \partial^2_{x_i} \ln \mathbb{Q}_x - \mathbb{E}\big[\partial^2_{x_i} \ln \mathbb{Q}_x \;\big|\; x_1, x_2\big] - \mathrm{Var}[\partial_{x_i} \ln \mathbb{Q}_x \mid x_1, x_2]\,, \tag{B.4}$$

$$\partial_{x_1}\partial_{x_2} \ln \mathbb{P}_\omega = \partial_{x_1}\partial_{x_2} \ln \mathbb{Q}_x - \mathbb{E}[\partial_{x_1}\partial_{x_2} \ln \mathbb{Q}_x \mid x_1, x_2] - \mathrm{Cov}[\partial_{x_1} \ln \mathbb{Q}_x, \partial_{x_2} \ln \mathbb{Q}_x \mid x_1, x_2]\,. \tag{B.5}$$

Consequently, derivatives of g can be written as:

$$g_{x_1} = \mathbb{E}[f_{x_1} + f \cdot \partial_{x_1} \ln \mathbb{P}_\omega \mid x_1, x_2]\,, \tag{B.6}$$

$$g_{x_i x_i} = \mathbb{E}\Big[f_{x_i x_i} + 2 f_{x_i} \partial_{x_i} \ln \mathbb{P}_\omega + f \partial^2_{x_i} \ln \mathbb{P}_\omega + f \big(\partial_{x_i} \ln \mathbb{P}_\omega\big)^2 \;\Big|\; x_1, x_2\Big]\,, \tag{B.7}$$

$$g_{x_1 x_2} = \mathbb{E}[f_{x_1 x_2} + f_{x_1} \partial_{x_2} \ln \mathbb{P}_\omega + f_{x_2} \partial_{x_1} \ln \mathbb{P}_\omega + f\, \partial_{x_1}\partial_{x_2} \ln \mathbb{P}_\omega + f\, \partial_{x_1} \ln \mathbb{P}_\omega\, \partial_{x_2} \ln \mathbb{P}_\omega \mid x_1, x_2]\,. \tag{B.8}$$

Plugging (B.6)–(B.8) into (20) gives the HN elasticity.

I derive these terms more precisely:

$$f_{x_i} = S_i, \qquad f_{x_i x_i} = \frac{\rho}{s_R} S_1 S_2, \qquad f_{x_1 x_2} = -\frac{\rho}{s_R} S_1 S_2, \qquad \lambda \equiv \frac{\sigma - 1}{\sigma} < 1.$$

The sensitivities of log prices to input uses are

$$c_{ik} \;\equiv\; \partial_{x_i} p_k = \begin{cases} S_i + \dfrac{f_{x_i x_i}}{S_i} - 1 = S_i + \dfrac{\rho}{s_R} S_{-i} - 1, & k = i, \\ S_i + \dfrac{f_{x_i x_{-i}}}{S_{-i}} = S_i \left(1 - \dfrac{\rho}{s_R}\right), & k \neq i, \end{cases}$$

$$u_i \;\equiv\; \partial^2_{x_i} p_k = \frac{\rho}{s_R}\left(1 - \frac{\rho}{s_R}\right) S_i S_{-i} \quad (k = 1, 2), \qquad h \;\equiv\; \partial_{x_1}\partial_{x_2} p_k = -\frac{\rho}{s_R}\left(1 - \frac{\rho}{s_R}\right) S_1 S_2 \quad (k = 1, 2).$$

To streamline the exposition, define

$$g_k(\omega; x) \;\equiv\; \partial_{p_k} \ln \mathbb{Q}_p(\omega_1, \omega_2, p_1, p_2), \qquad H_{kl}(\omega; x) \;\equiv\; \partial_{p_k}\partial_{p_l} \ln \mathbb{Q}_p,$$

$$g \equiv (g_1, g_2)^\top, \quad H \equiv \begin{pmatrix} H_{11} & H_{12} \\ H_{21} & H_{22} \end{pmatrix}, \quad c_i \equiv (c_{i1}, c_{i2})^\top.$$

Thus, I can write:

$$\partial_{x_i} \ln \mathbb{Q}_x = c_i^\top g, \qquad \partial^2_{x_i} \ln \mathbb{Q}_x = c_i^\top H c_i + u_i\,(g_1 + g_2), \qquad \partial_{x_1}\partial_{x_2} \ln \mathbb{Q}_x = c_1^\top H c_2 + h\,(g_1 + g_2).$$

Define the projected price-scores and second-order blocks:

$$A := c_1^\top g, \quad B := c_2^\top g, \qquad Q_1 := c_1^\top H c_1, \quad Q_2 := c_2^\top H c_2, \quad Q_{12} := c_1^\top H c_2 + h\,(g_1 + g_2).$$

Solve all the derivatives:

$$\begin{aligned}
g_{x_1} &= \mathbb{E}[S_1] \;+\; \mathrm{Cov}(f,\; A),\\
g_{x_2} &= \mathbb{E}[S_2] \;+\; \mathrm{Cov}(f,\; B),\\
g_{x_1x_1} &= \mathbb{E}\Big[\tfrac{\rho}{s_R} S_1 S_2\Big] \;+\; 2\;\mathrm{Cov}(S_1,\; A) \;+\; \mathrm{Cov}\big(f,\; Q_1\big) \;+\; \mathrm{Cov}\big(f,\; (A - \mathbb{E}A)^2\big),\\
g_{x_2x_2} &= \mathbb{E}\Big[\tfrac{\rho}{s_R} S_1 S_2\Big] \;+\; 2\;\mathrm{Cov}(S_2,\; B) \;+\; \mathrm{Cov}\big(f,\; Q_2\big) \;+\; \mathrm{Cov}\big(f,\; (B - \mathbb{E}B)^2\big),\\
g_{x_1x_2} &= -\,\mathbb{E}\Big[\tfrac{\rho}{s_R} S_1 S_2\Big] \;+\; \mathrm{Cov}(S_1,\; B) \;+\; \mathrm{Cov}(S_2,\; A) \;+\; \mathrm{Cov}\big(f,\; Q_{12}\big) \;+\; \mathrm{Cov}\big(f,\; (A - \mathbb{E}A)(B - \mathbb{E}B)\big).
\end{aligned}$$

Plugging them into the expression of HN substitutability estimate yields equation 21, where

$$R_{11} := \mathrm{Cov}(f,\; Q_1) + \mathrm{Cov}\big(f,\; (A - \mathbb{E}A)^2\big), \quad R_{22} := \mathrm{Cov}(f,\; Q_2) + \mathrm{Cov}\big(f,\; (B - \mathbb{E}B)^2\big)$$

$$R_{12} := \mathrm{Cov}(f,\; Q_{12}) + \mathrm{Cov}\big(f,\; (A - \mathbb{E}A)(B - \mathbb{E}B)\big).$$

B.3 General case

Proof in progress

Consider well-behaviored production function $y = f(\omega_1 + x_1, \omega_2 + x_2, z)$ where z is a vector of fixed inputs, and (f, y, ω, x, z, p) are log transformed (F, Y, Ω, X, Z, P).[44] Firms maximize profit in x_1, x_2, given $p_1, p_2, \omega_1, \omega_2, z$. Hicks-neutral productivity is omitted, simulating the case that it is perfectly addressed in estimation, and the means of ωs are normalized to zero. Firm behavior x_1, x_2 are characterized by the following FOCs:

$$p_i = ln\frac{dF}{dX_i} = ln\frac{dlnF}{dlnX_i}\frac{F}{X_i} = lnf_{x_i} + f(\omega_1 + x_1, \omega_2 + x_2, z) - x_i \qquad \forall i \in \{1, 2\}$$

Suppose that $\omega_1, \omega_2, p_1, p_2, z$ have known continuous and differentiable joint density

$$\mathbb{Q}_p(\omega_1, \omega_2, p_1, p_2, z).$$

Assume the mapping $(\omega_i, x_i, z) \mapsto (\omega_i, p_i = lnF_{X_i}, z)$ is a local diffeomorphism on the domain of interest, so the Jacobian determinant below does not vanish.

[44] Precisely, f is the function of lnY in log inputs.

Again, I omit the time subscript t. Under HN, suppose production function $f(\omega_1+x_1, \omega_2+x_2, z)$ is estimated as $g(x_1, x_2, z)$ by

$$\min_{g(.)} \mathbb{E}\left[(Y_i - g(x_{1,i}, x_{2,i}, z_i))^2\right]$$

The HN producton function estimate is

$$g(x_1, x_2, z) = \mathbb{E}\left[f(\omega_{1,i}x_{1,i}, \omega_{2,i}x_{2,i}, z_i)|(x_i, z_i) = (x, z)\right] \tag{B.9}$$

I intend to study the bias of elasticity of substitution between X_1 and X_2 under HN.

1. Conditional density of productivities In joint distribution $\mathbb{Q}$, prices can be replaced by inputs and productivities based on the FOCs. Thus, the distribution of inputs and productivities is the Jacobian formula:

$$\mathbb{Q}_x(\omega_1, \omega_2, x_1, x_2, z) = \mathbb{Q}_p\big(\omega_1, \omega_2,\ lnF_{X_1},\ lnF_{X_2}, z\big)\ \big|\Delta(\omega, x, z)\big| \tag{B.10}$$

where

$$\begin{aligned}
\Delta(\omega, x, z) &= \frac{\partial lnF_{X_1}}{\partial x_1}\frac{\partial lnF_{X_2}}{\partial x_2} - \frac{\partial lnF_{X_1}}{\partial x_2}\frac{\partial lnF_{X_2}}{\partial x_1} \\
&= \frac{(f_{x_1x_1} + f_{x_1}^2)(f_{x_2x_2} + f_{x_2}^2) - (f_{x_1x_2} + f_{x_1}f_{x_2})^2}{f_{x_1}f_{x_2}} \\
&= \frac{\det(f_{x_ix_j} + f_{x_i}f_{x_j})}{f_{x_1}f_{x_2}} = \frac{\det(\mathbb{H}(H))}{f_{x_1}f_{x_2}F^2}
\end{aligned} \tag{B.11}$$

where $\mathbb{H}$ is the Hessian matrix, and H is the output function in log inputs: $H(\omega_1+x_1, \omega_2+x_2, z) \equiv F(\Omega_1X_1, \Omega_2X_2, Z)$. Thus, $\Delta(\omega, x, z)$ measures the local curvature of the production surface in log-input space, relative to product of output elasticities of X_1, X_2. It is, to some extent, linked to the elasticity of substituition. When $\Delta(\omega, x, z) = 0$, elasticity is a constant, i.e., X_1, X_2 form a CES formula within $F(.)$. Otherwise, its sign helps determine whether small perturbations in input ratios change marginal ratw of substituition in a way consistent with local concavity/convexity or produce saddle behavior.

The marginal density of inputs is

$$\mathbb{P}_X(x_1, x_2, z) = \iint_{\mathbb{R}^2} \mathbb{Q}_x(w_1, w_2, x_1, x_2, z)\, dw_1 dw_2$$

By Bayes' rule, I compute the conditional density of productivities:

$$\mathbb{P}_\omega(\omega_1,\omega_2 \mid x_1,x_2,z)=\frac{\mathbb{Q}_x(\omega_1,\omega_2,x_1,x_2,z)}{\mathbb{P}_X(x_1,x_2,z)}=\frac{\mathbb{Q}_p\big(\omega_1,\omega_2,\ lnF_{X_1},\ lnF_{X_2},z\big)\ \big|\Delta(\omega,x,z)\big|}{\iint_{\mathbb{R}^2}\mathbb{Q}_p\big(w_1,w_2,\ lnF_{X_1},\ lnF_{X_2},z\big)\ \big|\Delta(w,x,z)\big|\,dw_1dw_2}. \tag{B.12}$$

I now compute the derivative of conditional density, staring by solving the derivative of $ln\mathbb{Q}_x$:

$$\begin{aligned}\partial_{x_1}ln\mathbb{Q}_x&=\partial_{x_1}ln\mathbb{Q}_p\big(\omega_1,\omega_2,\ lnF_{X_1},\ lnF_{X_2},z\big)+\partial_{x_1}ln\big|\Delta(w,x,z)\big|\\&=\sum_{i=1}^{2}\frac{\partial ln\mathbb{Q}_p}{\partial lnF_{X_i}}\cdot\partial_{x_1}lnF_{X_i}+\partial_{x_1}ln\big|\Delta(w,x,z)\big|\end{aligned} \tag{B.13}$$

$$\partial^2_{x_1}ln\mathbb{Q}_x=\sum_{i=1}^{2}\sum_{j=1}^{2}\frac{\partial^2 ln\mathbb{Q}_p}{\partial lnF_{X_i}\partial lnF_{X_j}}\cdot\partial_{x_1}lnF_{X_i}\cdot\partial_{x_1}lnF_{X_j}+\frac{\partial ln\mathbb{Q}_p}{\partial lnF_{X_i}}\cdot\partial^2_{x_1}lnF_{X_i}+\partial^2_{x_1}ln\big|\Delta(w,x,z)\big| \tag{B.14}$$

$$\partial_{x_2}\partial_{x_1}ln\mathbb{Q}_x=\sum_{i=1}^{2}\sum_{j=1}^{2}\frac{\partial^2 ln\mathbb{Q}_p}{\partial lnF_{X_i}\partial lnF_{X_j}}\cdot\partial_{x_1}lnF_{X_i}\cdot\partial_{x_2}lnF_{X_j}+\frac{\partial ln\mathbb{Q}_p}{\partial lnF_{X_i}}\cdot\partial_{x_1}\partial_{x_2}lnF_{X_i}+\partial_{x_1}\partial_{x_2}ln\big|\Delta(w,x,z)\big| \tag{B.15}$$

Thus, by Fisher's identity, the derivative of conditional density is

$$\partial_{x_1}ln\mathbb{P}_\omega=\partial_{x_1}ln\mathbb{Q}_x-\mathbb{E}\left[\partial_{x_1}ln\mathbb{Q}_x\mid x_1,x_2,z\right] \tag{B.16}$$

$$\partial^2_{x_1}ln\mathbb{P}_\omega=\partial^2_{x_1}ln\mathbb{Q}_x-\mathbb{E}\left[\partial^2_{x_1}ln\mathbb{Q}_x\mid x_1,x_2,z\right]-\mathrm{Var}\left[\partial_{x_1}ln\mathbb{Q}_x\mid x_1,x_2,z\right] \tag{B.17}$$

$$\partial_{x_2}\partial_{x_1}ln\mathbb{P}_\omega=\partial_{x_2}\partial_{x_1}ln\mathbb{Q}_x-\mathbb{E}\left[\partial_{x_2}\partial_{x_1}ln\mathbb{Q}_x\mid x_1,x_2,z\right]-\mathrm{Cov}\left[\partial_{x_1}ln\mathbb{Q}_x,\partial_{x_2}ln\mathbb{Q}_x\mid x_1,x_2,z\right] \tag{B.18}$$

2. Derivatives of HN production function As a preparetion for solving elasticities, I compute the derivitives of $g(.)$

$$g_{x_1}=\mathbb{E}\left[f_{x_1}\mid x_1,x_2,z\right]+\mathbb{E}\left[f\cdot\partial_{x_1}ln\mathbb{P}_\omega\mid x_1,x_2,z\right] \tag{B.19}$$

$$\begin{aligned}g_{x_1x_1}&=\mathbb{E}\left[f_{x_1x_1}\mid x_1,x_2,z\right]+2\mathbb{E}\left[f_{x_1}\cdot\partial_{x_1}ln\mathbb{P}_\omega\mid x_1,x_2,z\right]\\&\quad+\mathbb{E}\left[f\cdot\partial^2_{x_1}ln\mathbb{P}_\omega\mid x_1,x_2,z\right]+\mathbb{E}\left[f\cdot\left(\partial_{x_1}ln\mathbb{P}_\omega\right)^2\mid x_1,x_2,z\right]\end{aligned} \tag{B.20}$$

$$\begin{aligned}g_{x_1x_2}&=\mathbb{E}\left[f_{x_1x_2}\mid x_1,x_2,z\right]+\mathbb{E}\left[f_{x_1}\cdot\partial_{x_2}ln\mathbb{P}_\omega\mid x_1,x_2,z\right]\\&\quad+\mathbb{E}\left[f_{x_2}\cdot\partial_{x_1}ln\mathbb{P}_\omega\mid x_1,x_2,z\right]+\mathbb{E}\left[f\cdot\partial_{x_1}\partial_{x_2}ln\mathbb{P}_\omega\mid x_1,x_2,z\right]+\mathbb{E}\left[f\cdot\partial_{x_1}ln\mathbb{P}_\omega\cdot\partial_{x_2}ln\mathbb{P}_\omega\mid x_1,x_2,z\right]\end{aligned} \tag{B.21}$$

3. Productivity effect vs. price effect Conditional on input uses, taking total derivative on FOCs yields:

$$dp_i = \left(\frac{f_{x_i x_1}}{f_{x_i}} + f_{x_1}\right) d\omega_1 + \left(\frac{f_{x_i x_2}}{f_{x_i}} + f_{x_2}\right) d\omega_2$$

Thus, the conditional marginal impacts of productivities on prices are

$$\frac{\partial p_i}{\partial \omega_1} + \frac{\partial p_i}{\partial \omega_2} = \frac{f_{x_i x_1} + f_{x_i x_2}}{f_{x_i}} + f_{x_1} + f_{x_2} \tag{B.22}$$

Conditional on output and input uses, taking total derivative on the production function yields:

$$0 = f_{x_1} d\omega_1 + f_{x_2} d\omega_2 \tag{B.23}$$

This proof is restricted to the context without a strong negative correlation between the price effect and the productivity effect. In the following, I focus on the context of calculating the elasticity of substitution, i.e., fixing output level and other inputs. By definition of the elasticity, I have:

$$\begin{aligned} & \frac{d(\omega_1 + x_1 - \omega_2 - x_2)}{d(p_2 - \omega_2 - p_1 + \omega_1)} \mid y, z = \sigma \\ \iff\ & d(x_1 - x_2) = \underbrace{-\sigma d(p_1 - p_2)}_{\text{Price effect}} + \underbrace{(\sigma - 1) d(\omega_1 - \omega_2)}_{\text{Productivity effect}} \\ \text{Where}\quad & \sigma = \frac{f_{x_1} + f_{x_2}}{f_{x_1} + f_{x_2} + 2f_{x_1 x_2} - f_{x_1 x_1}\frac{f_{x_2}}{f_{x_1}} - f_{x_2 x_2}\frac{f_{x_1}}{f_{x_2}}} \end{aligned}$$

4. HN bias The true point elasticity and the HN-estimate are:

$$\hat{\sigma}^{HN} = \frac{g_{x_1} + g_{x_2}}{g_{x_1} + g_{x_2} + 2g_{x_1 x_2} - g_{x_1 x_1}\frac{g_{x_2}}{g_{x_1}} - g_{x_2 x_2}\frac{g_{x_1}}{g_{x_2}}} \tag{B.24}$$

I want to prove that without a strong negative correlation between the price effect and the productivity effect, HN over-estimates substitutability between X_1, X_2 ($\frac{\sigma-1}{\sigma}$). That is, there exist (x_1, x_2, z) such that locally

$$\frac{g_{x_1 x_1}\frac{g_{x_2}}{g_{x_1}} + g_{x_2 x_2}\frac{g_{x_1}}{g_{x_2}} - 2g_{x_1 x_2}}{g_{x_1} + g_{x_2}} > \frac{f_{x_1 x_1}\frac{f_{x_2}}{f_{x_1}} + f_{x_2 x_2}\frac{f_{x_1}}{f_{x_2}} - 2f_{x_1 x_2}}{f_{x_1} + f_{x_2}} \quad \exists(\omega_1, \omega_2) \mid \mathbb{P}_\omega > 0 \tag{B.25}$$

I start by discussing the bias in $g_{x_1} + g_{x_2}$ beyond the average local scale $\mathbb{E}\left[f_{x_1} + f_{x_2} \mid x_1, x_2, z\right]$, i.e.,

$$\begin{aligned} g_{x_1} + g_{x_2} - \mathbb{E}\left[f_{x_1} + f_{x_2} \mid x_1, x_2, z\right] &= \mathbb{E}\left[f_{x_1} + f_{x_2} + f \cdot (\partial_{x_1} ln\mathbb{P}_\omega + \partial_{x_2} ln\mathbb{P}_\omega) \mid x_1, x_2, z\right] \\ &= \mathrm{Cov}\left[f, \partial_{x_1} ln\mathbb{Q}_x + \partial_{x_2} ln\mathbb{Q}_x \mid x_1, x_2, z\right] \end{aligned} \tag{B.26}$$

Note that conditional on x_1, x_2, z, f is positively correlated with productivities (and thus prices). Since X_i effect $ln\mathbb{Q}_x$ only through changing $\Omega_i X_i$,

$$\begin{aligned} &\partial_{x_1} ln\mathbb{Q}_x + \partial_{x_2} ln\mathbb{Q}_x = \sum_{i=1}^{2} \frac{\partial ln\mathbb{Q}_p}{\partial lnF_{X_i}} \cdot (\partial_{x_1} lnF_{X_i} + \partial_{x_2} lnF_{X_i}) + \partial_{x_1} ln\big|\Delta(w,x,z)\big| + \partial_{x_2} ln\big|\Delta(w,x,z)\big| \\ &= \sum_{i=1}^{2} \frac{\partial ln\mathbb{Q}_p}{\partial p_i} \cdot (\partial_{\omega_1} p_i + \partial_{\omega_2} p_i) + \frac{\partial_{x_1}\Delta(w,x,z) + \partial_{x_2}\Delta(w,x,z)}{\Delta(w,x,z)} \\ &\underset{(B.11)}{=} \underbrace{\sum_{i=1}^{2} \frac{\partial ln\mathbb{Q}_p}{\partial p_i} \cdot \left(\frac{f_{x_i x_1} + f_{x_i x_2}}{f_{x_i}} + f_{x_1} + f_{x_2}\right)}_{\text{Scale-effect selection bias}} + \underbrace{\frac{\partial_{x_1}\Delta(w,x,z) + \partial_{x_2}\Delta(w,x,z)}{\Delta(w,x,z)}}_{\text{Jacobian part}} \end{aligned}$$

where the Jacobian part is a constant if F is a CES.

The scale-effect selection term $\mathrm{Cov}\left(f, \sum_{i=1}^{2} \frac{\partial ln\mathbb{Q}_p}{\partial p_i} \cdot \left[\frac{f_{x_i x_1} + f_{x_i x_2}}{f_{x_i}} + f_{x_1} + f_{x_2}\right]\right)$ is the crucial source of the bias. Under, e.g., a log-concave price law, pirce score $\sum_{i=1}^{2} \frac{\partial ln\mathbb{Q}_p}{\partial p_i}$ is decreasing in the overall price level. Conditional on observed input uses, higher factor-based productivity raises output f and raises marginal products, making f and $\sum_{i=1}^{2} \frac{\partial ln\mathbb{Q}_p}{\partial p_i}$ negatively correlated. However, in addition, an overall growth in production technology has diminishing own marginal products and moderate cross effects, then $\frac{f_{x_i x_1} + f_{x_i x_2}}{f_{x_i}} + f_{x_1} + f_{x_2}$ can be mildly negative. Thus, $\mathrm{Cov}\left(f, \sum_{i=1}^{2} \frac{\partial ln\mathbb{Q}_p}{\partial p_i} \cdot \left[\frac{f_{x_i x_1} + f_{x_i x_2}}{f_{x_i}} + f_{x_1} + f_{x_2}\right]\right)$ overall tends to give a nonnegative contribution to the covariance term. Since the output's covariance with the Jacobian part is typically bounded by a small value, the scale elasticity of two inputs $g_{x_1} + g_{x_2}$ is expected to be over-estimated relative to the local mean scale.

C Translog estimation procedure

I estimate the value–added technology under the assumptions in Section 6, allowing for factor–biased productivities on flexible inputs and a translog structure (heterogeneous elasticity of substitution).

Under a translog value–added aggregator with constant returns to scale restrictions imposed,[45] the FOC for any flexible input implies the log-linear share equation:

$$\frac{E_{X_i,jt}}{s_R\, VA_{jt}} \;=\; \alpha_{X_i} \;+\; \sum_{X_\ell \neq K} \alpha_{X_i X_\ell}\, \ln\!\left(\frac{\omega_{X_\ell,jt}\, X_{\ell,jt}}{K_{jt}}\right) \tag{C.1}$$

where α_{X_i} and $\alpha_{X_i X_\ell}$ are translog coefficients. I observe $K_{jt}, X_{i,jt}$, expenditures , $E_{X_i,jt}$ and an "observed" value added VA^O_{jt} that includes an idiosyncratic shock.

Stacking (C.1) and solving for the factor-biased productivity parameters yields:

$$\ln \omega_{X_i,jt} \;=\; \ln\!\left(\frac{K_{jt}}{X_{i,jt}}\right) \;+\; \sum_{X_\ell \in \boldsymbol{X_f}} \beta_{X_i X_\ell}\left(\frac{E_{X_\ell,jt}}{VA_{jt}} - s_R\, \alpha_{X_\ell}\right) \tag{C.2}$$

where the symmetric matrix $\beta_{X_i X_\ell} \equiv \beta_{X_\ell X_i}$ is a function of the translog quadratic coefficients scaled by $1/s_R$.

Substituting (C.2) into the translog value–added index implies a simple quadratic representation for Hicks-neutral productivity

$$\ln \omega_{R,jt} \;=\; \ln R_{jt} - s_R \ln K_{jt} \;+\; \frac{1}{2} \sum_{X_i \in \boldsymbol{X_f}} \sum_{X_\ell \in \boldsymbol{X_f}} \beta_{X_i X_\ell}\left[s_R^2\, \alpha_{X_i}\alpha_{X_\ell} \;-\; \frac{E_{X_i,jt}}{VA_{jt}}\frac{E_{X_l,jt}}{VA_{jt}}\right]. \tag{C.3}$$

First stage (control). Equation (C.2) shows that each productivity parameter is a function of observable shares $\frac{E_{X_l,jt}}{VA_{jt}}$, relative intensities $K_{jt}/X_{i,jt}$, and unknown parameters. Following the proxy/control tradition (Olley and Pakes, 1996; Levinsohn and Petrin, 2003; Ackerberg *et al.*, 2015), I absorb productivity into a nonparametric function of observables and estimate

$$\ln VA^O_{jt} \;=\; \Phi\Big(K_{jt},\, \boldsymbol{X}_{f,jt},\, \boldsymbol{P}_{f,jt},\; \frac{\boldsymbol{E_{X_f,jt}}}{VA_{jt}},\, m_{jt}\Big) \;+\; \varepsilon_{jt}, \tag{C.4}$$

I estimate $\Phi(\cdot)$ as 2-digit NACE industry code-specific second-level polynomial interactions. $\Phi(\cdot)$ recovers the expected value added conditional on observables and ε_{jt} is a measurement error. The estimated $\widehat{\Phi}$ allows us to form empirical counterparts of (C.2)–(C.3) as functions of data and

[45] CRS restrictions eliminate the coefficients on K in the share equations and impose the usual adding–up constraints on first– and second–order terms.

parameters $(s_R, \beta_{X_i X_\ell})$.[46]

Second stage (GMM). Impose Assumption 10 on the (vector) productivity process, so that the unpredictable productivity shocks are orthogonal to the information set $\mathcal{I}_{jt}$ (second-level polynomial interaction of current capital and lagged choices). I assume the productivity evolution process as VAR(1). Using the model–implied ω's from (C.2) and (C.3), estimate (s_R, α, β) by GMM from moments of the form

$$\mathbb{E}\big[\big(\omega_{jt} - \mathbb{E}[\omega_{jt} \mid \omega_{j,<t}]\big)\, Z_{jt}\big] = 0, \qquad Z_{jt} \in \mathcal{I}_{jt},$$

thereby identifying the value–added function and the evolution of productivity. Elasticity of substitution and productivity parameters can be recovered accordingly.

[46] The α's can be identified as average expenditure shares.

D Additional figures and tables

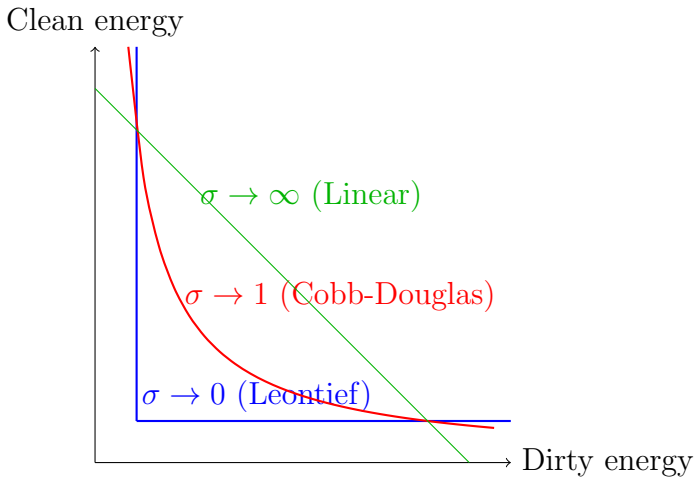


Figure A.1: Isoquants for clean vs. dirty energy under different elasticities

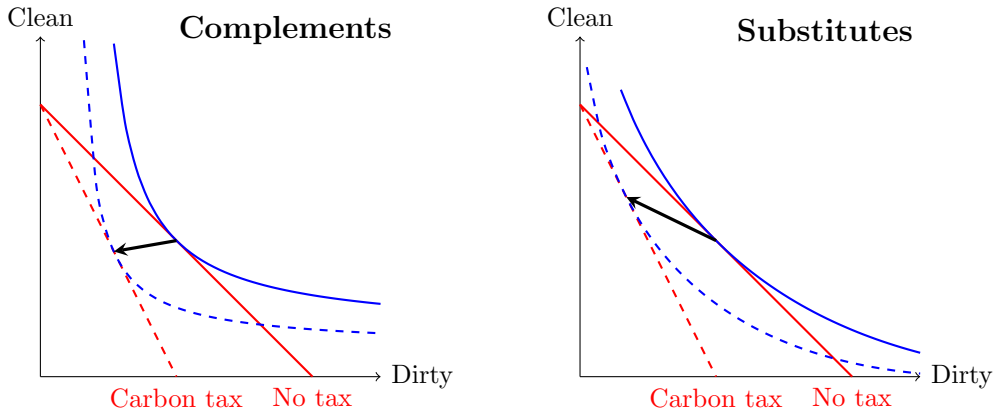


Figure A.2: Carbon tax: Weaker substitutes $\Longrightarrow$ less dirty-substitution effect

Notes: The figure shows the impact of carbon tax on energy demands within a simple context with financial constraints. Given a higher elasticity of substitution between clean and dirty, dirty energy use (emission) drops more due to the stronger substitution effect.

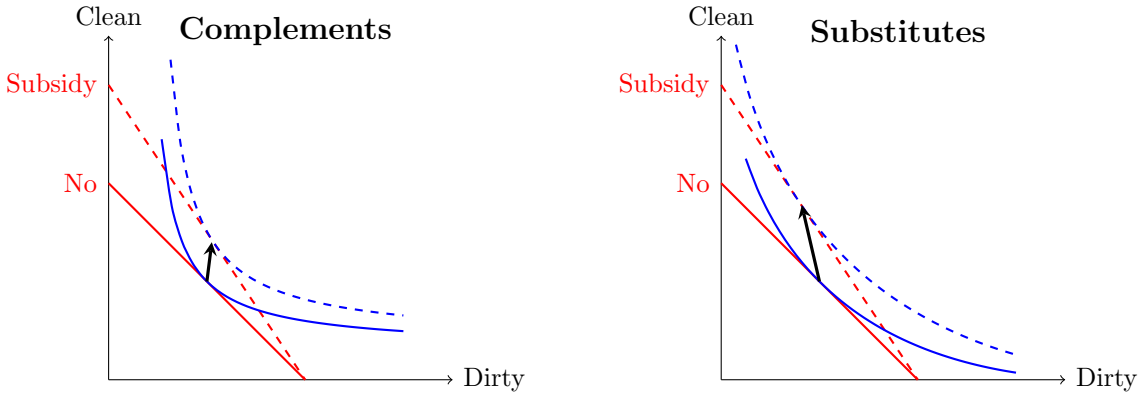


Figure A.3: Clean subsidy: weaker substitutes $\Longrightarrow$ less dirty-substitution effect

Notes: The figure shows the impact of clean subsidy on energy demands within a simple context with financial constraints. The scale effect increases dirty energy use. Given a higher elasticity of substitution between clean and dirty, dirty energy use (emission) drops more (or increases less) due to the stronger substitution effect.

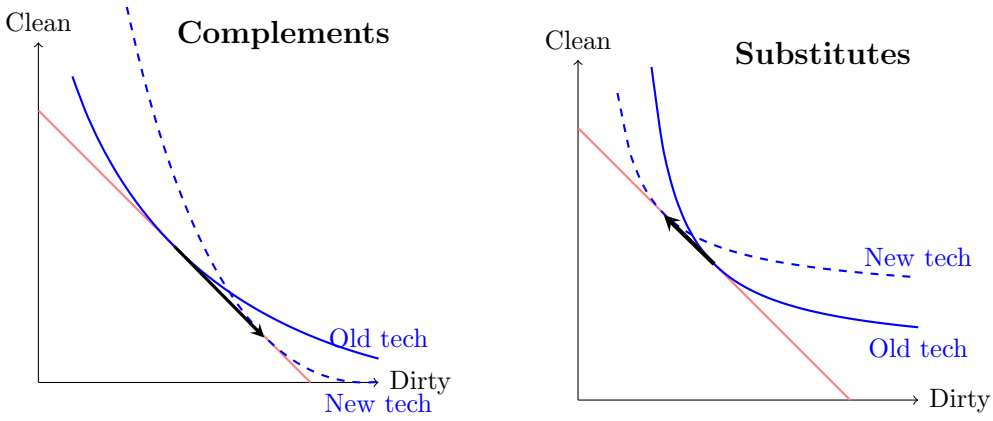


Figure A.4: Clean-augmenting technical change: weaker substitutes $\Longrightarrow$ less dirty-substitution effect

Notes: The figure shows the impact of augmenting on energy demands within a simple context with financial constraints. The scale effect increases dirty energy use. Given a higher elasticity of substitution between clean and dirty, dirty energy use (emission) drops more (or increases less) due to the stronger substitution effect.

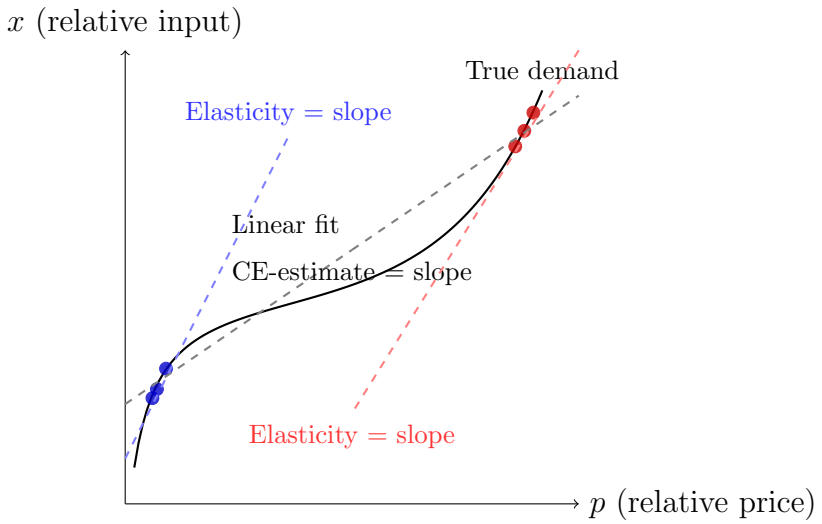


Figure A.5: Illustration of negative weight concern in CE estimation

Notes: The figure considers discrete distribution with two groups of firms with low (blue) and high (red) relative prices. Individual elasticity is the point derivative of the demand curve. Elasticity is first decreasing, then increasing within the gap. Accordingly, the weight on low-price firms are negative, and CE-estimate is outside of the range of individual point elasticity.

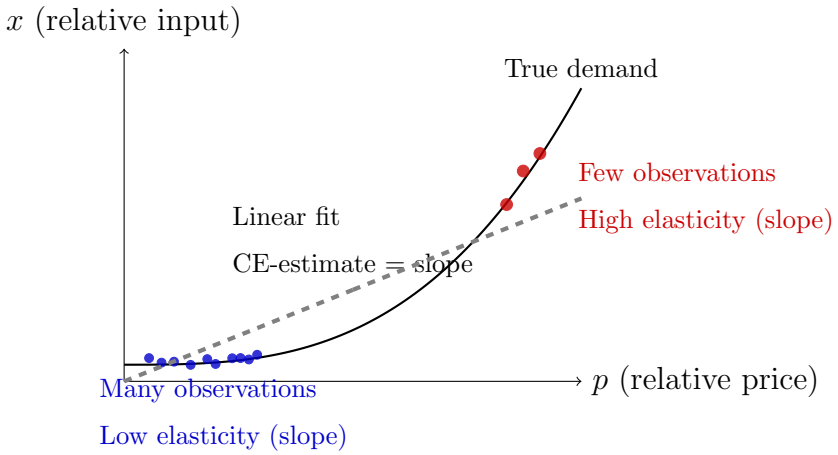


Figure A.6: Illustration of outlier concern in CE estimation

Notes: The figure considers discrete distribution with many low-price firms (blue) and several high-price firms (red, outliers). Individual elasticity is the point derivative of the demand curve. The elasticity of outliers are significantly higher, and they receive more weights in CE-estimates.

Table A.1: Translog coefficient estimates.

	Mine	ACF	TWFE		Mine	ACF	TWFE		Mine	ACF	TWFE
s_R	1.466***	0.999***	0.877***	α_{KK}	-0.028***	0.048**	-0.011	α_{KL}	-0.004	-0.033*	0.000
	(0.021)	(0.026)	(0.020)		(0.006)	(0.023)	(0.011)		(0.006)	(0.020)	(0.010)
α_K	0.407***	-0.033*	0.005	α_{LL}	0.026**	0.032	0.009	α_{KO}	-0.006	-0.008	-0.023**
	(0.008)	(0.020)	(0.027)		(0.011)	(0.026)	(0.013)		(0.07)	(0.025)	(0.011)
α_L	0.322***	0.006	0.075***	α_{OO}	-0.017*	0.010	0.101***	α_{KC}	0.052***	-0.008	-0.027***
	(0.005)	(0.022)	(0.012)		(0.009)	(0.046)	(0.021)		(0.006)	(0.021)	(0.011)
α_O	0.242***	0.861***	0.706***	α_{CC}	0.068***	0.029	0.085***	α_{KD}	-0.015***	0.001	0.007***
	(0.003)	(0.056)	(0.018)		(0.008)	(0.031)	(0.018)		(0.002)	(0.004)	(0.003)
α_C	0.020***	0.158***	0.151***	α_{DD}	0.011***	0.011	0.014***	α_{LO}	0.078***	0.026	0.023*
	(0.000)	(0.044)	(0.015)		(0.001)	(0.008)	(0.002)		(0.008)	(0.031)	(0.013)
α_D	0.009***	0.003	0.012***	α_{OC}	-0.066***	-0.010	-0.083***	α_{LC}	-0.073***	-0.021	-0.028**
	(0.000)	(0.009)	(0.003)		(0.008)	(0.008)	(0.016)		(0.009)	(0.026)	(0.011)
α_{CD}	0.019***	0.010	-0.000	α_{OD}	0.011***	-0.017	-0.018***	α_{LD}	-0.027***	-0.004	-0.004
	(0.003)	(0.007)	(0.003)		(0.003)	(0.012)	(0.004)		(0.003)	(0.007)	(0.003)

Notes: The table reports translog coefficient estimates of my method, "ACF" (control-function approach of Ackerberg *et al.* (2015)), and "TWFE" (two-way fixed effects used to absorb hicks-neutral productivity heterogeneity) in food & beverage industry.

Table A.2: Reduced from CE-based estimates.

	OLS					IV (lagged price)				
σ_{CD}	-0.66	-0.50	-0.66	-0.50	-0.50	-0.84	-0.86	-0.84	-0.73	-0.65
	(0.03)	(0.05)	(0.03)	(0.05)	(0.05)	(0.05)	(0.23)	(0.05)	(0.18)	(0.20)
N	3403	3403	3403	3403	3403	2103	2103	2103	2103	2103
R^2	0.51	0.80	0.52	0.81	0.82	0.26	0.67	0.25	0.67	0.67
Firm FE	N	Y	N	Y	Y	N	Y	N	Y	Y
Year FE	N	N	Y	Y	Y	N	N	Y	Y	Y
Industry·Year FE	N	N	N	N	Y	N	N	N	N	Y

Notes: The table reports reduced from CE-based estimates in food & beverage industry. In IV, lagged prices are used as instrument for current prices. All estimates are negative, underscore the efficiency of reduced form estimation. Standard errors clustered by firm are reported in parentheses.

Table A.3: Correlation between Log productivities.

	$\omega_{C,jt} - \omega_{D,jt}$				
$\omega_{R,jt}$	-0.007		-0.004	-0.006	-0.003
	(0.001)		(0.004)	(0.002)	(0.004)
$\omega_{L,jt}$		0.25	-0.24	-0.00	-0.10
		(0.09)	(0.11)	(0.12)	(0.12)
N	3476	3476	3476	3476	3476
R^2	0.03	0.01	0.65	0.04	0.65
Firm FE	N	N	Y	N	Y
Year FE	N	N	N	Y	Y

Notes: The table reports OLS estimates of regression log relative productivity between clean and dirty $\omega_{C,jt} - \omega_{D,jt}$ to hicks-neutral productivity $\omega_{R,jt}$ and labor-augmenting productivity $\omega_{L,jt}$. Standard errors clustered by firm are reported in parentheses.

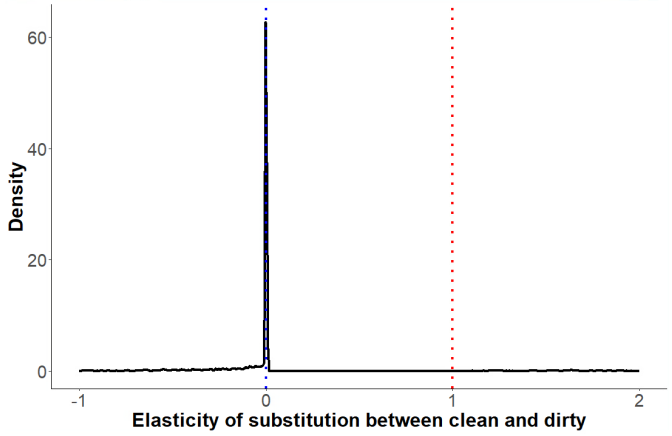


Figure A.7: Clean–dirty elasticity of substitution estimates based on my method: transport fuels as intermediate material

Notes: Distribution plot of the clean–dirty elasticity of substitution of each observation in food & beverage industry based on the production function estimates of my method when clarifying transport fuels as intermediate material.