

Microfoundations and the Causal Interpretation of Price-Exposure Designs

Luca Moreno-Louzada

*Stanford
University*

Guilherme Figueira

*São Paulo
School of Economics*

Pedro Picchetti

*Pontificia Universidad
Católica de Chile*

Work in progress

This version: December 2025

Abstract

This paper studies regional exposure designs that use commodity prices as instruments to study local effects of aggregate shocks. Unlike standard shift-share designs that leverage differential exposure to many shocks, the price-exposure relies on exogenous variation from a single shock, leading to challenges for both identification and inference. We motivate the design using a multi-sector labor model. Under the model and a potential outcomes framework, we characterize the 2SLS and TWFE estimands as weighted averages of region- and sector-specific effects plus contamination terms driven by the covariance structure of prices and by general-equilibrium output responses. We derive conditions under which these estimands have a clear causal interpretation and provide simple sensitivity analysis procedures for violations. Finally, we show that standard inference procedures suffer from an overrejection problem in price-exposure designs. We derive a new standard error estimator and show its desirable finite-sample properties through Monte Carlo simulations. In an application to gold mining and homicides in the Amazon, the price-exposure standard errors are roughly twice as large as conventional clustered standard errors, making the main effect statistically insignificant.

We thank Luis Alvarez, Bruno Ferman, Matthew Gentzkow, Isaac Sorkin, and Laura Weiwu for helpful feedback and suggestions. This is a work in progress; any comments are welcome. Emails: lucamlouzada@gmail.com; guilhermefigueira@live.com; pedropicchetti@gmail.com

1 Introduction

A pervasive empirical strategy in economics is to use regional-exposure designs with panel data to exploit heterogeneous exposure across locations facing a common shock.² This paper studies a particularly influential class of these strategies that leverages international commodity prices to identify local causal effects of a specific economic activity. The key empirical variable is typically constructed as:

$$Z_{it} = A_i \times P_t$$

where A_i denotes regional exposure which is fixed in time (e.g., mineral endowments), and P_t is a time-varying aggregate price which is the same for all units (e.g., the log international gold price). Unlike the typical shift-share design, which aggregates several shocks using weights for differential exposure, the *price-exposure design* focuses on a single shock.

This setup has been used, for example, to study the effect of different commodities on conflict (Dube and Vargas, 2013; Berman et al., 2017) and of cocaine production on child outcomes (Sviatschi, 2022).³ However, while this approach is ubiquitous, there is little formal discussion about its underlying identification assumptions and econometric properties. In particular, it is often unclear what causal parameter is recovered in a two-stage least squares (2SLS) or two-way-fixed-effects (TWFE) regressions, what assumptions about the first stage justify its causal interpretation, and how local labor-market dynamics or general equilibrium responses affect identification.

The first contribution of this paper is to provide microfoundations for the price–exposure design using a stylized multi-sector labor model. In our model, changes in the world price of a commodity in the sector of interest (we denote it as sector q) affect local wages and the allocation of labor across the other sectors, which determines the first stage relationship even for outcomes that are not themselves labor-market variables (e.g., human capital, crime, or conflict). This stylized environment makes explicit that a regression of Y_{it} on Z_{it} (or of Y_{it} on X_{it} instrumented with Z_{it}) implicitly imposes

²See Borusyak and Hull (2023) and Majerovitz and Sastry (2023)

³Other applications span the effects of oil exploitation on public finances (Martínez, 2023), health spending (Acemoglu, Finkelstein and Notowidigdo, 2013), crime (Soares and Souza, 2025), human capital (Balza et al., 2025), and political outcomes (Carreri and Dube, 2017; Fetzer and Kyburz, 2024); coffee production on schooling (Carrillo, 2020); maize on drug production (Dube, García-Ponce and Thom, 2016); and mining on human capital (Mejía, 2020), crime (Vásquez-Cortés, 2025), household consumption (Bazillier and Girard, 2020), deforestation (Girard, Molina-Millán and Vic, 2025), and proto-state formation (Sánchez De La Sierra, 2020).

strong restrictions on region-specific labor supply elasticities, sectoral demand parameters, and the way in which prices in different sectors move together. We derive closed-form expressions for the sector- q first stage and show that the slope of the relationship between Z_{it} and sector- q output is heterogeneous across regions and can even change sign when local labor supply is inelastic and exposure to co-moving sectors is high. This yields a transparent monotonicity condition for price-exposure designs.

Second, we place price-exposure regressions into a potential-outcomes framework and characterize the finite-population estimands of common estimators. We follow [Adão, Kolesár and Morales \(2019\)](#) — hereafter AKM — and treat the observed panel of regions and periods as the population of interest, with uncertainty coming from the stochastic process for prices. Under a linear potential-outcomes model in sectoral outputs, we show that the 2SLS estimand obtained by instrumenting sector- q output X_{iqt} with $Z_{iqt} = A_{iq}p_{qt}$ can be decomposed into two terms: (i) a weighted average of region-specific causal effects from an increase in the output of sector q , with weights proportional to the elasticity in the first-stage, exposure and the variance of the price in sector q ; and (ii) a contamination term that aggregates effects of other sectors β_{is} weighted by the elasticity from the other sectors, exposures, and the covariance of prices from other sectors with the price in sector q .

The IV estimand is therefore a convex average of sector- q effects only when (a) the first stage is monotone (the price-elasticity of the output in sector q is nonnegative for all regions) and (b) the focal price is uncorrelated with other prices. Otherwise, it is a “weakly causal” object in the sense of [Blandhol et al. \(2022\)](#), combining multiple channels in a model-dependent way. A similar contamination term appears even in the absence of price co-movement if the potential outcomes depend on the output of multiple sectors and general equilibrium responses lead to changes in sector- q price affecting outputs in other sectors.

Third, we clarify the interpretation of TWFE specifications that regress Y_{it} on $Z_{iqt} = A_{iq}p_{qt}$ with region and time fixed effects, a specification often described as a continuous-treatment differences-in-differences (DiD) or a reduced-form intention-to-treat (ITT). We show that the TWFE estimand can also be written as a weighted average of the same region-specific β_{iq} , plus a contamination term involving other sectors, but with weights that depend only on the variance and covariance of changes in prices. In contrast to IV, the TWFE estimand can be given a weakly causal decomposition under a weaker exogeneity assumption: mean-independence of Δp_{qt} from changes in potential outcomes and sectoral outputs, conditional on the finite-population filtration of effects and first

stage elasticities. This assumption resembles a parallel-trends restriction in first differences, but it operates on the common price shocks rather than on outcome trends. Our derivation highlights that this design is not the same as difference-in-differences and hence does not require the same identification assumptions. Under additional restrictions on the first stage (for instance, homogeneous elasticities equal to 1 for all regions), together with monotonicity and absence of cross-price comovement, the TWFE estimand coincides with a convex weighted average of sector- q effects. We revisit the reduced form specifications in [Dube and Vargas \(2013\)](#) and [Sviatschi \(2022\)](#) to discuss how including controls for other price shocks can affect the magnitudes of estimated effects.

Building on these decompositions, we derive sensitivity formulas that quantify how much the 2SLS estimand can be shifted by plausible values of the contamination term. We show how assumptions about the magnitude of the contamination term map into an identified set for the weighted average of interest, and we highlight “breakdown” values at which qualitative conclusions (e.g., the sign of the effect) change.

Finally, we study estimation and inference for price-exposure designs. Within the finite-population framework, we show the 2SLS estimator based on $Z_{iqt} = A_{iq}p_{qt}$ is consistent and asymptotically normal at rate \sqrt{NT} , with an asymptotic variance that takes a shift-share-type randomization form analogous to AKM, but exploiting time-series rather than cross-sectional variation. We use this characterization to construct a feasible “price-exposure” (PE) variance estimator that aggregates residuals across regions using only the variation in p_{qt} , and we show that it delivers valid inference under heterogeneous effects and arbitrary correlation in the outcome shocks. In contrast, conventional Eicker-Huber-White standard errors can severely understate sampling uncertainty when residuals inherit a shift-share structure from other prices, leading to systematic overrejection.

We document the magnitude of these distortions in Monte Carlo simulations under different combinations of N and T , and show that the PE standard errors achieve close-to-nominal coverage once the panel is moderately rich in the time dimension. To illustrate the practical relevance of these issues, we carry out a stylized empirical application to gold mining and homicides in the Brazilian Amazon. In this example, conventional clustered standard errors suggest a precisely estimated positive effect, whereas the PE standard errors are roughly twice as large and render the estimate statistically insignificant, highlighting how it can be important to take the specific sampling structure of price-exposure designs into account in empirical settings.

We conclude with brief guidance for applied work, emphasizing the importance of transparent discussion of the potential outcomes and identifying assumptions underlying causal claims. At a minimum, empirical researchers using price–exposure designs should state explicitly how the price process enters the relevant potential outcomes in their setting. Our goal is to offer a framework that clarifies the target parameters in price–exposure designs and serves as a starting point for these discussions.

Our analysis relates to long-standing questions in econometrics about the interpretation of causal parameters. Angrist, Graddy and Imbens (2000) showed that, under standard IV assumptions, the 2SLS estimator identifies a weighted average of marginal treatment effects. Recent work has extended these insights to characterizes how these weights depend on treatment effect heterogeneity (Blandhol et al., 2022; Alvarez and Toneto, 2024). A parallel literature studies difference-in-differences and two-way fixed-effects estimands as weighted averages of underlying causal effects, documenting the possibility of negative or counterintuitive weights in designs with a continuous treatment (de Chaisemartin and d’Haultfœuille, 2018; de Chaisemartin, d’Haultfœuille and Vazquez-Bare, 2024; Callaway, Goodman-Bacon and Sant’Anna, 2024). We build on these perspectives by showing that price-exposure estimands can be written as weighted averages of region- and sector-specific causal effects, with weights that have a transparent economic interpretation, and by characterizing when these weights are convex.

A second strand of the literature that we contribute to studies exposure designs, in which a common aggregate shock is transmitted through heterogeneous exposure shares (e.g., Goldsmith-Pinkham, Sorkin and Swift, 2020; Borusyak and Hull, 2023). Within this literature, we focus on the special case in which the aggregate shock is a single commodity price and develop a microfoundation and design-based inference strategy tailored to this setting. We build on results on shift-share identification with exogenous shocks (Adão, Kolesár and Morales, 2019; Borusyak, Hull and Jaravel, 2022), emphasizing a many-time-period rather than many-sector asymptotic framework, and on work showing how inference is complicated when exposure instruments load on a single aggregate shock (Arkhangelsky and Korovkin, 2019; Majerovitz and Sastry, 2023), exploiting the structure of prices to microfound and interpret the parameters of interest.

The rest of the article is organized as follows. We begin by discussing motivating examples of how the design has been implemented in previous empirical work in Section 1.1. Section 2 presents the economic model which microfounds the design. Section 3 develops a potential-outcomes framework

and studies the causal interpretation of IV and TWFE estimands. Section 4 discusses estimation and inference. Section 5 reports Monte Carlo simulations. Section 6 discusses examples from empirical applications, and Section 7 concludes.

1.1 Price Exposure Designs in the Applied Literature

Applied researchers are often interested in estimating the causal effect of some economic activity on an outcome. Price–exposure designs proxy for endogenous or unobserved local activity by interacting a time-invariant exposure A_i (e.g., geological endowments, suitability, historic production) with an aggregate price P_t common to all regions, building $Z_{it} = A_i \times P_t$. Two implementations are typical: (i) a reduced-form TWFE regression of outcomes on Z_{it} , interpreted as an ITT effect, and often called a variant of DiD; and (ii) an IV specification that instruments a measure of local activity X_{it} with Z_{it} . Both approaches implicitly assume a positive and sufficiently homogeneous first stage linking Z_{it} to X_{it} , an assumption rarely discussed in detail.

Below, we summarize a few examples of how this design has been used in applied work.⁴

Dube and Vargas (2013) The authors study the effects of commodity price shocks on violence in Colombia. Using a municipality-year panel with data from 1988 to 2005, they construct two price-exposure variables, $Z_{it}^{\text{oil}} = q_{i,1988}^{\text{oil}} \times p_t^{\text{oil}}$ and $Z_{it}^{\text{coffee}} = q_{i,1997}^{\text{coffee}} \times p_t^{\text{coffee}}$. Here $q_{i,0}^{\text{oil}}$ is used as the pre-period production, and Z_{mt}^{coffee} is instrumented with an interaction of local temperature and rainfall and production in other countries. They estimate it with 2SLS as a single equation including both oil and coffee, and call it a difference-in-differences estimator where time variation stems from movements in annual prices.

Berman et al. (2017) The study analyzes 55 km×55 km grid cells across Africa, defining treatment as the interaction $Z_{it} = \mathbb{1}\{\text{number of mines}_i > 0\} \times p_t^m$, where p_t^m is the international price of the mineral extracted in cell i . The authors estimate different specifications with time and cell-fixed effects, and do not elaborate on what the strategy is called. They find that conflict increases when the price rises.

⁴We simplify and adapt the original notation of the papers to standardize methodologies across different settings.

Carrillo (2020) The paper exploits municipality-level heterogeneity in pre-existing coffee cultivation to build the exposure term $Z_{it} = q_{i,0}^{\text{coffee}} \times p_t^{\text{coffee}}$, where $q_{i,0}^{\text{coffee}}$ is baseline coffee area and p_t^{coffee} is the world price of coffee. A cohort-specific reduced form specification compares children whose school-age years coincided with coffee booms to those who faced busts, controlling for municipality and birth-cohort fixed effects. The paper finds that higher coffee prices lower completed years of schooling and depress adult earnings.

Sviatschi (2022) The author studies Peruvian districts whose geographical conditions make them suitable for coca. Treatment is $Z_{it} = s_i^{\text{coca}} \times \widehat{p_t^{\text{coca}}}$, with s_i^{coca} fixed suitability and $\widehat{p_t^{\text{coca}}}$ the price of coca leaves in Peru instrumented by production in Colombia. A TWFE specification shows that price spikes raise child labor in coca farming, reduce schooling, and increase the probability of incarceration three decades later.

Mejía (2020) This work combines satellite-based measures of gold mining pits with Colombian administrative school records and defines exposure as $Z_{it} = s_i^{\text{gold}} \times p_t^{\text{gold}}$, where s_i denotes the standardized number of gold deposits in the vicinity of school i . Using school and year fixed effects, the author shows both reduced form TWFE estimates and 2SLS where the number of cells with detected mining activity is instrumented by Z_{it} . The paper finds that mining raises primary enrollment and lowers dropout, but harms test scores and tertiary enrollment.

Balza et al. (2025) The authors assemble well-level drilling data and micro-school enrollment records for Colombia. Their instrument is $Z_{it} = s_i^{\text{oil}} \times p_t^{\text{oil}}$, where s_i^{oil} is a proxy for local oil endowment i and p_t^{oil} is the global oil price. A 2SLS specification, with the first stage predicting the number of wells drilled, shows that oil exploitation reduces higher-education enrollment and delays college entry.

2 Economic Model

In this section, our goal is to derive a microfoundation for the first stage relationship between Z_{it} , the price-exposure instrument, and X_{it} , the local activity that researchers are interested in. We develop a simple multi-sector model with region-specific labor to illustrate how omitting explicit equilibrium conditions can complicate interpretation of estimated coefficients due to heterogeneous

elasticities between regions, general equilibrium responses, and co-movements in commodity prices. The model draws heavily from Appendix C in AKM.

Regions and sectors. Consider an economy composed of multiple regions, indexed by $i = 1, \dots, N$. Each region i is endowed with a fixed labor supply \bar{L}_i , which is fully employed and can be reallocated *across sectors* within a region but is *immobile across regions*. Sectors are indexed by $s \in \mathcal{S} \equiv \mathcal{S}_T \cup \{0\}$, where \mathcal{S}_T are tradable sectors with exogenous world prices P_{st} and $s = 0$ is a nontradable (internal) sector whose price is normalized to 1. Throughout, it is useful to think of a tradable sector as a sector producing a commodity such as gold; $P_{gold,t}$ is then the international price of gold at time t .

Production. In tradable sector $s \in \mathcal{S}_T$ and region i , output at time t is

$$X_{ist} = A_{is} E_{ist} K_{is}^{1-\theta_s} L_{ist}^{\theta_s}, \quad 0 < \theta_s < 1, \quad (1)$$

where $A_{is} > 0$ is a time-invariant productivity/capacity shifter (e.g. the number of mineral gold deposits in our running example), $E_{ist} > 0$ is a time-varying sector-region productivity shock, $K_{is} > 0$ is a sector-region capital stock (time-invariant within the short-run window), and L_{ist} is labor in sector s . The internal sector produces

$$I_{it} = B_i V_{it} K_{i0}^{1-\theta_0} N_{it}^{\theta_0}, \quad 0 < \theta_0 < 1.$$

with notation analogous to (1) and labor N_{it} .

Preferences/demand and market structure. Within each sector s , a standard CES aggregator across varieties with elasticity of substitution $\sigma_s > 1$ yields the usual isoelastic revenue schedule; firms are price takers up to a constant markup (which is absorbed into sectoral constants).

Wage equalization and labor supply. Within region i , labor is mobile across sectors, so a single wage w_{it} clears markets. Reduced-form labor supply in region i is

$$L_{it}^S = v_{it} w_{it}^{\phi_i}, \quad \phi_i > 0, \quad (2)$$

where v_{it} is a (possibly slowly varying) shifter capturing demographics, participation, etc. Total labor demand is $L_{it}^D \equiv \sum_{s \in \mathcal{S}_T} L_{ist} + N_{it}$. Equilibrium requires $L_{it}^S = L_{it}^D$.

Sectoral Labor Demand Profit maximization in (1) with CES demand implies, for each tradable $s \in \mathcal{S}_T$,

$$L_{ist} = \psi_{is} w_{it}^{-\sigma_s \delta_s} (A_{is} E_{ist} P_{st})^{(\sigma_s - 1) \delta_s} K_{is}^{(1 - \theta_s)(\sigma_s - 1) \delta_s}, \quad (3)$$

where $\psi_{is} > 0$ is a time-invariant constant that collects markups and sectoral demand primitives, and $\delta_s \equiv \frac{1}{1 + (\sigma_s - 1)(1 - \theta_s)}$. $\delta_s \in (0, 1)$ is a reduced-form parameter governing the sensitivity of sector- s labor demand to the revenue shifter $A_{is} E_{ist} P_{st}$, which depends on technology through θ_s and on the elasticity of substitution σ_s .

Taking logs and deviations from a reference equilibrium $(\cdot)^*$,

$$d \ln L_{ist} = (\sigma_s - 1) \delta_s d \ln P_{st} + (\sigma_s - 1) \delta_s d \ln E_{ist} - \sigma_s \delta_s d \ln w_{it}, \quad (4)$$

where all other terms in (3) are time-invariant and thus drop out in first differences.⁵

For the internal sector $s = 0$ (price 1) we obtain the same form with $P_{0t} \equiv 1$ and $E_{i0t} \equiv V_{it}$:

$$d \ln N_{it} = (\sigma_0 - 1) \delta_0 d \ln V_{it} - \sigma_0 \delta_0 d \ln w_{it}. \quad (5)$$

Normalization and Baselines Let $\ell_{is,0}^* \equiv L_{is}^*/L_i^*$ denote baseline employment shares (with $s = 0$ covering the internal sector). Fix baselines $\bar{P}_s, \bar{E}_{is}, \bar{v}_i$ and define centered shocks

$$p_{st} \equiv \ln P_{st} - \ln \bar{P}_s, \quad e_{ist} \equiv \ln E_{ist} - \ln \bar{E}_{is}, \quad v_{it} \equiv \ln v_{it} - \ln \bar{v}_i,$$

with $\mathbb{E}_t[p_{st}] = \mathbb{E}_t[e_{ist}] = \mathbb{E}_t[v_{it}] = 0$.

Equilibrium Wage Feedback With the normalizations above, linearizing the market-clearing condition $L_{it}^S = L_{it}^D$ using (2), (4), and (5) yields

$$\begin{aligned} \phi_i d \ln w_{it} + d \ln v_{it} &= \sum_{s \in \mathcal{S}_T} \ell_{is,0}^* \left[(\sigma_s - 1) \delta_s d \ln P_{st} + (\sigma_s - 1) \delta_s d \ln E_{ist} - \sigma_s \delta_s d \ln w_{it} \right] \\ &\quad + \ell_{i0}^* \left[(\sigma_0 - 1) \delta_0 d \ln V_{it} - \sigma_0 \delta_0 d \ln w_{it} \right]. \end{aligned} \quad (6)$$

Collecting the wage terms defines the key denominator

$$\lambda_i \equiv \phi_i + \sum_{s \in \mathcal{S}} \ell_{is,0}^* \sigma_s \delta_s > 0,$$

so that

$$d \ln w_{it} = \frac{1}{\lambda_i} \left\{ \sum_{s \in \mathcal{S}_T} \ell_{is,0}^* (\sigma_s - 1) \delta_s d p_{st} + \sum_{s \in \mathcal{S}} \ell_{is,0}^* (\sigma_s - 1) \delta_s d e_{ist} - d v_{it} \right\}. \quad (7)$$

⁵If capital K_{is} varies at horizon t , an additional term $(\sigma_s - 1) \delta_s (1 - \theta_s) d \ln K_{is}$ appears; we treat K_{is} as fixed over the short run.

Price co-movement. Allow other tradable prices to co-move with sector q :

$$dp_{st} = \rho_{qs} dp_{qt} + du_{st}, \quad \rho_{qq} = 1, \quad (8)$$

where du_{st} is orthogonal to dp_{qt} . Holding de and $d\nu$ orthogonal to dp_{qt} , the wage elasticity to the sector- q price is

$$\frac{\partial \ln w_{it}}{\partial \ln P_{qt}} = \frac{1}{\lambda_i} \sum_{s \in \mathcal{S}_T} \ell_{is,0}^* (\sigma_s - 1) \delta_s \rho_{qs}. \quad (9)$$

Proposition 1 (Heterogeneous first stage). *Under the model described above, sector- q output in region i admits the local levels approximation*

$$X_{iqt} \approx \alpha_{iq} + \kappa_{iq} [A_{iq} p_{qt}] + \varepsilon_{iqt}, \quad (10)$$

where the heterogeneous coefficient multiplying A_{iq} is

$$\kappa_{iq} \equiv \underbrace{\bar{E}_{iq} K_{iq}^{1-\theta_q} (L_{iq,0})^{\theta_q}}_{\text{pre-period capacity (fixed at baseline)}} \times (\sigma_q - 1) \theta_q \delta_q \left[1 - \frac{\sigma_q}{\sigma_q - 1} \cdot \frac{1}{\lambda_i} \sum_{s \in \mathcal{S}_T} \ell_{is,0}^* (\sigma_s - 1) \delta_s \rho_{qs} \right]. \quad (11)$$

with $\lambda_i = \phi_i + \sum_{s \in \mathcal{S}} \ell_{is,0}^* \sigma_s \delta_s$ and ρ_{qs} from (8).

Proof: See Appendix A.

Remark 1 (Monotonicity). Define the exposure index

$$S_{iq} \equiv \frac{1}{\lambda_i} \sum_{s \in \mathcal{S}_T} \ell_{is,0}^* (\sigma_s - 1) \delta_s \rho_{qs},$$

so that the elasticity of sector- q output with respect to its own price is

$$\frac{\partial \ln X_{iqt}}{\partial \ln P_{qt}} = (\sigma_q - 1) \theta_q \delta_q \left[1 - \frac{\sigma_q}{\sigma_q - 1} S_{iq} \right].$$

Since $A_{iq} > 0$ and the capacity term in (11) is strictly positive, monotonicity (a positive first stage) is equivalent to

$$\kappa_{iq} > 0 \iff \frac{\partial \ln X_{iqt}}{\partial \ln P_{qt}} > 0 \iff S_{iq} < \frac{\sigma_q - 1}{\sigma_q}.$$

If this inequality fails, the general-equilibrium wage effect dominates and the first stage slope $A_{iq} \kappa_{iq}$ becomes negative.

Intuition: when prices in other tradable sectors co-move with p_{qt} (large ρ_{qs}), those sectors also expand and bid up the local wage. This wage-drag offsets the direct effect of p_{qt} on sector q , and can overturn it if the co-moving bundle is strong enough.

The wage-drag term $\frac{\sigma_q}{\sigma_q-1} S_{iq}$ is large when: i) Local labor supply is inelastic (ϕ_i small), so λ_i is small and the wage moves a lot; ii) the region is highly exposed to co-moving tradables: $\sum_s \ell_{is,0}^* \rho_{qs}$ is large (e.g., concentrated mining hubs; many sectors rise together); iii) demand/congestion elasticities are high in those sectors (large σ_s or δ_s), amplifying their labor pull; and iv) co-movement is tight and positive ($\rho_{qs} \approx 1$ for many s). Conversely, the first stage remains positive when the region is diversified (small $\ell_{is,0}^*$ on co-moving sectors), labor supply is elastic (large $\phi_i \Rightarrow \lambda_i$ large), or co-movement is weak/negative (small or negative ρ_{qs}). If one allows nontradables to co-move as well (a local-demand channel), this enters exactly like an additional term in the sum (with $\rho_{q0} > 0$) and pushes toward violation.

Remark 2 (Economic interpretation of α_{iq}). With the normalization $p_{st} = \ln P_{st} - \ln \bar{P}$ and $e_{ist} = \ln E_{ist} - \ln \bar{E}_{is}$, the linearization in (10) is taken around the pre-period baseline (\bar{P}, \bar{E}) . The intercept equals the baseline level of sector- q output in region i :

$$\alpha_{iq} = X_{iq,0} = A_{iq} \bar{E}_{iq} K_{iq}^{1-\theta_q} (L_{iq,0})^{\theta_q}.$$

It is the quantity produced when the centered log price satisfies $p_{qt} = 0$ and the centered productivity shock satisfies $e_{iq,t} = 0$ (i.e., at baseline prices and productivity). Economically, α_{iq} captures time-invariant *capacity* combining geology A_{iq} , installed capital K_{iq} , and the baseline labor allocation $L_{iq,0}$ determined in the pre-period equilibrium across all sectors.

3 Potential Outcomes Framework and Causal Interpretation

In this section, we build on the results from the model presented in Section 2 and derive the causal interpretation of point estimates from two popular approaches towards the price-exposure design.

The starting point is to define potential outcomes in the price-exposure setting. The economic model with multiple sectors motivates a potential outcomes model which depends on multiple shifters. Let $Y_{it}(x_1, \dots, x_S)$ denote the potential outcome of region i in period t that would occur if the outputs of the S sectors were set to (x_1, \dots, x_S) . We assume that potential outcomes are linear functionals of the outputs in each sector:

$$Y_{it}(x_1, \dots, x_S) = \sum_{s=1}^S x_s \beta_{is} + \eta_{it} \quad (12)$$

Under this model, we have $\eta_{it} = Y_{it}(0, \dots, 0)$ as the potential outcome associated to output 0 across all sectors and β_{iq} as the causal effect of an increase of the output from sector q by one unit:

$$\frac{\partial Y_{it}(x_1, \dots, x_S)}{\partial x_q} = \beta_{iq}$$

We follow AKM and take a finite-population approach towards identification and inference. That is, we define the population of interest to be the observed set of N regions over T periods of time, instead of focusing on a large superpopulation over from which the observations are drawn. Uncertainty in the setting arises from the stochastic nature of the shock, while causal effects and potential outcomes are taken as given. Define $\mathcal{F}_0 \equiv \{\eta_{it}, \varepsilon_{ist}, \beta_{is}, \alpha_{is}, \kappa_{is}\}_{i=1, t=1, s=1}^{N, T, S}$ as the collection of parameters that determine the panel of causal effects and potential outcomes. We study the property of different estimators conditional on \mathcal{F}_0 , under the assumption that the price of the sector of interest is exogenous. Our setup can be summarized by the two following assumptions:

Assumption 1 (Potential outcomes and first stage).

(i) *The observed outcomes are generated by*

$$Y_{it} = Y_{it}(x_1, \dots, x_q, \dots, x_S),$$

where $Y_{it}(\cdot)$ follows the specification in equation (12).

(ii) *Output in each sector is determined according to the structure in Section 2, so that Proposition 1 holds.*

Assumption 1 can be seen as the usual exclusion restriction in this setting. Part (i) says that prices do not have a direct effect on Y_{it} . Instead, they only affect the outcome to the extent that they affect outputs. In particular, we rule out direct regional effects of p_{qt} on the outcome even holding the sectoral outputs fixed, such as direct local political responses to commodity price shocks. This corresponds to the common applied interpretation that the price-exposure instrument identifies the effect of “sector q activity” as long as there are no other direct channels of the price on the outcome. Thus Assumption 1(i) is purely a restriction on the way outcomes depend on sectoral activity; it does not by itself impose any statistical exogeneity condition on the price process.⁶

⁶Applied work often invokes small-open-economy intuition or external instruments for the commodity price to argue that the country is a price taker and local conditions do not feed back into p_{qt} . That kind of argument speaks to a separate exogeneity condition (in Assumption 2 below), not to the exclusion restriction as stated in Assumption 1(i).

What Assumption 1(i) does not restrict is how a shock to p_{qt} propagates across sectors: our multi-sector model implies that a change in p_{qt} will generally move outputs in many sectors, through both general-equilibrium labor reallocation and co-movement in other tradable prices. Whenever Y_{it} depends on these additional sectors (i.e., $\beta_{is} \neq 0$ for some $s \neq q$), standard price-exposure regressions implicitly combine the effect of sector- q activity with effects operating through other sectors. Characterizing how this multi-sector structure contaminates the usual estimands is one of our goals in this paper.

At the same time, in principle Assumption 1(i) need not require that there is a single economic mechanism from sector q to Y_{it} . In practice, the outcome might be affected by several economic channels related to sector q (such as employment or income in the gold mining sector), and not only by the actual output X_{iqt} (the amount of gold produced). The outcome is allowed to depend on employment, income, or other variables in sector q , as long as these are (possibly complicated) functions of X_{iqt} ; formally, these channels are absorbed into the response function $Y_{it}(\cdot)$. The restrictive component of the assumption, which we adopt for simplicity, is the linear form in equation (12), which compresses all such mechanisms for each sector into a one-dimensional sufficient statistic X_{iqt} and an additive effect β_{iq} . This is usually acknowledged by applied researchers studying the effect of “sector- q activity”, but it underscores the need to think carefully about how different choices of proxies for X_{iqt} (e.g., mined area, physical output, or sectoral employment) map into the economic effect that the resulting estimates are taken to represent.

Assumption 2 (Price exogeneity). *The price in sector q is as good as randomly assigned, in the sense that for all $t \in \{1, \dots, T\}$,*

$$\mathbb{E}[p_{qt} | \mathcal{F}_0] = 0$$

Assumption 2 imposes that the price shock in sector q is mean-independent from potential outcomes, given the mean-independence from η_{it} and the vector $(\beta_{i1}, \dots, \beta_{is})$, and also mean-independence from the potential “treatments” in each sector, which follows from mean-independence from $\{(\alpha_{is}, \kappa_{is}, \varepsilon_{ist})\}_{i,s,t}$. Thus, Assumption 2 can be interpreted as the standard exogeneity assumption in IV settings that the instrument is mean-independent from potential quantities, but expressed in terms of the parameters from the stylized economic model.⁷ In

⁷This assumption is substantively strong. Our multi-sector model suggests that general-equilibrium labor reallocation will cause outputs in other sectors to respond to changes in p_{qt} , even when the other sectoral prices $\{p_{st}\}_{s \neq q}$ are

applications, this is typically motivated by viewing the country as a small open economy that takes p_{qt} as given, or by constructing an external instrument for the commodity price (such as with shocks to global demand or to supply in other countries).

3.1 IV Estimands

We begin with the causal interpretation of IV estimates, which exploit the price-exposure as an instrument for output in sector q . We define the instrument as $Z_{iqt} \equiv A_{iq} \times p_{qt}$, where $p_{qt} \equiv \ln P_{qt} - \ln \bar{P}_{qt}$.⁸ In this setting, the 2SLS estimator stacked across time-periods of the coefficient on X_{iqt} is given by

$$\hat{\beta}_q^{2SLS} = \frac{\sum_{i=1}^N \sum_{t=1}^T Z_{iqt} Y_{it}}{\sum_{i=1}^N \sum_{t=1}^T Z_{iqt} X_{iqt}}$$

Given the finite-population framework, the 2SLS estimand is defined as the population analogue of the 2SLS estimator under repeated sampling of the shocks:

$$\beta_q^{2SLS} = \frac{\sum_{i=1}^N \sum_{t=1}^T \mathbb{E}[Z_{iqt} Y_{it} | \mathcal{F}_0]}{\sum_{i=1}^N \sum_{t=1}^T \mathbb{E}[Z_{iqt} X_{iqt} | \mathcal{F}_0]}$$

The following proposition establishes the causal interpretation of the 2SLS estimand:

Proposition 2. *Suppose that Assumption 2 holds and the potential outcomes model is correctly specified (i.e. Assumption 1 holds). Then,*

$$\beta_q^{2SLS} = \frac{\sum_{i=1}^N \sum_{t=1}^T \pi_{iqt} \beta_{iq}}{\sum_{i=1}^N \sum_{t=1}^T \pi_{iqt}} + \frac{\sum_{i=1}^N \sum_{t=1}^T \{\sum_{s \neq q} \pi_{isqt} \beta_{is}\}}{\sum_{i=1}^N \sum_{t=1}^T \pi_{iqt}}$$

where $\pi_{iqt} = \kappa_{iq} A_{iq}^2 \mathbb{V}(p_{qt} | \mathcal{F}_0)$ and $\pi_{isqt} = \kappa_{is} A_{is} A_{iq} \text{Cov}(p_{qt}, p_{st} | \mathcal{F}_0)$.

statistically independent of p_{qt} . Requiring mean-independence from the entire vector $\{\varepsilon_{ist}\}_{i,s,t}$ effectively ignores this channel. It amounts to imposing $\mathbb{E}[p_{qt} \varepsilon_{ist} | \mathcal{F}_0] = 0$ for all i, s, t , so any component of X_{ist} that moves systematically with p_{qt} must be captured in the structured part of the first stage rather than in ε_{ist} . In other words, we allow the outcome to move with the output in sector $s \neq q$, but only through co-movement of p_{st} with p_{qt} , ignoring the channel of output changes via changes in wages. We derive our main results under this assumption for expositional simplicity. In Section 3.1.1 we show how allowing such general-equilibrium terms to enter ε_{ist} generates additional contamination even when sectoral prices are mutually independent.

⁸Results are equivalent without demeaning the price, as long as time fixed effects are included.

Proof: See Appendix B.

Proposition 2 decomposes the IV estimand in two terms. The first is a weighted average of region-specific parameters associated to an increase in the output of sector q . The weights are increasing in the region-specific elasticity κ_{iq} and the region-specific exposure A_{iq} . The weights are positive as long as $\kappa_{iq} \geq 0$, which can be interpreted as the monotonicity condition from IV settings mapped to the economic model from Section 2 (see Remark 1).

The second term is a linear combination of region-specific parameters associated to increases in the output of all the other sectors stemming from movements in other prices. The sign of the coefficients associated to these causal parameters depends on the region-specific elasticity of the sectors, $\{\kappa_{is}\}_{s \neq q}$ and the covariances between the price of sector q and prices of other sectors. The 2SLS estimand can only be interpreted as a weighted average of causal effects in the case where the prices of all sectors are independent from the price in sector q . If all elasticities and price covariances are positive — and thus the monotonicity condition in Remark 1 holds —, then the 2SLS estimand holds a weakly causal interpretation in the Blandhol et al. (2022) sense.

In the case of panel data it is common to include time fixed effects to account for changes in the mean of the shift over time. Adding time fixed effects is equivalent to demeaning the price shock. Also, note that serial correlation in international prices does not affect the causal interpretation of the estimand, as the causal model rules out dynamic effects.

3.1.1 Causal Interpretation under General Equilibrium Effects

Proposition 2 provides the causal interpretation of the 2SLS estimand under the light of the economic model from Section 2. It follows from the causal decomposition that the contamination from causal effects of the output from other sectors affects the interpretation of the 2SLS estimand in the presence of price co-movement. The model, however, rules general equilibrium effects outside of price co-movement which might be relevant in applied settings.

In this section, we derive the causal interpretation of the 2SLS estimand considering a general equilibrium channel in which the prices of sector q affect the output in other sectors. To that end, we write the variable ε_{ist} as $\varepsilon_{ist} = \tilde{\varepsilon}_{ist} + \sum_{s' \neq s} \gamma_{iss'} p_{s't}$, where $\tilde{\varepsilon}_{ist}$ is idiosyncratic and the vector of parameters $\gamma_{iss'}$ captures how the output in sector s at time t responds to changes in prices from other sectors at the same time period.

Under these new specification of the model, we define the filtration as

$$\tilde{\mathcal{F}}_0 = \left\{ \eta_{it}, \alpha_{is}, \beta_{is}, (\gamma_{iss'})', \kappa_{is}, A_{is}, \tilde{\varepsilon}_{ist} \right\}_{i=1, t=1, s=1}^{N, T, S}$$

The exogeneity assumption we invoke is $\mathbb{E}[p_{st}|\tilde{\mathcal{F}}_0] = 0$. In the next proposition, we show that even when prices are independent, the 2SLS estimand can be decomposed as a weighted average of causal effects associated to the output in sector q and a contamination term.

Proposition 3. *Suppose Assumption 1 holds and that $\mathbb{E}[p_{st}|\tilde{\mathcal{F}}_0] = 0$ for all $s = 1, \dots, S$. Then,*

$$\beta_q^{2SLS} = \frac{\sum_{i=1}^N \sum_{t=1}^T \tilde{\pi}_{iqt} \beta_{iq}}{\sum_{i=1}^N \sum_{t=1}^T \tilde{\pi}_{iqt}} + \frac{\sum_{i=1}^N \sum_{t=1}^T \left\{ \sum_{s \neq q} \tilde{\pi}_{isqt} \beta_{is} \right\}}{\sum_{i=1}^N \sum_{t=1}^T \tilde{\pi}_{iqt}} + \frac{\sum_{i=1}^N \sum_{t=1}^T \left\{ \sum_{s \neq q} A_{iq} \gamma_{isq} \mathbb{V}(p_{qt}|\tilde{\mathcal{F}}_0) \beta_{is} \right\}}{\sum_{i=1}^N \sum_{t=1}^T \tilde{\pi}_{iqt}}$$

where $\tilde{\pi}_{iqt} = \kappa_{iq} A_{iq}^2 \mathbb{V}(p_{qt}|\tilde{\mathcal{F}}_0)$ and $\tilde{\pi}_{isqt} = \kappa_{is} A_{is} A_{iq} \text{Cov}(p_{qt}, p_{st}|\tilde{\mathcal{F}}_0)$.

Proof: See Appendix B.

Proposition 3 shows that if the prices of sector q directly affect the output in other sectors, then the 2SLS estimand does not hold a straightforward causal interpretation even in the absence of price co-movement. The estimand holds a weakly causal interpretation if for all $s \neq q$, the parameter γ_{isq} is nonnegative. That is, if an increase in the price of sector q affects positively the output in all sectors (or at least does not affect any sector negatively), then the estimand holds a weakly causal interpretation.

3.1.2 Sensitivity Analysis

Propositions 2 and 3 imply that the exogeneity of the price of sector q is not sufficient to provide a clear causal interpretation to the IV estimand. The two possible sources of contamination are (i) causal effects from other sectors affecting the estimand, which is driven by price co-movement and (ii) the effect of the price in sector q affecting the output of other sectors.

In the case where the prices and outputs from other sectors are observed by the applied researcher, the natural solution is to control for such quantities and close the channels of potential contamination of the estimand. Nevertheless, it is often the case in price-exposure designs that these quantities are not available for the researcher. Many applications, especially in the development economics literature, focus on the interaction between informal or illegal markets and legal markets, and is often the case that prices and outputs in illegal markets are not available.

Our framework motivates a simple procedure for sensitivity analysis based on the 2SLS estimand. When the sources of contamination cannot be observed in the data, researchers may place assumptions on the magnitude of the bias and obtain identified sets for the weighted average of causal effects from sector q and conduct inference on the partially identified causal estimand β_q .

Concretely, assume the contamination terms lie in the interval $[\underline{b}, \bar{b}]$ ⁹. Propositions 2 and 3 imply that β_q lies in the interval $[\beta_q^{LB}, \beta_q^{UB}]$, where

$$\beta_q^{LB} = \beta_q^{2SLS} - \bar{b}, \quad \beta_q^{UB} = \beta_q^{2SLS} - \underline{b}$$

A natural estimation procedure is to plug-in the 2SLS estimator $\hat{\beta}_q^{2SLS}$ in the expression above, which yields unbiased estimates for the bounds. By choosing different values for \underline{b} and \bar{b} , researchers can assess the robustness of the qualitative takeaways of the study.

3.2 Two-Way Fixed Effects Estimands

Another pervasive approach in the applied literature is to specify a two-way fixed effect (TWFE) regression, with time (θ_t) and individual (γ_i) fixed effects, and the price-exposure variable:

$$Y_{it} = \theta_t + \gamma_i + \beta^{TWFE} (A_{iq} p_{qt}) + u_{it}$$

For the sake of simplicity, we consider the case of panel data with two periods of time, so that the regression model can be expressed as

$$\Delta Y_i = (\theta_2 - \theta_1) + \beta^{TWFE} A_{iq} \Delta p_q + \Delta u_i$$

and the TWFE coefficient associated to the price-exposure variable as

$$\hat{\beta}_q^{TWFE} = \frac{\sum_{i=1}^N \Delta Z_{iq} \Delta Y_i}{\sum_{i=1}^N (\Delta Z_{iq})^2}$$

The TWFE estimand is defined as the population analogue of the TWFE estimator under repeated sampling of the shocks:

⁹The interpretation of these sensitivity parameters is straightforward in the case where only one of the channels of contamination poses a threat to the causal interpretation of the estimand. Nevertheless, the procedure is valid in all cases.

$$\beta^{TWFE} = \frac{\sum_{i=1}^N \mathbb{E}[\Delta Z_{iq} \Delta Y_i | \mathcal{F}_0]}{\sum_{i=1}^N \mathbb{E}[(\Delta Z_{iq})^2 | \mathcal{F}_0]}$$

It is common for applied researchers to claim a 'reduced form' or Intention-to-Treat (ITT) interpretation for the TWFE estimand. The proposition below provides the causal interpretation of the estimand:

Proposition 4. *Suppose that Assumption 2 holds and the potential outcomes model is correctly specified. Then,*

$$\beta^{TWFE} = \frac{\sum_{i=1}^N \kappa_{iq} (A_{iq})^2 \mathbb{V}(\Delta p_q | \mathcal{F}_0) \beta_{iq}}{\sum_{i=1}^N (A_{iq})^2 \mathbb{V}(\Delta p_q | \mathcal{F}_0)} + \frac{\sum_{i=1}^N \{ \sum_{s \neq q} \kappa_{is} A_{is} A_{iq} \text{Cov}(\Delta p_s, \Delta p_q | \mathcal{F}_0) \beta_{is} \}}{\sum_{i=1}^N (A_{iq})^2 \mathbb{V}(\Delta p_q | \mathcal{F}_0)}$$

Proof: See Appendix B.

Proposition 4 shows that the TWFE can be decomposed as a linear combination of causal parameters. The first term is a linear combination of region-specific effects of an increase in the output of sector q . The linear combination holds a weakly causal interpretation as long as monotonicity in the first-stage holds. However, it can be interpreted as a convex weighted average of causal effects only in the case when the A_{iq}^2 -weighted average of κ_{iq} equals one (which holds, for example, if $\kappa_{iq} = 1$ for all regions). The second term is a contamination term similar to the one in Proposition 2, but depends on the covariance of the variation of prices across time.

Hence, it follows that the TWFE estimand holds the interpretation of a weighted average of causal effects only under additional restrictions on the first-stage (which is why researchers usually claim a reduced form interpretation towards this estimand) and an additional assumption regarding the covariance of prices from different sectors.

Nevertheless, it is worth noting that the causal decomposition from Proposition 4 can be obtained under an exogeneity assumption weaker than Assumption 2. By exploiting within-regions variation across time, the time-invariant parameters from the model such as $\{\alpha_{is}\}_{s=1}^S$ are cancelled out. Moreover, define $\Delta \mathcal{F}_0 \equiv \{\Delta \eta_i, \Delta \varepsilon_{is}, \beta_{is}, \kappa_{is}\}_{i=1, s=1}^{N, S}$ to invoke the following exogeneity assumption:

Assumption 3 (Price exogeneity in trends). *The price in sector q is randomly assigned in the sense that*

$$\mathbb{E}[\Delta p_q | \Delta \mathcal{F}_0] = 0$$

Assumption 3 is weaker than Assumption 2, and can be interpreted as a parallel trends assumption for potential outcomes and potential output. It is different than the usual parallel trends assumption, however, as it assumes exogeneity conditional on the parameters $\{\beta_{is}\}_{s=1}^S$ and $\{\kappa_{is}\}_{s=1}^S$, which usually is not the case in TWFE/Two-stage TWFE settings. This is necessary because unlike the standard DiD settings, there is no time-period in which regions are fully untreated nor a time-period where the instrument has value zero for all regions. Contrary to a common interpretation in empirical work, the reduced form price-exposure design allows for non-parallel trends in Y_{it} across regions with different exposure levels, and the exposures A_i may be endogenous as long as innovations in p_{qt} are mean-independent of the finite-population primitives. Below, we show that the estimand β^{TWFE} maintains its causal decomposition under the weaker exogeneity assumption.

Corollary 1. *Suppose Assumption 3 holds. Then,*

$$\beta^{TWFE} = \frac{\sum_{i=1}^N \kappa_{iq}(A_{iq})^2 \mathbb{V}(\Delta p_q | \Delta \mathcal{F}_0) \beta_{iq}}{\sum_{i=1}^N (A_{iq})^2 \mathbb{V}(\Delta p_q | \Delta \mathcal{F}_0)} + \frac{\sum_{i=1}^N \left\{ \sum_{s \neq q} \kappa_{is} A_{is} A_{iq} \text{Cov}(\Delta p_s, \Delta p_q | \Delta \mathcal{F}_0) \beta_{is} \right\}}{\sum_{i=1}^N (A_{iq})^2 \mathbb{V}(\Delta p_q | \Delta \mathcal{F}_0)}$$

Corollary 1 provides a straightforward motivation for the TWFE approach towards the price-exposure design. Despite the fact that in general it only holds a weakly causal interpretation, the estimand holds its causal interpretation under a weaker exogeneity assumption than the one necessary for the IV estimand. In general, researchers choose the TWFE approach when they have access to the price-exposure variable but not to the actual variable of interest, which is why it is often referred to as a reduced form approach. Proposition 4 and Corollary 1 formalize such an interpretation in terms of the stylized economic model from Section 2.

4 Estimation and Inference

In this section, we analyze the statistical properties of the 2SLS estimator $\hat{\beta}_q^{2SLS}$. We consider large-sample properties of $\hat{\beta}_q^{2SLS}$ as the number of regions grows to infinity ($N \rightarrow \infty$), and the number of time periods grows to infinity ($T \rightarrow \infty$). Here we describe the main assumptions, technical regularity conditions are stated in Appendix C. Alongside Assumptions 1 and 2, we invoke the following assumption regarding the panel of sector prices:

Assumption 4 (Price independence).

- (i) For all $t \in \{1, \dots, T\}$, the prices (p_{1t}, \dots, p_{St}) are independent conditional on \mathcal{F}_0 ;
- (ii) For all $s \in \{1, \dots, S\}$, the prices (p_{s1}, \dots, p_{sT}) are independent conditional on \mathcal{F}_0 .

Assumption 4 (i) requires sector prices to be independent. Under this assumption, it follows that $\text{Cov}(p_{qt}, p_{st} | \mathcal{F}_0) = 0$ for all s , and thus the 2SLS estimand captures the weighted average of sector q causal effects without the contamination terms that affect its causal interpretation. Assumption 4 (ii) requires the sector prices to be independent across time periods. These assumptions combined extend to our setting the assumption underlying randomization style inference in randomized controlled trials that the treatment assignment is independent across units.

We now establish the asymptotic properties of the 2SLS is consistent and asymptotically normal, with \sqrt{NT} convergence rate.

Theorem 1. Suppose Assumptions 1, 2, 4, and Assumptions 5 and 6 from Appendix C hold. Then,

$$\hat{\beta}_q^{2SLS} = \beta_q^{2SLS} + o_p(1)$$

and

$$\sqrt{NT}(\hat{\beta}_q - \beta_q) \xrightarrow{d} \mathcal{N}\left(0, \frac{\mathcal{V}_{NT}}{\left(\frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T Z_{iqt} X_{iqt}\right)^2}\right)$$

where

$$\mathcal{V}_{NT} = \frac{1}{NT} \mathbb{V}\left(\sum_{i=1}^N \sum_{t=1}^T (A_{iq} p_{qt}) u_{it} | \mathcal{F}_0\right)$$

with u_{it} being the residual in the structural equation.

Proof: See Appendix C.

Theorem 1 shows that the asymptotic variance has the usual 'sandwich' format. To construct a consistent estimator for the standard error, it is sufficient to construct a consistent estimator for \mathcal{V}_{NT} . Consider first the case of homogeneous treatment effects, $\beta_{iq} = \beta$. Then it follows under Assumptions 2 and 4 that

$$\mathcal{V}_{NT} = \frac{1}{NT} \sum_{t=1}^T \mathbb{V}(p_{qt} | \mathcal{F}_0) \left(\sum_{i=1}^N A_{iq} u_{it}\right)^2$$

which motivates the following plug-in estimator for the variance:

$$\hat{V}_{PE}(\hat{\beta}) = \frac{\hat{V}_{NT}}{\left(\frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T Z_{iqt} X_{iqt}\right)^2}$$

where $\hat{V}_{NT} = \frac{1}{NT} \sum_{t=1}^T p_{qt}^2 \left(\sum_{i=1}^N A_{iq} \hat{u}_{it}\right)^2$. The estimator for the variance is similar to the one proposed by AKM, but instead of exploring variation across sectors, we exploit variation across time periods. The variance estimate leads to valid inference in the case of constant treatment effects, and inference remains valid under heterogeneous effects under additional regularity conditions stated in Appendix C.

As in AKM, our asymptotic results point to a possible overrejection problem associated to standard inference procedures. Take, for instance, the Eicker-Huber-White (or cluster-robust) standard errors. The expected difference between our proposed variance estimator and the robust variance estimator is a function of the correlation between residuals across regions and the exposures of these regions. Thus, if residuals are positively correlated, then it leads to an overrejection problem for the EHW variance estimator.

As AKM point out, "standard inference methods may overreject even when the unobserved shifters (prices from other sectors in our settings) contained in the residual vary along a different dimension than the shift-share (price-exposure in our setting) covariate of interest" if the variance of prices from other sectors is large and the correlation structure from the exposures in the residual is similar to the structure of the exposure in sector q . That seems to be the main channel driving overrejection in our setting. In the next section, we conduct a series of Monte Carlo simulations to assess how large N and T need to be for these asymptotics to provide a good approximation to the finite sample distribution of the 2SLS estimator.

4.1 Inference for Sensitivity Analysis

The results from Theorem 1 can be combined with existing results from the partial identification literature to produce valid confidence intervals for the partially identified causal parameter β_q . We adapt the confidence intervals from Imbens and Manski (2004) for our setting. Let $\mathcal{B} = \bar{b} - \underline{b}$ denote the length of the identified set. The Imbens-Manski confidence interval in the price-exposure setting takes the form

$$\left[\hat{\beta}_q^{2SLS} - C\sqrt{\hat{V}_{PE}(\hat{\beta})}, \hat{\beta}_q^{2SLS} + C\sqrt{\hat{V}_{PE}(\hat{\beta})} \right]$$

where C is a constant that solves $\Phi\left(\frac{\mathcal{B}}{\sqrt{\hat{V}_{PE}}} + C\right) - \Phi(-C) = 1 - \alpha$, with $\Phi(\cdot)$ being the standard normal CDF. The Imbens-Manski confidence interval can be used to conduct breakdown analysis in order to assess the robustness of qualitative takeaways from the study.

The breakdown point b^* is the value of contamination required to overturn a causal conclusion. For instance, how large must the contamination be in order for the Imbens-Manski interval to contain a null effect.

Standard error estimates for $\hat{\beta}_q^{2SLS}$ are weakly conservative, and thus Imbens-Manski CIs for the identified set will be conservative as well as estimated breakdown points. See Appendix D of [Rambachan and Roth \(2025\)](#) for a thorough discussion on inference in identified sets under design-based uncertainty.

5 Monte Carlo Simulations

In this section, we conduct a series of Monte Carlo experiments in order to study the finite sample properties of the 2SLS estimator and different inference procedures in price-exposure settings. The main objective of this exercise are (i) to verify that under Assumptions 1, 2, and 4 the 2SLS yields an unbiased estimate for the estimand β_q and (ii) to assess how large the panel must be in order for the asymptotic results from Theorem 1 to provide a good approximation for the finite sample distribution of the 2SLS estimator.

In order to do so, we consider the performance of the 2SLS estimator in terms of its mean bias, and the performance of the proposed standard error estimator (we label it *P-E*) and the Eicker-Huber-White estimator (we label it *Robust*) in terms of its empirical 95% coverage probability and mean standard error size under different sizes of the panel sample.

We consider a simple data generating process (DGP) with only two tradable sectors. For each of the time periods, the prices from each sector are drawn from uniform distributions with mean zero independent across the sectors and time periods, ensuring that Assumptions 1, 2, and 4 hold.

The idiosyncratic shock in the output of each sector, ε_{ist} , follows a Gaussian distribution. The exposures A_{is} are drawn from uniform distributions in which the lower limit is greater than zero.

The terms α_{is} are generated by multiplying A_{is} with a positive constant. Finally, the elasticity terms κ_{is} are drawn from a standard uniform distribution, which ensures that monotonicity holds.

The untreated potential outcomes η_{it} are drawn from a Gaussian distribution. The causal effects β_{is} follow a uniform distribution between 0 and 2. Under this definition for the values that comprise \mathcal{F}_0 , the estimand β_q^{2SLS} is equal to 1.

We document the performance of the standard error estimators over the randomization distribution. We first generate the filtration which generates potential outcomes and then simulate the data through different simulations of price sectors, holding potential treatments fixed. We conduct 1.000 Monte Carlo simulations for each exercise.

The finite sample properties of the 2SLS estimator and the inference procedures are compared under two scenarios. In the first, the number of regions is fixed at $N = 1.000$ and we study the performance of the estimators under different numbers of time periods. The results are summarized in Table 1.

In terms of bias, the 2SLS estimator exhibits desirable finite sample properties across all panel specifications. Naturally, the average bias exhibits its small value in magnitude when $T = 100$. Yet, even in the worst case scenario with a single time period available, the mean bias is smaller than 1% of the true causal parameter. In terms of inference, the simulation results illustrate the overrejection problem associated to standard inference procedures. The robust standard errors are substantially smaller than our proposed standard error estimator across all panel sizes. The empirical coverage of the 95% confidence interval goes from 0 in the cross-sectional case to 62% in the largest panel. The proposed $P-E$ estimator outperforms the robust standard error under in all panels, but only exhibits empirical coverage close to 95% when the number of time periods reaches triple digits.

In the second scenario, the number of time-periods is fixed at $T = 1.000$ and we study the performance of the standard error estimators under different numbers of regions. The results are displayed in Table 2.

Once again, the average bias of the 2SLS estimator is negligible across the different sizes of the panel. When it comes to inference, the robust standard error estimators yield severely biased confidence intervals across all specifications. The overrejection problem appears to be more dramatic than in the case of a large cross-section and a growing number of time periods. In the best case, the empirical coverage of the 95% confidence interval is 26%, almost half of what is observed in the best case under Scenario 1. The simulation results show that the $P-E$ standard error estimator only

Table 1: Monte Carlo Simulations - Scenario 1

	Mean Bias	Mean SE		Coverage	
		Robust	P-E	Robust	P-E
$N = 1.000, T = 1$	0.008	0.0	0.002	0	0
$N = 1.000, T = 5$	-0.009	0.048	0.519	0.369	0.560
$N = 1.000, T = 10$	0.007	0.032	0.376	0.463	0.717
$N = 1.000, T = 20$	0.001	0.051	0.382	0.471	0.818
$N = 1.000, T = 100$	-0.002	0.015	0.310	0.617	0.938

Note: Simulations based on 1.000 Monte Carlo experiments with sample size $N = 1.000$ under the DGP described in the Section. "Mean bias", "Mean SE" and "Coverage" stand for the mean simulated bias, mean simulated size of standard errors and 95% coverage

Table 2: Monte Carlo Simulations - Scenario 2

	Mean Bias	Mean SE		Coverage	
		Robust	P-E	Robust	P-E
$N = 1, T = 1.000$	-0.027	0.004	0.006	0	0
$N = 5, T = 1.000$	-0.001	0.023	0.049	0	0
$N = 10, T = 1.000$	0.010	0.042	0.099	0.214	0.554
$N = 20, T = 1.000$	-0.001	0.031	0.126	0.227	0.591
$N = 100, T = 1.000$	0.001	0.015	0.115	0.263	0.961

Note: Simulations based on 1.000 Monte Carlo experiments with sample size $N = 1.000$ under the DGP described in the Section. "Mean bias", "Mean SE" and "Coverage" stand for the mean simulated bias, mean simulated size of standard errors and 95% coverage

exhibits desirable finite-sample properties when $N = 100$. Under the other time periods, it exhibits a severe overrejection problem as well.

6 Empirical Applications

6.1 Other price shocks: evidence from existing reduced form applications

Our model implies that when multiple sectoral prices move together, local exposure to a "focal" price is typically correlated with exposure to other prices. In that case, the coefficient on a single price-exposure term conflates several channels unless other price exposures are included explicitly. A few

prominent applications already add such controls, but mostly present them as robustness checks. This section documents what those specifications do in [Dube and Vargas \(2013\)](#) and [Sviatschi \(2022\)](#), and how the magnitudes of their key coefficients change once other price shocks are included. Table 3 summarizes the main numbers.

Dube and Vargas (2013). [Dube and Vargas \(2013\)](#) study how international coffee and oil prices affect civil conflict in Colombian municipalities. Their baseline specification (Table 2) interacts log coffee and log oil prices with pre-determined measures of coffee intensity and oil production, and estimates effects on four measures of violence: guerrilla attacks, paramilitary attacks, clashes, and casualties. For paramilitary attacks, the coefficient on coffee intensity $\times \log p_t^{\text{coffee}}$ is -0.160 and the coefficient on oil production $\times \log p_t^{\text{oil}}$ is 0.726 .¹⁰

In Appendix Table A.VII, they extend this regression by adding interactions of log prices for other agricultural commodities (sugar, banana, palm, and tobacco) with corresponding pre-determined intensity measures. The coffee coefficient for paramilitary attacks remains negative and highly significant but becomes smaller in magnitude, at -0.125 . For guerrilla attacks, clashes, and casualties, the coffee coefficients move from -0.611 , -0.712 , and -1.828 in the baseline to -0.463 , -0.639 , and -1.192 once sugar, banana, palm, and tobacco price exposures are included. These are non-trivial adjustments—roughly a 20–30% reduction in the size of the coffee effect across outcomes—but they leave the qualitative pattern unchanged: higher coffee prices reduce violence in coffee-intensive areas, even after controlling for other commodity price shocks. In the original paper, the multi-commodity specification is presented as a robustness exercise about “other crops.” In our framework, it is exactly the sort of specification that recovers a cleaner causal interpretation of the coffee-price exposure coefficient in an environment with multiple relevant price processes.

Sviatschi (2022): child labor. [Sviatschi \(2022\)](#) studies how exposure to coca price shocks during childhood affects later outcomes. Her main child-labor specification (Table I, column (1)) regresses labor-force participation on a district-level coca price-exposure index, $PriceShock_{dt}$, interacted with age-group indicators. The coefficient on $PriceShock_{dt} \times \mathbf{1}\{11-14\}$ is 0.144 , capturing the additional increase in child labor at ages 11–14 relative to older ages. The baseline coefficient on $PriceShock_{dt}$ (the effect for the omitted age group) is 0.244 , so the implied total effect at ages 11–14 is $0.244 +$

¹⁰Table 2, column (2). The coefficients for guerrilla attacks, clashes, and casualties on coffee intensity $\times \log p_t^{\text{coffee}}$ are -0.611 , -0.712 , and -1.828 , respectively.

0.144 = 0.388.

Appendix Table A.II repeats this regression but adds exposure to price shocks for gold, coffee, cacao, cotton, and sugar, each constructed analogously by interacting international prices with pre-determined suitability or deposit measures. Once these additional commodity exposures are included, the coefficient on $PriceShock_{dt} \times \mathbf{1}\{11-14\}$ remains 0.144 (with a slightly smaller standard error), while the baseline $PriceShock_{dt}$ coefficient falls from 0.244 to 0.201. The implied total effect at ages 11–14 moves from 0.388 to 0.345. In contrast, some of the coefficients on other commodities are sizable and significant for child labor.

Sviatschi (2022): adult incarceration. For adult criminal outcomes, [Sviatschi \(2022\)](#) aggregates exposure to coca prices at ages 11–14 into a single index, $PriceShockAge11-14_{dc}$, and relates it to adult incarceration. In her main specification (Table III, column (1)), the coefficient on $PriceShockAge11-14_{dc}$ for the all-crimes incarceration rate is 3.607. Table V then adds exposure to price shocks for gold, coffee, cacao, cotton, and sugar (again interacted with pre-determined suitability or deposits). The coefficient on $PriceShockAge11-14_{dc}$ declines slightly to 3.523, while the other commodity exposures are small and statistically indistinguishable from zero. As in the child-labor regressions, the coca exposure coefficient at ages 11–14 is very stable, but the additional commodity shocks matter enough to be worth including if one seeks a structural interpretation of the coca coefficient.

Taken together, these examples deliver two messages that our framework clarifies. First, in both applications the focal price-exposure coefficient remains large and significant once other commodity price exposures are controlled for. This is reassuring about the robustness of their substantive conclusions. Second, the fact that the coefficients do move, sometimes by 10–30%, highlights that omitting other prices changes the object estimated by the focal price-exposure coefficient. In [Dube and Vargas \(2013\)](#) and [Sviatschi \(2022\)](#), these multi-commodity specifications are presented as robustness checks. Our analysis suggests that, whenever the outcome is plausibly affected by several correlated sectoral prices, such specifications should instead be treated as the preferred empirical implementation of a price-exposure design.

Table 3: Price-Exposure Coefficients With and Without Controls for Other Price Shocks

	Baseline specification (1)	+ Other price shocks (2)
<i>Panel A: Dube and Vargas (2013), Colombia</i>		
Guerrilla attacks: coffee int. $\times \log p_t^{\text{coffee}}$	−0.611** (0.249)	−0.463*** (0.163)
Paramilitary attacks: coffee int. $\times \log p_t^{\text{coffee}}$	−0.160*** (0.061)	−0.125*** (0.036)
Clashes: coffee int. $\times \log p_t^{\text{coffee}}$	−0.712*** (0.246)	−0.639*** (0.195)
Casualties: coffee int. $\times \log p_t^{\text{coffee}}$	−1.828* (0.987)	−1.192* (0.639)
<i>Panel B: Sviatschi (2022), Peru</i>		
Child labor (age 11–14): coca price shock $\times 1\{11-14\}$	0.144*** (0.028)	0.144*** (0.036)
Adult incarceration: coca price shock at ages 11–14	3.607*** (0.917)	3.523*** (0.901)

Note: Each entry reports the coefficient and standard error on the focal price-exposure term in the cited specification, taken directly from the original tables. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Panel A uses the coefficients on coffee intensity $\times \log$ coffee price from Dube and Vargas (2013). Column (1) reproduces Table 2, columns (1)–(4), which include exposure to coffee and oil price shocks. Column (2) uses Appendix Table A.VII, which augments the regression with exposure to price shocks for sugar, banana, palm, and tobacco. Panel B uses Sviatschi (2022). For child labor, column (1) reports the coefficient on $PriceShock_{dt} \times 1\{11-14\}$ from Table I, column (1); column (2) reports the same coefficient from Appendix Table A.II, which adds exposure to price shocks for gold, coffee, cacao, cotton, and sugar. For adult incarceration, column (1) reports the coefficient on $PriceShock_{Age11-14_{dc}}$ from Table III, column (1), and column (2) reports the corresponding coefficient from Table V, column (1), which adds the same set of other commodity price exposures.

6.2 Inference in 2SLS: an application to gold mining and violence

In this section, we present a stylized empirical application using data on gold mining and homicides in the Brazilian Amazon to compare conventional cluster-robust standard errors with the price-exposure standard errors developed in Section 4. This application is inspired by the setting in Pereira and Pucci (2024), who study the impact of gold market deregulation on violence. Their paper does not use a price-exposure design; instead, they implement a DiD design exploiting a 2013 deregulation law that weakened monitoring of gold transactions. They show that deregulation increased illegal gold mining in protected areas and that municipalities more exposed to illegal mining experienced

large increases in homicides after the deregulation.

Because their analysis clearly documents a mechanism linking gold-related rents, illegal mining, and local violence, this setting is a natural context in which to illustrate the implications of using a price–exposure design. In what follows, we implement a deliberately naive price–exposure specification that interacts the number of gold deposits in a municipality with the international gold price. Our goal is purely illustrative: unlike [Pereira and Pucci \(2024\)](#), we abstract from the institutional details of the 2013 reform, omit additional controls, and do not interpret the resulting coefficients as new causal estimates.

We built a panel for 591 municipalities, observed between 2000 and 2022. For each municipality i and year t , we observe the homicide rate per 100,000 inhabitants, Y_{it} , the area of gold mining within the municipality in square kilometers, X_{it} , and the number of identified gold deposits, A_i .¹¹ We match these data to the international gold price p_t and construct the standard price–exposure instrument

$$Z_{it} \equiv (A_i) \log p_t.$$

We estimate the canonical price–exposure specification with municipality and year fixed effects. Specifically, we begin with a first-stage regression of mining area on the price–exposure instrument,

$$X_{it} = \alpha_i^X + \lambda_t^X + \pi Z_{it} + v_{it},$$

and a reduced-form regression of homicide rates on the same instrument,

$$Y_{it} = \alpha_i^Y + \lambda_t^Y + \rho Z_{it} + \eta_{it}.$$

Second, we estimate an ordinary least squares (OLS) regression of Y_{it} on X_{it} ,

$$Y_{it} = \alpha_i + \lambda_t + \beta^{\text{OLS}} X_{it} + u_{it},$$

and a 2SLS specification that instruments X_{it} with Z_{it} in the same two-way fixed-effects framework.

With fitted values

$$\hat{X}_{it} = \hat{\alpha}_i^X + \hat{\lambda}_t^X + \hat{\pi} Z_{it},$$

the second stage takes the form

$$Y_{it} = \alpha_i + \lambda_t + \beta^{2\text{SLS}} \hat{X}_{it} + \varepsilon_{it}.$$

¹¹ For more details about the data and summary statistics, see [Appendix D](#).

All regressions include municipality and year fixed effects and cluster standard errors by municipality. The results are summarized in Table 4. Higher gold prices are strongly associated with increases in mining activity, and the reduced-form and 2SLS estimates suggest that expansions in mining are accompanied by higher homicide rates, consistent with [Pereira and Pucci \(2024\)](#).

Table 4: Gold Prices, Mining, and Homicides in the Brazilian Amazon

	Mining area		Homicide Rate	
	FS: X on Z	RF: Y on Z	OLS: Y on X	2SLS: Y on \hat{X}
	(1)	(2)	(3)	(4)
Price exposure Z_{it}	0.161*** (0.023)	0.086*** (0.026)		
Mining area			0.320*** (0.114)	0.535** (0.220)
Municipality FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Clusters	Municipality	Municipality	Municipality	Municipality
Observations	13,593	13,593	13,593	13,593
R^2	0.807	0.374	0.374	0.373

Note: This table reports first-stage, reduced-form, OLS, and 2SLS estimates relating gold prices, mining area, and homicide rates in Brazilian Amazon municipalities. All specifications include municipality and year fixed effects. Standard errors, clustered by municipality, are reported in parentheses. See Appendix D for details about the variables. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

We then focus on the 2SLS specification and compare standard errors. Using the same underlying regressions, we compute both the usual cluster-robust variance estimator and the price–exposure variance derived in Section 4. Table 5 reports the 2SLS coefficient on mining area together with (i) the municipality-clustered standard error and (ii) the price–exposure standard error. Numerically, the point estimate of β is around 0.54: an additional square kilometer of mining area is associated with roughly 0.5 additional homicides per 100,000 inhabitants. The cluster-robust standard error is about 0.22, implying statistical significance at conventional levels, whereas the price–exposure standard

error is roughly twice as large (about 0.48) and renders the estimate statistically indistinguishable from zero at the 5% level.

Table 5: 2SLS Estimate with Clustered and Price–Exposure Standard Errors

	Homicide rate	
	Clustered SE	Price–Exposure SE
	(1)	(2)
Mining area	0.535** (0.220)	0.535 (0.476)
Observations	13,593	13,593
R ²	0.373	0.373

Note: This table reports 2SLS estimates of the effect of mining area on homicide rates in Brazilian Amazon municipalities. Column (1) reports the conventional standard error clustered by municipality; column (2) reports the corresponding price–exposure standard error. See Appendix D for details about the variables. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

This simple exercise illustrates our main message. A standard empirical implementation of a price–exposure design, using off-the-shelf 2SLS and clustered standard errors, would lead to the conclusion that mining has a precisely estimated effect on violence in this setting. Once we account for the sampling structure implied by the price–exposure design and use the corresponding price–exposure variance, the point estimate is unchanged but the uncertainty around it increases substantially, and the result is no longer statistically significant. Interpreting these estimates substantively would require exactly the kind of careful work done by [Pereira and Pucci \(2024\)](#): tracing the institutional details of deregulation, documenting the mechanisms linking rents to illegal mining and violence, and specifying a research design that aligns with those mechanisms.

7 Conclusion

It is a common empirical strategy in applied microeconomics to estimate treatment effects using price-exposure designs.

In this paper, we develop a multi-sector labor model that microfound the approach and clarifies what 2SLS and continuous-treatment TWFE estimates recover. Within a finite-population potential-outcomes framework, we show that both estimands can be written as weighted averages of region- and sector-specific causal effects of the focal price, plus contamination terms that depend on causal effects of other sectoral outputs whenever prices co-move or general-equilibrium channels are active. We characterize transparent conditions under which these estimands admit a causal interpretation.

We then study estimation and inference in these designs and quantify the consequences of ignoring the shift-share structure of the residual. We derive a randomization-style asymptotic variance for the 2SLS estimator, show in simulations that conventional Eicker-Huber-White and cluster-robust standard errors can substantially overreject in empirically relevant settings, and propose a feasible price-exposure variance estimator that achieves better coverage. Re-examining prominent applications through this lens, we conclude that price-exposure designs remain powerful tools for applied work, but that credible causal interpretation requires explicit modeling of mechanisms, careful discussion of price co-movement and general-equilibrium forces, and design-consistent inference.

Our results have several implications for applied work. In practice, we recommend that researchers (i) make the theory of change and the associated potential-outcomes model explicit, clarifying which sectoral outputs are allowed to enter the outcome and how the price shock is assumed to operate; (ii) discuss co-movement and general equilibrium channels, and control for other relevant price shocks or sectoral outputs whenever possible, especially in settings where multiple commodities or sectors plausibly matter; (iii) report first stages whenever a measure of sectoral activity is available, and discuss the plausibility of monotonicity and the sources of heterogeneity in the price-output relationship; (iv) be explicit about the validity of exogeneity of the focal price process, which is the central identifying assumption in most settings, rather than parallel trends in outcomes or the exogeneity of exposure shares; and (v) treat short panels with particular caution when conducting inference, given that standard errors that ignore the shift-share structure of residuals can exhibit substantial overrejection in designs with limited time-series variation.

This is still a work in progress. For the next steps, we intend to derive explicit equivalencies to results from previous work, improve the empirical applications, study inference for TWFE and in settings with serial correlation or co-movement in prices, examine implications of adding covariates

to the estimation, and expand the labor model to allow for migration and time-varying changes in capital or other structural parameters.

References

- Acemoglu, Daron, Amy Finkelstein, and Matthew J Notowidigdo.** 2013. “Income and health spending: Evidence from oil price shocks.” *Review of Economics and Statistics*, 95(4): 1079–1095.
- Adão, Rodrigo, Michal Kolesár, and Eduardo Morales.** 2019. “Shift-share designs: Theory and inference.” *The Quarterly Journal of Economics*, 134(4): 1949–2010.
- Alvarez, Luis A. F., and Rodrigo Toneto.** 2024. “The interpretation of 2SLS with a continuous instrument: A weighted LATE representation.” *Economics Letters*, 237: 111658.
- Angrist, Joshua D, Kathryn Graddy, and Guido W Imbens.** 2000. “The interpretation of instrumental variables estimators in simultaneous equations models with an application to the demand for fish.” *Review of Economic Studies*, 67(3): 499–527.
- Arkhangelsky, Dmitry, and Vasily Korovkin.** 2019. “On policy evaluation with aggregate time-series shocks.” arXiv preprint 1905.13660.
- Balza, Lenin H, Camilo De Los Rios, Raul Jimenez Mori, and Osmel Manzano.** 2025. “The local human capital costs of oil exploitation.” *Journal of Development Economics*, 173: 103410.
- Bazillier, Remi, and Victoire Girard.** 2020. “The gold digger and the machine. Evidence on the distributive effect of the artisanal and industrial gold rushes in Burkina Faso.” *Journal of Development Economics*, 143: 102411.
- Berman, Nicolas, Mathieu Couttenier, Dominic Rohner, and Mathias Thoenig.** 2017. “This Mine is Mine! How Minerals Fuel Conflicts in Africa.” *American Economic Review*, 107(6): 1564–1610.
- Blandhol, Christine, John Bonney, Magne Mogstad, and Alexander Torgovitsky.** 2022. “When is TSLS actually late?” NBER Working Paper 29709.
- Borusyak, Kirill, and Peter Hull.** 2023. “Nonrandom exposure to exogenous shocks.” *Econometrica*, 91(6): 2155–2185.

- Borusyak, Kirill, Peter Hull, and Xavier Jaravel.** 2022. “Quasi-experimental shift-share research designs.” *Review of Economic Studies*, 89(1): 181–213.
- Callaway, Brantly, Andrew Goodman-Bacon, and Pedro HC Sant’Anna.** 2024. “Difference-in-differences with a continuous treatment.” NBER Working Paper 32117.
- Carreri, Maria, and Oeindrila Dube.** 2017. “Do natural resources influence who comes to power, and how?” *The Journal of Politics*, 79(2): 502–518.
- Carrillo, Bladimir.** 2020. “Present bias and underinvestment in education? Long-run effects of childhood exposure to booms in Colombia.” *Journal of Labor Economics*, 38(4): 1127–1265.
- de Chaisemartin, Clément, and Xavier d’Haultfœuille.** 2018. “Fuzzy differences-in-differences.” *Review of Economic Studies*, 85(2): 999–1028.
- de Chaisemartin, Clément, Xavier d’Haultfœuille, and Gonzalo Vazquez-Bare.** 2024. “Difference-in-difference estimators with continuous treatments and no stayers.” *AEA Papers and Proceedings*, 114: 610–613.
- Dube, Oeindrila., and Juan F. Vargas.** 2013. “Commodity price shocks and civil conflict: Evidence from Colombia.” *Review of Economic Studies*, 80(4): 1384–1421.
- Dube, Oeindrila, Omar García-Ponce, and Kevin Thom.** 2016. “From maize to haze: Agricultural shocks and the growth of the mexican drug sector.” *Journal of the European Economic Association*, 14(5): 1181–1224.
- Fetzer, Thiemo, and Stephan Kyburz.** 2024. “Cohesive institutions and political violence.” *Review of Economics and Statistics*, 106(1): 133–150.
- Girard, Victoire, Teresa Molina-Millán, and Guillaume Vic.** 2025. “Artisanal mining in Africa. Green for Gold?” *Economic Journal*, 135(672): 2578–2597.
- Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift.** 2020. “Bartik instruments: What, when, why, and how.” *American Economic Review*, 110(8): 2586–2624.
- Imbens, Guido W., and Charles F. Manski.** 2004. “Confidence Intervals for Partially Identified Parameters.” *Econometrica*, 72(6): 1845–1857.

- Majerovitz, Jeremy, and Karthik Sastry.** 2023. “How Much Should We Trust Regional-Exposure Designs?” SSRN preprint 4433676.
- Martínez, Luis R.** 2023. “Natural Resource Rents, Local Taxes, and Government Performance: Evidence from Colombia.” *Review of Economics and Statistics*, 1: 1–28.
- Mejía, Leonardo Bonilla.** 2020. “Mining and human capital accumulation: Evidence from the Colombian gold rush.” *Journal of Development Economics*, 145: 102471.
- Pereira, Leila, and Rafael Pucci.** 2024. “A Tale of Gold and Blood: The Consequences of Market Deregulation on Local Violence.” FEA-USP Economics WP 2024-18.
- Rambachan, Ashesh, and Jonathan Roth.** 2025. “Design-Based Uncertainty for Quasi-Experiments.” *Journal of the American Statistical Association*, Forthcoming.
- Sánchez De La Sierra, Raúl.** 2020. “On the origins of the state: Stationary bandits and taxation in eastern congo.” *Journal of Political Economy*, 128(1): 000–000.
- Soares, Rodrigo R, and Danilo Souza.** 2025. “Too much of a good thing: Accelerated growth and crime.” *Journal of Development Economics*, 103499.
- Sviatschi, Maria Micaela.** 2022. “Making a narco: Childhood exposure to illegal labor markets and criminal life paths.” *Econometrica*, 90(4): 1835–1878.
- Vásquez-Cortés, Mateo.** 2025. “Peer effects and recidivism: wartime connections and criminality among Colombian ex-combatants.” *American Political Science Review*, 119(2): 637–652.

Appendices

A Labor Model Derivation

In this appendix we verify the expressions used in Section 2 with an expanded microfoundation and derive the first-stage equation linking tradable-sector output to the exogenous price in Proposition 1.

A.1 Microfoundation of sectoral labor demand

To microfound (3), we enrich the environment in Section 2 with standard assumptions on within-sector differentiation and price setting, while keeping the regional structure and production technology in (1) unchanged.¹² Start from a firm producing a differentiated variety $\omega \in \Omega$ in sector $s \in \mathcal{S}_T$ and region i at time t with Cobb–Douglas technology

$$x_{ist}(\omega) = A_{is}E_{ist} k_{ist}(\omega)^{1-\theta_s} l_{ist}(\omega)^{\theta_s}, \quad 0 < \theta_s < 1, \quad (13)$$

where $k_{ist}(\omega)$ and $l_{ist}(\omega)$ are capital and labor for variety ω . The sector–region capital stock is $K_{is} \equiv \int_{\Omega_s} k_{ist}(\omega) d\omega$, which is fixed in the short run.

Unit cost. Given wages w_{it} and rental rate r_{ist} , a firm that wants to produce one unit of output solves

$$\min_{k,l} w_{it}l + r_{ist}k \quad \text{s.t.} \quad A_{is}E_{ist}k^{1-\theta_s}l^{\theta_s} \geq 1.$$

The Lagrangian

$$\mathcal{L} = w_{it}l + r_{ist}k + \lambda \left[1 - A_{is}E_{ist}k^{1-\theta_s}l^{\theta_s} \right]$$

yields the first–order conditions

$$\begin{aligned} \partial_l \mathcal{L} = 0 &\Rightarrow w_{it} = \lambda A_{is}E_{ist} \theta_s k^{1-\theta_s} l^{\theta_s-1}, \\ \partial_k \mathcal{L} = 0 &\Rightarrow r_{ist} = \lambda A_{is}E_{ist} (1 - \theta_s) k^{-\theta_s} l^{\theta_s}. \end{aligned}$$

Taking the ratio,

$$\frac{w_{it}}{r_{ist}} = \frac{\theta_s}{1 - \theta_s} \frac{k}{l} \Rightarrow \frac{k}{l} = \frac{1 - \theta_s}{\theta_s} \frac{w_{it}}{r_{ist}}.$$

At the sector level (under symmetry of varieties),

$$\frac{K_{is}}{L_{ist}} = \iota_s \frac{w_{it}}{r_{ist}}, \quad \iota_s \equiv \frac{1 - \theta_s}{\theta_s} > 0. \quad (14)$$

Substituting the optimal factor demands into the cost function $c_{ist}(w_{it}, r_{ist}, A_{is}E_{ist})$ gives the dual of (13):

$$c_{ist} = \chi_s w_{it}^{\theta_s} r_{ist}^{1-\theta_s} (A_{is}E_{ist})^{-1}, \quad (15)$$

where $\chi_s > 0$ depends only on θ_s .

¹²This step is identical in spirit to Appendix C.1 of AKM. We follow similar algebra but adapt all notation to our setting and allow for the additional multiplicative shock E_{ist} . Our equation (3) is analogous to their equation C.4.

Demand and revenue. Within sector s , a standard CES aggregator with elasticity of substitution $\sigma_s > 1$ delivers isoelastic demand for each differentiated variety. Let P_{st} denote the sectoral CES price index and C_{st} total *nominal* expenditure on sector- s varieties in the world economy. For a variety ω produced in region i ,

$$x_{ist}(\omega) = \left(\frac{p_{ist}(\omega)}{P_{st}} \right)^{-\sigma_s} C_{st}.$$

Revenue of variety ω is

$$\begin{aligned} R_{ist}(\omega) &\equiv p_{ist}(\omega) x_{ist}(\omega) \\ &= \mu_s^{1-\sigma_s} C_{st} P_{st}^{\sigma_s-1} c_{ist}^{1-\sigma_s} \\ &= \tilde{\psi}_{is} P_{st}^{\sigma_s-1} (A_{is} E_{ist})^{\sigma_s-1} w_{it}^{\theta_s(1-\sigma_s)} r_{ist}^{(1-\theta_s)(1-\sigma_s)}, \end{aligned} \quad (16)$$

where $\tilde{\psi}_{is} \equiv \mu_s^{1-\sigma_s} C_{st} \chi_s^{1-\sigma_s} > 0$ absorbs sectoral constants. We interpret the economy as a collection of small regions embedded in a large world market. Sectoral prices $\{P_{st}\}_s$ and nominal expenditures $\{C_{st}\}_s$ are determined in the world economy and are taken as exogenous by each region, so firms treat (P_{st}, C_{st}) as given when choosing inputs. Under symmetry of varieties, sectoral revenue is

$$R_{ist} \equiv \int_{\Omega_s} R_{ist}(\omega) d\omega,$$

which inherits the same dependence on $(A_{is} E_{ist}, w_{it}, r_{ist}, P_{st})$ up to a constant that we absorb into $\tilde{\psi}_{is}$.

Labor demand at the variety level. With Cobb–Douglas technology, labor’s share in variable cost is θ_s . Total variable cost per variety equals revenue divided by the markup, so

$$w_{it} l_{ist}(\omega) = \theta_s \frac{R_{ist}(\omega)}{\mu_s} \Rightarrow l_{ist}(\omega) = \underbrace{\frac{\theta_s}{\mu_s}}_{\equiv v_s} w_{it}^{-1} R_{ist}(\omega).$$

Under symmetry, sectoral labor and revenue satisfy $L_{ist} \propto l_{ist}(\omega)$ and $R_{ist} \propto R_{ist}(\omega)$, so we can write

$$L_{ist} = \bar{v}_s w_{it}^{-1} R_{ist}, \quad (17)$$

for some constant $\bar{v}_s > 0$ that absorbs v_s and the mass of varieties.

Eliminating capital prices. Use (14) to express r_{ist} in terms of w_{it} , L_{ist} , and K_{is} :

$$\frac{K_{is}}{L_{ist}} = \iota_s \frac{w_{it}}{r_{ist}} \Rightarrow r_{ist} = \iota_s w_{it} \frac{L_{ist}}{K_{is}}.$$

Substituting into (16),

$$\begin{aligned}
R_{ist} &= \tilde{\psi}_{is} P_{st}^{\sigma_s-1} (A_{is} E_{ist})^{\sigma_s-1} w_{it}^{\theta_s(1-\sigma_s)} \left[l_s w_{it} \frac{L_{ist}}{K_{is}} \right]^{(1-\theta_s)(1-\sigma_s)} \\
&= \tilde{\psi}_{is} l_s^{(1-\theta_s)(1-\sigma_s)} P_{st}^{\sigma_s-1} (A_{is} E_{ist})^{\sigma_s-1} w_{it}^{\theta_s(1-\sigma_s)+(1-\theta_s)(1-\sigma_s)} \\
&\quad \times L_{ist}^{(1-\theta_s)(1-\sigma_s)} K_{is}^{-(1-\theta_s)(1-\sigma_s)} \\
&= \bar{\psi}_{is} P_{st}^{\sigma_s-1} (A_{is} E_{ist})^{\sigma_s-1} w_{it}^{1-\sigma_s} L_{ist}^{(1-\theta_s)(1-\sigma_s)} K_{is}^{-(1-\theta_s)(1-\sigma_s)},
\end{aligned} \tag{18}$$

where $\bar{\psi}_{is} \equiv \tilde{\psi}_{is} l_s^{(1-\theta_s)(1-\sigma_s)} > 0$.

Solving for sectoral labor demand. Plugging (18) into (17) gives

$$\begin{aligned}
L_{ist} &= \bar{v}_s w_{it}^{-1} R_{ist} \\
&= \bar{v}_s \bar{\psi}_{is} P_{st}^{\sigma_s-1} (A_{is} E_{ist})^{\sigma_s-1} w_{it}^{-1} w_{it}^{1-\sigma_s} L_{ist}^{(1-\theta_s)(1-\sigma_s)} K_{is}^{-(1-\theta_s)(1-\sigma_s)} \\
&= \tilde{\xi}_{is} P_{st}^{\sigma_s-1} (A_{is} E_{ist})^{\sigma_s-1} w_{it}^{-\sigma_s} L_{ist}^{(1-\theta_s)(1-\sigma_s)} K_{is}^{-(1-\theta_s)(1-\sigma_s)},
\end{aligned} \tag{19}$$

where $\tilde{\xi}_{is} \equiv \bar{v}_s \bar{\psi}_{is} > 0$.

Rearranging (19), we isolate the power of L_{ist} on the left:

$$\begin{aligned}
L_{ist} &= \tilde{\xi}_{is} P_{st}^{\sigma_s-1} (A_{is} E_{ist})^{\sigma_s-1} w_{it}^{-\sigma_s} L_{ist}^{(1-\theta_s)(1-\sigma_s)} K_{is}^{-(1-\theta_s)(1-\sigma_s)} \\
\Rightarrow L_{ist}^{1-(1-\theta_s)(1-\sigma_s)} &= \tilde{\xi}_{is} P_{st}^{\sigma_s-1} (A_{is} E_{ist})^{\sigma_s-1} w_{it}^{-\sigma_s} K_{is}^{-(1-\theta_s)(1-\sigma_s)}.
\end{aligned} \tag{20}$$

Define

$$\begin{aligned}
\Delta_s &\equiv 1 - (1 - \theta_s)(1 - \sigma_s) \\
&= 1 - [1 - \theta_s - \sigma_s + \theta_s \sigma_s] \\
&= \theta_s + \sigma_s - \theta_s \sigma_s \\
&= 1 + (\sigma_s - 1)(1 - \theta_s).
\end{aligned} \tag{21}$$

Since $\sigma_s > 1$ and $0 < \theta_s < 1$, we have $\Delta_s > 1$, and we set

$$\delta_s \equiv \frac{1}{\Delta_s} = \frac{1}{1 + (\sigma_s - 1)(1 - \theta_s)} \in (0, 1). \tag{22}$$

Raising both sides of (20) to the power δ_s then yields

$$\begin{aligned}
L_{ist} &= \tilde{\xi}_{is}^{\delta_s} P_{st}^{(\sigma_s-1)\delta_s} (A_{is} E_{ist})^{(\sigma_s-1)\delta_s} w_{it}^{-\sigma_s \delta_s} K_{is}^{-(1-\theta_s)(1-\sigma_s)\delta_s} \\
&= \tilde{\xi}_{is}^{\delta_s} (A_{is} E_{ist} P_{st})^{(\sigma_s-1)\delta_s} w_{it}^{-\sigma_s \delta_s} K_{is}^{(1-\theta_s)(\sigma_s-1)\delta_s},
\end{aligned} \tag{23}$$

where we used $1 - \sigma_s = -(\sigma_s - 1)$ to rewrite the exponent on K_{is} . Defining

$$\psi_{is} \equiv \tilde{\zeta}_{is}^{\delta_s} > 0, \quad (24)$$

we obtain the sectoral labor demand equation

$$L_{ist} = \psi_{is} w_{it}^{-\sigma_s \delta_s} (A_{is} E_{ist} P_{st})^{(\sigma_s - 1) \delta_s} K_{is}^{(1 - \theta_s)(\sigma_s - 1) \delta_s}, \quad (25)$$

which coincides with (3).

A.2 Proof of Proposition 1

Sector- q Employment and Output Responses Fix a tradable sector $q \in \mathcal{S}_T$. Substituting the wage-price elasticity (9) into (4) with $s = q$ gives the exact log response of sector- q employment to its own price:

$$\begin{aligned} \frac{\partial \ln L_{iqt}}{\partial \ln P_{qt}} &= (\sigma_q - 1) \delta_q - \sigma_q \delta_q \cdot \frac{1}{\lambda_i} \sum_{s \in \mathcal{S}_T} \ell_{is,0}^* (\sigma_s - 1) \delta_s \rho_{qs} \\ &= (\sigma_q - 1) \delta_q \left[1 - \frac{\sigma_q}{\sigma_q - 1} \cdot \frac{1}{\lambda_i} \sum_{s \in \mathcal{S}_T} \ell_{is,0}^* (\sigma_s - 1) \delta_s \rho_{qs} \right]. \end{aligned} \quad (26)$$

Sector- q output is, from (1),

$$X_{iqt} = A_{iq} E_{iqt} K_{iq}^{1 - \theta_q} L_{iqt}^{\theta_q}. \quad (27)$$

Taking differentials,

$$d \ln X_{iqt} = d \ln A_{iq} + d \ln E_{iqt} + (1 - \theta_q) d \ln K_{iq} + \theta_q d \ln L_{iqt}. \quad (28)$$

Evaluating (28) at a baseline where A_{iq}, K_{iq} are fixed (short-run) and focusing on the semi-elasticity w.r.t. the exogenous price P_{qt} ,

$$\frac{\partial \ln X_{iqt}}{\partial \ln P_{qt}} = \theta_q \cdot \frac{\partial \ln L_{iqt}}{\partial \ln P_{qt}} = (\sigma_q - 1) \theta_q \delta_q \left[1 - \frac{\sigma_q}{\sigma_q - 1} \cdot \frac{1}{\lambda_i} \sum_{s \in \mathcal{S}_T} \ell_{is,0}^* (\sigma_s - 1) \delta_s \rho_{qs} \right]. \quad (29)$$

Pre-period equilibrium objects. We'll use the semi elasticity to derive a first order expansion of the equilibrium output. For that, we treat baseline objects as fixed and observed in the pre-period:

$$L_{iq,0} \equiv L_{iq}^*, \quad X_{iq,0} \equiv X_{iq}^*,$$

$$X_{iq,0} = A_{iq} \bar{E}_{iq} K_{iq}^{1-\theta_q} (L_{iq,0})^{\theta_q}.$$

and use this notation below to emphasize they are time-invariant conditioning variables.

Thus, the *levels* semi-elasticity (evaluate at baseline X_{iq}^*) is

$$\tilde{\kappa}_{iq} \equiv \left. \frac{\partial X_{iqt}}{\partial \ln P_{qt}} \right|_0 = X_{iq,0} \cdot (\sigma_q - 1) \theta_q \delta_q \left[1 - \frac{\sigma_q}{\sigma_q - 1} \cdot \frac{1}{\lambda_i} \sum_{s \in \mathcal{S}_T} \ell_{is,0}^* (\sigma_s - 1) \delta_s \rho_{qs} \right]. \quad (30)$$

For expositional purposes, also define the scaled elasticity:

$$\kappa_{iq} \equiv \frac{\tilde{\kappa}_{iq}}{A_{iq}}, \quad (31)$$

We also need to calculate the semi elasticity of the random shock. Let e_{iqt} be a centered (log) idiosyncratic efficiency shock defined by

$$e_{iqt} \equiv \ln E_{iqt} - \ln \bar{E}_{iq},$$

so that

$$E_{iqt} = \bar{E}_{iq} \exp(e_{iqt}), \quad e_{iqt} = 0 \text{ at the baseline } E_{iqt} = \bar{E}_{iq}.$$

Since $\partial \ln E_{iqt} / \partial e_{iqt} = 1$, the semi elasticities below are invariant to this normalization. From sectoral labor demand in (4) applied to sector q :

$$d \ln L_{iqt} = (\sigma_q - 1) \delta_q d \ln P_{qt} + (\sigma_q - 1) \delta_q d \ln E_{iqt} - \sigma_q \delta_q d \ln w_{it},$$

and holding the common price fixed for this derivative ($d \ln P_{qt} = 0$) while $d \ln E_{iqt} = de_{iqt}$, we get

$$\frac{\partial \ln L_{iqt}}{\partial e_{iqt}} = (\sigma_q - 1) \delta_q - \sigma_q \delta_q \frac{\partial \ln w_{it}}{\partial e_{iqt}}.$$

The market clearing condition in (6) implies the wage semi-elasticity with respect to a *local* shock in (i, q) :

$$\frac{\partial \ln w_{it}}{\partial e_{iqt}} = \frac{1}{\lambda_i} \ell_{iq,0}^* (\sigma_q - 1) \delta_q, \quad \lambda_i \equiv \phi_i + \sum_{s \in \mathcal{S}} \ell_{is,0}^* \sigma_s \delta_s,$$

where $\ell_{is,0}^*$ are baseline employment shares. Hence

$$\frac{\partial \ln L_{iqt}}{\partial e_{iqt}} = (\sigma_q - 1) \delta_q \left[1 - \frac{\sigma_q \delta_q}{\lambda_i} \ell_{iq,0}^* \right].$$

Using production $X_{iqt} = A_{iq} E_{iqt} K_{iq}^{1-\theta_q} L_{iqt}^{\theta_q}$ gives

$$\frac{\partial \ln X_{iqt}}{\partial e_{iqt}} = \underbrace{\frac{\partial \ln E_{iqt}}{\partial e_{iqt}}}_{=1} + \theta_q \frac{\partial \ln L_{iqt}}{\partial e_{iqt}} = 1 + (\sigma_q - 1) \theta_q \delta_q \left[1 - \frac{\sigma_q \delta_q}{\lambda_i} \ell_{iq,0}^* \right].$$

Evaluating at the baseline and multiplying by $X_{iq,0}$ yields

$$\kappa_{iq}^E = \left. \frac{\partial X_{iqt}}{\partial e_{iqt}} \right|_0 = X_{iq,0} \left\{ 1 + (\sigma_q - 1) \theta_q \delta_q \left[1 - \frac{\sigma_q \delta_q}{\lambda_i} \ell_{iq,0}^* \right] \right\}.$$

First-Order Linear Approximation in Levels From Taylor's Theorem, a first-order approximation of (27) around the baseline yields the empirical first stage

$$X_{iqt} \approx X_{iq,0} + \left. \frac{\partial X_{iqt}}{\partial p_{qt}} \right|_0 p_{qt} + \left. \frac{\partial X_{iqt}}{\partial e_{iqt}} \right|_0 e_{iqt}$$

Then, using (30):

$$X_{iqt} \approx \alpha_{iq} + \tilde{\kappa}_{iq} p_{qt} + \kappa_{iq}^E e_{iqt} + \varepsilon_{iqt}, \quad (32)$$

Define κ_{iq} as in (11) so that $\tilde{\kappa}_{iq} = A_{iq} \kappa_{iq}$. Substituting into (32) and collecting all shocks orthogonal to p_{qt} (including $\kappa_{iq}^E e_{iqt}$) into ε_{iqt} yields

$$X_{iqt} \approx \alpha_{iq} + [A_{iq} \kappa_{iq}] p_{qt} + \varepsilon_{iqt}, \quad (33)$$

which is exactly (10), concluding the proof.

B Other Proofs

B.1 Proof of Proposition 2

$$\begin{aligned}
\beta_q^{2SLS} &= \frac{\sum_{i=1}^N \sum_{t=1}^T \mathbb{E}[Z_{iqt} Y_{it} | \mathcal{F}_0]}{\sum_{i=1}^N \sum_{t=1}^T \mathbb{E}[Z_{iqt} X_{iqt} | \mathcal{F}_0]} = \frac{\sum_{i=1}^N \sum_{t=1}^T \mathbb{E}[(A_{iq} p_{qt}) Y_{it} | \mathcal{F}_0]}{\sum_{i=1}^N \sum_{t=1}^T \mathbb{E}[(A_{iq} p_{qt}) X_{iqt} | \mathcal{F}_0]} \\
&= \frac{\sum_{i=1}^N \sum_{t=1}^T \mathbb{E}[(A_{iq} p_{qt}) (\sum_{s=1}^S \beta_{is} X_{ist} + \eta_{it}) | \mathcal{F}_0]}{\sum_{i=1}^N \sum_{t=1}^T \mathbb{E}[(A_{iq} p_{qt}) (\alpha_{iq} + \kappa_{iq} [A_{iq} p_{qt}] + \varepsilon_{iqt}) | \mathcal{F}_0]} \\
&= \frac{\sum_{i=1}^N \sum_{t=1}^T \mathbb{E}[(A_{iq} p_{qt}) (\sum_{s=1}^S \beta_{is} \{\alpha_{is} + \kappa_{is} [A_{is} p_{st}] + \varepsilon_{ist}\} + \eta_{it}) | \mathcal{F}_0]}{\sum_{i=1}^N \sum_{t=1}^T \mathbb{E}[(A_{iq} p_{qt}) (\alpha_{iq} + \kappa_{iq} [A_{iq} p_{qt}] + \varepsilon_{iqt}) | \mathcal{F}_0]} \\
&= \frac{\sum_{i=1}^N \sum_{t=1}^T \mathbb{E}[\kappa_{iq} (A_{iq} p_{qt})^2 \beta_{iq} | \mathcal{F}_0]}{\sum_{i=1}^N \sum_{t=1}^T \mathbb{E}[\kappa_{iq} (A_{iq} p_{qt})^2 | \mathcal{F}_0]} + \frac{\sum_{i=1}^N \sum_{t=1}^T \mathbb{E}[\sum_{s \neq q} \kappa_{is} A_{is} A_{iq} p_{st} p_{qt} \beta_{is} | \mathcal{F}_0]}{\sum_{i=1}^N \sum_{t=1}^T \mathbb{E}[\kappa_{iq} (A_{iq} p_{qt})^2 | \mathcal{F}_0]} \\
&= \frac{\sum_{i=1}^N \sum_{t=1}^T \kappa_{iq} (A_{iq})^2 \mathbb{V}(p_{qt} | \mathcal{F}_0) \beta_{iq}}{\sum_{i=1}^N \sum_{t=1}^T \kappa_{iq} (A_{iq})^2 \mathbb{V}(p_{qt} | \mathcal{F}_0)} + \frac{\sum_{i=1}^N \sum_{t=1}^T \sum_{s \neq q} \kappa_{is} A_{is} A_{iq} \text{Cov}(p_{st}, p_{qt} | \mathcal{F}_0) \beta_{is}}{\sum_{i=1}^N \sum_{t=1}^T \kappa_{iq} (A_{iq})^2 \mathbb{V}(p_{qt} | \mathcal{F}_0)} \\
&= \frac{\sum_{i=1}^N \sum_{t=1}^T \pi_{iqt} \beta_{iq}}{\sum_{i=1}^N \sum_{t=1}^T \pi_{iqt}} + \frac{\sum_{i=1}^N \sum_{t=1}^T \{\sum_{s \neq q} \pi_{isqt} \beta_{is}\}}{\sum_{i=1}^N \sum_{t=1}^T \pi_{iqt}}
\end{aligned}$$

B.2 Proof of Proposition 3

To obtain this result, we write the numerator of the reduced form as

$$\begin{aligned}
\sum_i \sum_t \mathbb{E}[Z_{iqt} Y_{it} | \tilde{\mathcal{F}}_0] &= \sum_i \sum_t \mathbb{E}[(A_{iq} p_{qt}) Y_{it} | \tilde{\mathcal{F}}_0] = \sum_i \sum_t \mathbb{E}\left[(A_{iq} p_{qt}) \left(\eta_{it} + \sum_{s=1}^S X_{ist} \beta_{is}\right) | \tilde{\mathcal{F}}_0\right] \\
&= \sum_i \sum_t \mathbb{E}\left[(A_{iq} p_{qt}) \left(\eta_{it} + \sum_{s=1}^S \{\alpha_{is} + \kappa_{is} (A_{is} p_{st}) + \varepsilon_{ist}\} \beta_{is}\right) | \tilde{\mathcal{F}}_0\right] \\
&= \sum_i \sum_t \tilde{\pi}_{iqt} \beta_{iq} + \sum_i \sum_t \sum_{s \neq q} \tilde{\pi}_{isqt} \beta_{is} + \sum_i \sum_t \sum_{s \neq q} \sum_{s' \neq s} A_{iq} \gamma_{iss'} \text{Cov}(p_{qt}, p_{s't} | \tilde{\mathcal{F}}_0) \beta_{is}
\end{aligned}$$

where $\tilde{\pi}_{iqt} = \kappa_{iq} A_{iq}^2 \mathbb{V}(p_{qt} | \tilde{\mathcal{F}}_0)$ and $\tilde{\pi}_{isqt} = \kappa_{is} A_{is} A_{iq} \text{Cov}(p_{qt}, p_{st} | \tilde{\mathcal{F}}_0)$.

B.3 Proof of Proposition 4

The proof is similar to the one from Proposition 2, the main difference comes from the decomposition of the denominator. Note that

$$\sum_{i=1}^N \mathbb{E} \left[(\Delta Z_{iq})^2 \mid \mathcal{F}_0 \right] = \sum_{i=1}^N \mathbb{E} \left[(A_{iq} \Delta p_q)^2 \mid \mathcal{F}_0 \right] = \sum_{i=1}^N A_{iq}^2 \mathbb{V}(\Delta p_q \mid \mathcal{F}_0)$$

which concludes the proof.

C Asymptotics

In this section we provide additional details for the asymptotic results in Section 4.

C.1 Additional Assumptions

Assumption 5 (Regularity Conditions for Consistency).

1. *The support of β_{is} is bounded.*
2. *The second moments of η_{it} exist and are uniformly bounded over i, t .*
3. *The second moments of α_{is} are uniformly bounded over i, t .*
4. *The second moments of A_{is} are uniformly bounded over i, t .*
5. *The second moments of κ_{is} are uniformly bounded over i, t .*
6. *$\frac{1}{NT} \sum_i \sum_t \varepsilon_{ist} \rightarrow 0$, and $\frac{1}{(NT)^2} \sum_i \sum_t \varepsilon_{ist}^2 \rightarrow 0$.*
7. *$\frac{1}{NT} \sum_i \sum_t \kappa_{is} A_{is}^2 \mathbb{V}(p_{qs} \mid \mathcal{F}_0)$ converges in probability to a strictly positive non-random limit.*
8. *For some $\nu > 0$, $\mathbb{E} \left[|p_{st}|^{2+\nu} \mid \mathcal{F}_0 \right]$ exists and is uniformly bounded.*

Assumption 6 (Regularity Conditions for Normality).

1. *The fourth moments of η_{it} exist and are uniformly bounded over i, t .*
2. *The fourth moments of α_{is} are uniformly bounded over i, t .*
3. *The fourth moments of A_{is} are uniformly bounded over i, t .*
4. *The fourth moments of κ_{is} are uniformly bounded over i, t .*
5. *For some $\nu > 0$, $\mathbb{E} \left[|p_{st}|^{4+\nu} \mid \mathcal{F}_0 \right]$ exists and is uniformly bounded.*

C.2 Proof of Theorem 1

Consistency

Under Assumptions 1 and 5. Assuming that the price shocks (p_{q1}, \dots, p_{qT}) are independent conditional on \mathcal{F}_0 , we can show that

$$\frac{1}{NT} \sum_i \sum_t Z_{iqt} X_{iqt} = \frac{1}{NT} \sum_i \sum_t \kappa_{iq} A_{iq} \sigma_{qt}^2 + o_p(1)$$

Begin by writing

$$\begin{aligned} \frac{1}{NT} \sum_i \sum_t Z_{iqt} X_{iqt} &= \frac{1}{NT} \sum_i \sum_t (A_{iq} p_{qt}) (\alpha_{iq} + \kappa(A_{iq} p_{qt}) + \varepsilon_{iqt}) \\ &= \frac{1}{NT} \sum_i \sum_t (A_{iq} p_{qt}) \alpha_{iq} + \frac{1}{NT} \sum_i \sum_t \kappa_{iq} A_{iq}^2 (p_{qt}^2 - \sigma_{qt}^2) + \frac{1}{NT} \sum_i \sum_t A_{iq} p_{qt} \varepsilon_{iqt} \\ &\quad + \frac{1}{NT} \sum_i \sum_t \kappa_{iq} A_{iq} \sigma_{qt}^2 + o_p(1) \end{aligned}$$

The proof follows by showing that the first three terms of the second equality are $o_p(1)$.

The first term has mean zero conditional on \mathcal{F}_0 , and variance conditional on \mathcal{A}_q equal to

$$\begin{aligned} \mathbb{V}_{\mathcal{A}_q} \left(\frac{1}{NT} \sum_i \sum_t (A_{iq} p_{qt}) \alpha_{iq} \right) &= \mathbb{V}_{\mathcal{A}_q} \left(\frac{1}{NT} \sum_i A_{iq} \alpha_{iq} \sum_t p_{qt} \right) \\ &= \frac{1}{(NT)^2} \mathbb{E} \left[\sum_t \sigma_{qt}^2 \sum_i (A_{iq} \alpha_{iq})^2 \right] \leq \frac{1}{(NT)^2} \mathbb{E} \left[\sum_i (A_{iq} \alpha_{iq})^2 \right] \rightarrow 0 \end{aligned}$$

Convergence to zero follows from Assumption 5.1 and 5.2.

For the second term, note that

$$\begin{aligned} \mathbb{E} \left[(NT)^{-1} \left| \sum_t (p_{qt}^2 - \sigma_{qt}^2) \sum_i \kappa_{iq} A_{iq}^2 \right|^{1+\nu/2} \middle| \mathcal{F}_0 \right] &\leq \frac{2}{(NT)^{1+\nu/2}} \sum_i (\kappa_{iq} A_{iq}^2)^{1+\nu/2} \mathbb{E} \left[\sum_t |p_{qt}^2 - \sigma_{qt}^2|^{1+\nu/2} \middle| \mathcal{F}_0 \right] \\ &\leq \frac{1}{(NT)^{1+\nu/2}} \sum_i (\kappa_{iq} A_{iq}^2)^{1+\nu/2} \rightarrow 0 \end{aligned}$$

Where the first inequality follows from Assumption 5.6 and the second follows from Assumption 5.2 and 5.3. The third term has mean zero conditional on \mathcal{F}_0 and variance

$$\mathbb{V}_{\mathcal{A}_q} \left(\frac{1}{NT} \sum_i \sum_t A_{iq} p_{qt} \varepsilon_{iqt} \right) = \frac{1}{(NT)^2} \mathbb{E}_{\mathcal{A}_q} \left[\sum_i \sum_t \sigma_{qt}^2 A_{iq}^2 \varepsilon_{iqt}^2 \right] \leq \frac{1}{(NT)^2} \varepsilon_{iqt}^2 \rightarrow 0$$

Similarly, we show that

$$\frac{1}{NT} \sum_i \sum_t Z_{iq} Y_{it} = \frac{1}{NT} \sum_i \sum_t \kappa_{iq} A_{iq}^2 \sigma_{qt}^2 \beta_{iq} + o_p(1)$$

First, note that

$$\begin{aligned} \frac{1}{NT} \sum_i \sum_t Z_{iq} Y_{it} &= \frac{1}{NT} \sum_i \sum_t (A_{iq} p_{qt}) (\eta_{it} + \sum_{s=1}^s X_{ist} \beta_{is}) \\ &= \frac{1}{NT} \sum_i \sum_t A_{iq} p_{qt} \eta_{it} + \frac{1}{NT} \sum_i \sum_t (A_{iq} p_{qt}) X_{iq} \beta_{iq} + \frac{1}{NT} \sum_i \sum_t \sum_{s \neq q} (A_{iq} p_{qt}) X_{ist} \beta_{is} \end{aligned}$$

The first term has mean zero conditional on \mathcal{F}_0 and variance equal to

$$\begin{aligned} \mathbb{V}_{\mathcal{A}_q} \left(\frac{1}{NT} \sum_i \sum_t A_{iq} p_{qt} \eta_{it} \right) &= \frac{1}{(NT)^2} \mathbb{E}_{\mathcal{A}_q} \left[\sum_i \sum_t \sigma_{qt}^2 (A_{iq} \eta_{it})^2 \right] \\ &\leq \frac{1}{(NT)^2} \mathbb{E}_{\mathcal{A}_q} \left[\sum_i \sum_t (A_{iq} \eta_{it})^2 \right] \rightarrow 0 \end{aligned}$$

Asymptotic Distribution

Write the regression error associated to β_q as

$$\begin{aligned} u_{it} &= Y_{it} - X_{iq} \beta_q = \eta_{it} + X_{iq} (\beta_{iq} - \beta_q) + \sum_{s \neq q} X_{ist} \beta_{is} \\ &= \eta_{it} + \{ \alpha_{iq} + \kappa_{iq} (A_{iq} p_{qt}) + \varepsilon_{iq} \} (\beta_{iq} - \beta_q) + \sum_{s \neq q} X_{ist} \beta_{is} \\ &= \eta_{it} + (\alpha_{iq} + \varepsilon_{iq}) (\beta_{iq} - \beta_q) + \kappa_{iq} (A_{iq} p_{qt}) (\beta_{iq} - \beta_q) + \sum_{s \neq q} X_{ist} \beta_{is} \end{aligned}$$

From the consistency result, it follows that

$$\begin{aligned}
\hat{\beta}_q - \beta_q &= (1 + o_p(1)) \left(\frac{1}{NT} \sum_i \sum_t \kappa_{iq} A_{iq}^2 \sigma_{qt}^2 \right)^{-1} \times \left(\frac{1}{NT} \sum_i \sum_t (A_{iq} p_{qt}) Y_{it} - \kappa_{iq} A_{iq}^2 \sigma_{qt}^2 \beta_{iq} \right) \\
&= (1 + o_p(1)) \left(\frac{1}{NT} \sum_i \sum_t \kappa_{iq} A_{iq}^2 \sigma_{qt}^2 \right)^{-1} \times \left(\frac{1}{NT} \sum_i \sum_t (A_{iq} p_{qt}) (\eta_{it} + \sum_{s=1}^S X_{ist} \beta_{is}) - \kappa_{iq} A_{iq}^2 \sigma_{qt}^2 \beta_{iq} \right) \\
&= (1 + o_p(1)) \left(\frac{1}{NT} \sum_i \sum_t \kappa_{iq} A_{iq}^2 \sigma_{qt}^2 \right)^{-1} \times \left(\frac{1}{NT} \sum_i \sum_t (A_{iq} p_{qt}) u_{it} + (A_{iq} p_{qt}) X_{iq} \beta_q - \kappa_{iq} A_{iq}^2 \sigma_{qt}^2 \beta_{iq} \right) \\
&= (1 + o_p(1)) \left(\frac{1}{NT} \sum_i \sum_t \kappa_{iq} A_{iq}^2 \sigma_{qt}^2 \right)^{-1} \\
&\quad \times \left(\frac{1}{NT} \sum_i \sum_t (A_{iq} p_{qt}) u_{it} + (A_{iq} p_{qt}) (\alpha_{iq} + \varepsilon_{iq}) \beta_q + \kappa_{iq} A_{iq}^2 p_{qt}^2 \beta_q - \kappa_{iq} A_{iq}^2 \sigma_{qt}^2 \beta_{iq} \right) \\
&= (1 + o_p(1)) \left(\frac{1}{NT} \sum_i \sum_t \kappa_{iq} A_{iq}^2 \sigma_{qt}^2 \right)^{-1} \\
&\quad \times \left(\frac{1}{NT} \sum_i \sum_t (A_{iq} p_{qt}) u_{it} + (A_{iq} p_{qt}) (\alpha_{iq} + \varepsilon_{iq}) \beta_q + \kappa_{iq} A_{iq}^2 p_{qt}^2 \beta_q - \kappa_{iq} A_{iq}^2 \sigma_{qt}^2 (\beta_{iq} - \beta_q) - \kappa_{iq} A_{iq}^2 \sigma_{qt}^2 \beta_q \right)
\end{aligned}$$

Note that $\frac{1}{NT} \sum_i \sum_t \kappa_{iq} A_{iq}^2 \sigma_{qt}^2 (\beta_{iq} - \beta_q) = 0$ and hence, it follows that

$$\begin{aligned}
\hat{\beta}_q - \beta_q &= (1 + o_p(1)) \left(\frac{1}{NT} \sum_i \sum_t \kappa_{iq} A_{iq}^2 \sigma_{qt}^2 \right)^{-1} \\
&\quad \times \left(\frac{1}{NT} \sum_i \sum_t (A_{iq} p_{qt}) u_{it} + (A_{iq} p_{qt}) (\alpha_{iq} + \varepsilon_{iq}) \beta_q + \kappa_{iq} A_{iq}^2 (p_{qt}^2 - \sigma_{qt}^2) \beta_q \right)
\end{aligned}$$

And therefore,

$$\begin{aligned}
\sqrt{NT}(\hat{\beta}_q - \beta_q) &= (1 + o_p(1)) \left(\frac{1}{NT} \sum_i \sum_t \kappa_{iq} A_{iq}^2 \sigma_{qt}^2 \right)^{-1} \\
&\quad \times \left(\frac{1}{\sqrt{NT}} \sum_i \sum_t (A_{iq} p_{qt}) u_{it} + (A_{iq} p_{qt}) (\alpha_{iq} + \varepsilon_{iq}) \beta_q + \kappa_{iq} A_{iq}^2 (p_{qt}^2 - \sigma_{qt}^2) \beta_q \right) \\
&= (1 + o_p(1)) \left(\frac{1}{NT} \sum_i \sum_t \kappa_{iq} A_{iq}^2 \sigma_{qt}^2 \right)^{-1} \times \left(\frac{1}{\sqrt{NT}} \sum_i \sum_t (A_{iq} p_{qt}) u_{it} + o_p(1) \right)
\end{aligned}$$

In order to prove the result, it suffices to show that $\frac{1}{\sqrt{NT}} \sum_i \sum_t (A_{iq} p_{qt}) u_{it} \xrightarrow{d} N(0, V_{NT})$. We do so by showing that the Lindeberg conditions hold, and thus the convergence in distribution follows from the Lindeberg CLT.

We decompose $\frac{1}{\sqrt{NT}} \sum_i \sum_t (A_{iq} p_{qt}) u_{it}$ as

$$\begin{aligned}
\frac{1}{\sqrt{NT}} \sum_i \sum_t (A_{iq} p_{qt}) u_{it} &= \frac{1}{\sqrt{NT}} \sum_i \sum_t (A_{iq} p_{qt}) \eta_{it} + \frac{1}{\sqrt{NT}} \sum_i \sum_t (A_{iq} p_{qt}) (\alpha_{iq} + \varepsilon_{ist}) (\beta_{iq} - \beta_q) \\
&+ \frac{1}{\sqrt{NT}} \sum_i \sum_t \kappa_{iq} A_{iq}^2 p_{qt}^2 (\beta_{iq} - \beta_q) + \frac{1}{\sqrt{NT}} \sum_i \sum_t \sum_{s \neq q} (A_{iq} p_{qt}) (\alpha_{is} + \varepsilon_{ist}) \beta_{is} \\
&+ \frac{1}{\sqrt{NT}} \sum_i \sum_t \sum_{s \neq q} \kappa_{is} A_{is} A_{is} p_{st} p_{qt} \beta_{is} \\
&= \frac{1}{\sqrt{NT}} \sum_i \sum_t (A_{iq} p_{qt}) \eta_{it} + \frac{1}{\sqrt{NT}} \sum_i \sum_t (A_{iq} p_{qt}) (\alpha_{iq} + \varepsilon_{ist}) (\beta_{iq} - \beta_q) \\
&+ \frac{1}{\sqrt{NT}} \sum_i \sum_t \kappa_{iq} A_{iq}^2 (p_{qt}^2 - \sigma_{qt}^2) (\beta_{iq} - \beta_q) + \frac{1}{\sqrt{NT}} \sum_i \sum_t \sum_{s \neq q} (A_{iq} p_{qt}) (\alpha_{is} + \varepsilon_{ist}) \beta_{is} \\
&+ \frac{1}{\sqrt{NT}} \sum_i \sum_t \sum_{s \neq q} \kappa_{is} A_{is} A_{is} p_{st} p_{qt} \beta_{is} = \frac{1}{NT} \sum_i \sum_t y_{it}^{(1)} + y_{it}^{(2)} + y_{it}^{(3)} + y_{it}^{(4)} + y_{it}^{(5)} \\
&= \frac{1}{NT} \sum_i \sum_t y_{it}
\end{aligned}$$

To show that the Lindeberg condition holds, it suffices to show that for some $\nu > 0$,

$$\frac{1}{(NT)^{1+\nu/4}} \sum_i \sum_t \mathbb{E}_{\mathcal{A}_q} \left[y_{it}^{2+\nu/2} \right] \rightarrow 0.$$

The Lindeberg condition holds if

$$\begin{aligned}
\frac{1}{(NT)^2} \sum_i \sum_t \mathbb{E}_{\mathcal{A}_q} \left[\left(y_{it}^{(1)} \right)^4 \right] &\rightarrow 0, \quad \frac{1}{(NT)^2} \sum_i \sum_t \mathbb{E}_{\mathcal{A}_q} \left[\left(y_{it}^{(2)} \right)^4 \right] \rightarrow 0, \\
\frac{1}{(NT)^{1+\nu/4}} \sum_i \sum_t \mathbb{E}_{\mathcal{A}_q} \left[\left(y_{it}^{(3)} \right)^{2+\nu/2} \right] &\rightarrow 0, \quad \frac{1}{(NT)^2} \sum_i \sum_t \mathbb{E}_{\mathcal{A}_q} \left[\left(y_{it}^{(4)} \right)^4 \right] \rightarrow 0, \\
\frac{1}{(NT)^2} \sum_i \sum_t \mathbb{E}_{\mathcal{A}_q} \left[\left(y_{it}^{(5)} \right)^4 \right] &\rightarrow 0
\end{aligned}$$

The first follows from Assumptions 6.1 and 6.5. The second follows from Assumptions 5.1, 5.6, 6.1 and 6.5. The third follows from Assumptions 6.3, 6.4 and 6.5. The fourth follows from Assumptions 5.1, 5.6, 6.2 and 6.5. The fifth follows from Assumptions 5.1, 6.2 and 6.6. Therefore, it follows from the Lindeberg CLT that

$$\frac{1}{\sqrt{NT}} \sum_i \sum_t (A_{iq} p_{qt}) u_{it} \xrightarrow{d} N(0, V_{NT})$$

which concludes the proof.

Inference under Heterogeneous Effects

For valid, yet conservative, inference under heterogeneous effects, it must be the case that

$$\frac{\sum_t p_{qt}^2 \hat{R}_t^2}{NT} \geq \mathcal{V}_{NT} + o_p(1)$$

In order to assess the conditions under which the inequality holds, note that

$$\begin{aligned} \mathcal{V}_{NT} &= \frac{\mathbb{V}(\sum_t p_{qt} R_t | \mathcal{F}_0)}{NT} = \frac{\sum_t \mathbb{E}[p_{qt}^2 R_t^2 | \mathcal{F}_0]}{NT} + \frac{\sum_{t \neq t'} \text{Cov}(p_{qt} R_t, p_{qt'} R_{t'} | \mathcal{F}_0)}{NT} - \frac{\sum_t \mathbb{E}[p_{qt} R_t | \mathcal{F}_0]^2}{NT} \\ &= D_1 + D_2 - D_3 \end{aligned}$$

Note that

$$\begin{aligned} D_2 &= \frac{\sum_{t \neq t'} \sigma_{qt}^2 \sigma_{qt'}^2 \sum_{i,j} \kappa_{iq}^2 A_{iq}^4 (\beta_{iq} - \beta) \kappa_{jq}^2 A_{jq}^4 (\beta_{jq} - \beta_q)}{NT}, \\ D_3 &= \frac{\sum_t \left(\sum_i \kappa_{iq} A_{iq}^2 \sigma_{qt}^2 (\beta_{iq} - \beta_q) \right)^2}{NT} \end{aligned}$$

And thus, in the case of homogeneous effects $D_2 = D_3 = 0$, the standard error estimator consistently estimates D_1 . In order for inference to be valid, yet conservative, under treatment effect heterogeneity, it must be the case that $D_2 = o_p(1)$.

D Empirical application details

All data used in the empirical application come from publicly available sources: DATASUS¹³ (mortality microdata from SIM, considering ICD-10 codes X86-Y09), Serviço Geológico do Brasil¹⁴ (mineral deposit locations), IBGE¹⁵ (municipal population), World Bank¹⁶ (international gold prices), and MapBiomas¹⁷ (machine-learning based satellite mining data).

There are 591 municipalities in the panel, observed between 2000 and 2022. Table 6 reports summary statistics for the variables used. The average homicide rate is about 22 per 100,000

¹³<https://datasus.saude.gov.br/>

¹⁴<https://geoportal.cprm.gov.br/geosgb/>

¹⁵<https://www.ibge.gov.br/estatisticas/economicas/contas-nacionais.html>

¹⁶<https://www.worldbank.org/en/research/commodity-markets>

¹⁷https://mapbiomas.org/en/download-dos-atbds?cama_set_language=en

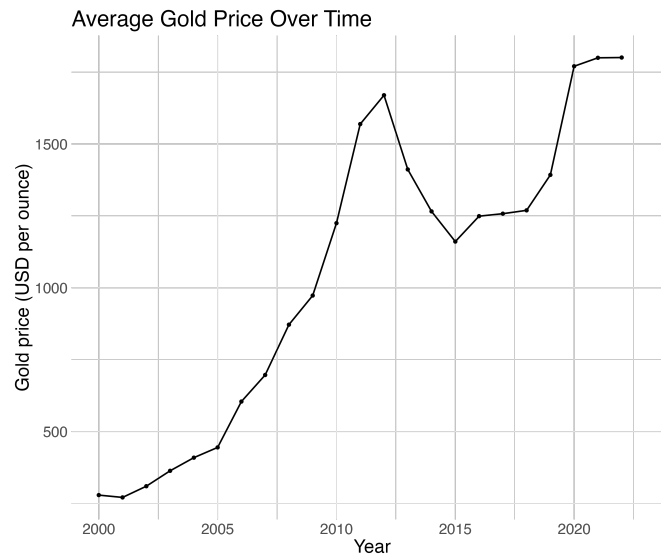
inhabitants, with substantial cross-sectional variation and a long right tail. Mining activity is highly concentrated: mean mining area is roughly 1 km², with most municipality-years at zero and a small number of locations with very large mines. The count of deposits and the price-exposure variable Z_{it} display similarly skewed distributions. Figures 1 and 2 show the evolution of gold prices, mining area, and homicide rates over time, illustrating the rise in gold prices was accompanied by an increase in mining activity (which motivates the first stage) and in homicides in municipalities with more gold deposits (which motivates the empirical analysis).

Table 6: Summary Statistics


Statistic	N	Mean	St. Dev.	Min	Max
Homicide rate (per 100k)	13,593	22.38	27.47	0.00	658.32
Mining area (km ²)	13,593	0.98	7.69	0.00	412.22
Gold deposits (count)	13,593	3.92	21.29	0	459
Gold price (USD/oz)	13,593	1,046.22	515.99	270.99	1,800.60
Price exposure Z	13,593	26.61	145.10	0.00	3,440.61

Note: This table reports summary statistics for the variables used in Tables 4 and 5. Statistics are calculated over the full panel of 591 municipalities and 23 years.

Figure 1: Gold Prices and Mining Area in the Brazilian Amazon



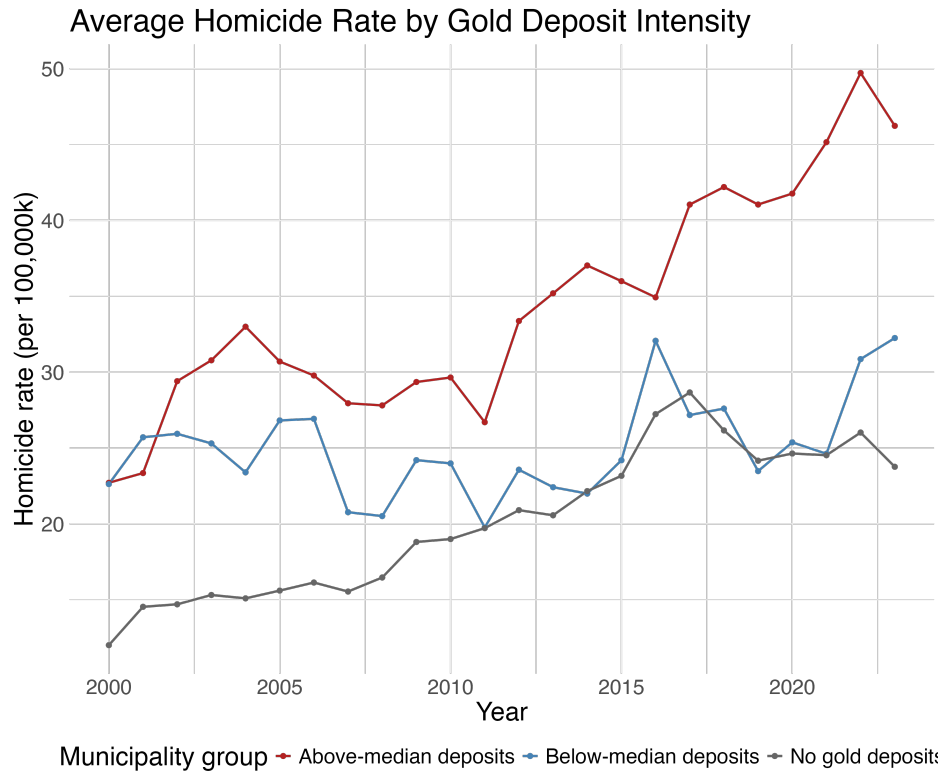
(a) Average gold price over time.



(b) Average gold mining area over time.

Note: Panel (a) plots the average international gold price (USD per ounce) by year. Panel (b) plots the average gold mining area (km^2) across municipalities by year. Both series are computed from the same municipality-year panel used in the empirical application.

Figure 2: Homicides and Gold Deposit Intensity in the Brazilian Amazon



Note: The figure plots average homicide rates per 100,000 inhabitants by municipality group, where groups are defined by the number of gold deposits (no deposits, below-median deposits, and above-median deposits among municipalities with positive deposits). Series are computed from the municipality-year panel used in the empirical application.