Domestic Trade Frictions: Crime*

Alfonso Cebreros† Jordán Mosqueda‡ Daniel Ramos-Menchelli§

Banco de México UC San Diego Johns Hopkins

May 2026

Abstract

Crime is a pervasive constraint on firms in many developing countries. Its economic reach extends beyond direct losses as crime could distort agents' decisions, economic activity, and prices. This paper studies how crime affects domestic trade and the economy in the context of Mexico, where the presence of organized crime is widespread. We develop a spatial trade model in which transport intermediaries choose routes subject to crime-related shipping costs, and combine it with reduced-form analysis that exploits confidential transportation-firm microdata, origin-destination wholesale agricultural prices, crop production data, and administrative crime records. We construct route-level measures of crime exposure and show that higher exposure raises transport firms' operating costs, with the clearest evidence for extortion. Wholesale price regressions provide suggestive, though noisier, evidence that route exposure is passed through to agricultural prices. We embed these estimates in the model and study a counterfactual that removes crime frictions across space. We find that removing this friction lowers mean bilateral trade costs by 2.9 percent, lowers destination tradable price indices by 6.5 percent, and raises welfare by 1.4 percent. These results suggest that crime is a substantial domestic trade friction with consequences for firms, prices, and real income.

---

*This project is cleared for the use of confidential microdata from INEGI. We are grateful to Dr. Natalia Volkow and her team in Mexico City's lab for their kind assistance and support.

†Banco de México, email: carlos.cebreros@banxico.org.mx

‡University of California San Diego, email: jmosqued@ucsd.edu

§Johns Hopkins University, email: danielramos@jhu.edu

1 Introduction

Crime makes locations economically farther apart. A producing location can be geographically close to a market but distant in economic terms if the corridor between them exposes carriers to extortion, theft, or violence. These frictions are particularly salient in developing countries, where weak institutions enable "middlemen" to extract rents (Atkin and Khandelwal, 2020). In such an environment, transport firms may reroute, travel longer distances, spend more on prevention, reduce service, or charge higher delivered prices. Thus, market integration, specialization, and real income may suffer as a result of domestic trade frictions (Donaldson, 2018; Sotelo, 2020; Gollin and Rogerson, 2014; Atkin and Khandelwal, 2020; Atkin and Donaldson, 2022).

This paper asks how crime along domestic shipping routes affects transport costs, goods prices, and welfare. We study these questions in the context of Mexico, where crime is a first-order concern: roughly one out of four firms reports to be a victim of crime.[1] In particular, we ask: What is the effect of crime on the costs of transportation firms and agricultural goods prices? What are the aggregate welfare effects of crime as a route-level trade friction? We present evidence that crime affects transportation firms costs, and, although noisier, we show evidence that crime affects prices. Plugging the empirical estimates in a spatial trade model with intermediaries, we find that removing this friction could increase aggregate welfare by 1.3%, with substantial heterogeneity across space. The gains are driven by reductions in trade costs and prices, thereby increasing trade integration across regions.

We contribute by connecting three strands of research that speak to different parts of the same economic mechanism. Domestic trade studies show that transport frictions shape prices, trade costs and flows, and welfare within countries (Donaldson, 2018; Sotelo, 2020). Transportation-market studies shows that trade costs are shaped by intermediaries, route choice, and network structure (Brancaccio et al., 2020; Allen et al., 2024). Work on crime, corruption, and extortion shows that predation can act as a cost of moving goods and be passed through to firms and consumers (Olken and Barron, 2009; Sequeira and Djankov, 2014; Brown et al., 2025). What remains less understood is how criminal risk on domestic roads affects transport firms, delivered prices, and aggregate welfare once shipments can reallocate across routes and markets.

We address this gap studying these questions theoretically and empirically. Theoretically, we develop a spatial trade framework with transportation intermediaries that ship goods across crime-exposed routes along the road network. Each route has a non-

[1] Encuesta Nacional de Victimización a Empresas (INEGI), 2023.

crime component, such as distance or travel time, and a crime component, which we empirically measure using the locations traversed by the route. Crime raises route-specific shipping costs, and intermediaries aggregate route services into an origin-destination transport cost. This structure makes two forces explicit: crime increases bilateral trade costs when it makes the routes connecting an origin and destination more expensive, and the magnitude depends on the ability of intermediaries to substitute across routes. If routes are poor substitutes, local crime shocks translate more directly into delivered prices.

The model serves three purposes. First, it formalizes route-level crime as a domestic trade friction in an economy where intermediaries choose among feasible routes. Second, it serves as a guide for the empirical analysis of transportation costs and agricultural prices. Third, it provides the quantitative structure for counterfactuals. Once crime is mapped into route-level transport costs using the empirical estimates, we can study how much trade costs fall, how trade flows reallocate, and how welfare changes when road insecurity is reduced or removed.

We combine several sources of Mexican administrative and spatial data for the empirical analysis. The core firm-level evidence comes from the confidential *Encuesta Anual de Transportes* (EAT) from *INEGI*, which consists of three panel series, 2014-2017, 2017-2021, 2021-2022, that report transport firms' outcomes such as revenue, cost, employment, and shipping outcomes such as trips, distance traveled, shipment values, vehicle stocks, and their main origin-destination markets. To the best of our knowledge this is the first paper that exploits this dataset. We combine these data with administrative municipality-level crime records from the Secretariado Ejecutivo del Sistema Nacional de Seguridad Pública. We also link transport firms to the *Encuesta Nacional de Victimización de Empresas* (ENVE), which reports direct victimization, economic losses, and prevention spending. Finally, to study the effect of crime on prices we exploit a 29-year panel of wholesale agricultural prices from *Sistema Nacional de Integración de Mercados* (SNIIM) that includes origin-destination prices of fruits and vegetables sold in the main wholesale markets in the country, as well as agricultural output data from *Servicio de Información Agroalimentaria y Pesquera* (SIAP) to identify producing municipalities.

We begin by constructing a shift-share crime exposure measure for the empirical analysis on transportation firms and agricultural prices. We focus on extortion, homicides, and transport-related robbery crimes because these crimes are the among the most reported firm-related crimes (ENVE), and likely capture the presence of organized crime. We construct exposure measures individually for each crime, and for a crime bundle. First, using the road network, we simulate potential route alternatives across all municipalities. The 'shares' are thus given by the probability of choosing a route, using the inverse of driving time, which is purely a function of the geography of the road network – arguably orthogonal to crime. We then aggregate crime counts at the municipality level to the route level using the municipalities that each route crosses. We compute the 'shifts' as percent changes in crime in a given year with respect to a baseline period, 2011-2013, which is a period where crime was at lower levels and more stable relative to the latter part of the decade. This shift-share measure as constructed, we argue, satisfies the exclusion restrictions that the applied econometrics literature has established for the validity of these instruments (Goldsmith-Pinkham et al., 2020; Borusyak et al., 2025).

For the firm-level analysis, we aggregate the origin-destination-year shift-share measure to the firm level using the firm's reported origin-destination pairs, ranked by the share of revenue. Then, we run firm-level regressions with firm and year fixed effects, that control for unobserved time-invariant firm characteristics and common temporal shocks across firms. The results point to a cost-side channel. Firms whose shipment portfolios become more 100% more exposed to crime experience higher total costs (2.9%) and higher fuel costs (2.4%), but do not appear to increase insurance costs. The magnitudes differ by the type of crime exposure, with extortion being the most precise estimate. The bundled crime measure, that includes extortion, homicides, and transport robbery, points in a similar direction for the main cost outcomes, although the mapping is not uniform across all outcomes. We find that extortion decreases shippment values and increases total distance traveled and the number of trips completed. These are margins most closely connected to rerouting, delays, and higher operating costs. We therefore interpret the firm-level evidence as support for the transportation-cost mechanism, with the clearest patterns coming from extortion and broader crime exposure rather than from every individual crime measure.

In addition to the reduced-form analysis, we provide complementary descriptive evidence of how crime affects firms. We do so by analyzing a subset of transportation firms that appear in both the transport EAT and victimization EAT surveys. While the EAT reveals how exposure is associated with operating outcomes, ENVE records whether the firm suffered any type of crime and the direct economic burden of victimization and prevention. Using the matched sample, we show that, in a given year, the typical firm: i) suffers extortion, cargo robbery, and vehicle theft with 0.16, 0.53, and 0.84 probability, respectively; ii) suffers between 2 and 3 incidents of extortion, cargo robbery, and vehicle theft— with the maximum counts being 12, 19, and 24 incidents, respectively; iii) that yearly direct economic losses represent 113% of total costs and 58% of total revenue; and that iv) crime-prevention spending represents 34% of total costs and 18% of total revenue. This evidence shows that the burden of crime for a typical firm can be substantial, potentially an order of magnitude larger than regular taxes.

The second part of the empirical analysis studies wholesale prices of fruits and vegetables. We estimate two main specifications that exploit spatial and temporal variation. The first specification includes crop-year, origin-crop, and destination-crop fixed effects, exploiting crime variation across origin-destination pairs. Intuitively, it compares the price of an avocado produced in Michoacán and sold in Mexico City relative to the price of the same avocado in Guadalajara. We find that doubling the crime bundle exposure leads to 1.2% higher prices. The second specification is directly implied by the model, and includes a much more restrictive set of fixed effects, effectively exploiting within origin-destination pair time variation — using origin-crop-year, destination-crop-year, and origin-destination fixed effects. These fixed effects remove changes in local supply and demand crop-level shocks, as well as time invariant characteristics of the corridors traversed. We find that increasing crime exposure leads to a similar-magnitude positive change in prices (0.6%), although the estimates in these highly saturated specifications are imprecise. We therefore interpret this evidence as consistent with some price pass-through from route exposure, although additional evidence from more goods, e.g. manufactured goods and non-tradables would be useful to establish a general channel.

We then move to quantifying the welfare effects of this friction with our structural framework. We discipline the route-crime-to-cost semi-elasticity with the estimates from the firm-level analysis, set the geography of the model using SNIIM and SIAP locations and the actual natioal road network, and calibrate fundamental parameters using labor, wages, output and value added data from the 2024 Economic Census (INEGI). We find that counterfactually removing crime from all routes raises average welfare by 1.35% by reducing mean bilateral trade costs 2.9% and destination tradable price indices by 6.5% percent, effectively increasing market access and integration. In terms of the spatial effects, these welfare gains are far from uniform: municipalities in the 90th percentile of gains experience up to 5% while in the 10th percentile gain 0.7%. The largest visible gains among producing municipalities appear in parts of Veracruz, Hidalgo, Sonora, Sinaloa, Durango, and Guerrero, consistent with the idea that route-crime reductions matter most where exposed corridors account for an important part of delivered costs.

Taken together, the paper contributes by showing empirically and structurally how crime can operate as a network trade friction, and by quantifying the aggregate effects of this friction. It extends the domestic trade-cost literature by studying an institutional and security-related barrier to market integration, rather than a purely physical or infrastructural one (Donaldson, 2018; Sotelo, 2020). It also contributes to the transport-intermediation literature by observing the firms that move goods and measuring the operating margins through which route risk appears, building on work that shows how intermediaries, route choice, and transport-market structure shape effective trade costs (Brancaccio et al., 2020; Allen et al., 2024). Finally, it contributes to the literature on crime and economic activity by studying the reverse of the more common question: instead of asking how economic shocks affect crime (Dell, 2015; Dell et al., 2019; Dix-Carneiro et al., 2018; Khanna et al., 2025), it asks how crime feeds back into the real economy by raising the cost of moving goods across space. This mechanism is motivated by evidence that predation and corruption distort transport and distribution: Olken and Barron (2009) document structured extortion on Indonesian trucking routes, Sequeira and Djankov (2014) show that port corruption induces firms to reroute shipments, and Brown et al. (2025) show that gang extortion in El Salvador is passed through to delivery fees, retail prices, and consumer welfare. This paper brings these insights into a domestic spatial-trade framework in which local insecurity can affect prices and welfare through the entire road network.

The policy implication is that road security should be treated as a form of trade facilitation. The results suggest that crime along freight corridors raises transport firms' operating costs and, through this channel, may affect prices and welfare. Policy should therefore prioritize economically central shipping routes, particularly corridors with high exposure to extortion and limited route substitutes, rather than focusing exclusively on crime at origins or destinations. In this sense, reducing predation on transport networks can operate like an improvement in domestic infrastructure by lowering the effective cost of connecting markets.

The rest of the paper proceeds as follows. Section 2 describes crime and road insecurity in Mexico. Section 3 develops the spatial trade model with transport intermediaries and route-level crime. Section 4 describes the firm, price, production, crime, and road-network data. Section 5 presents the exposure construction and empirical specifications. Section 6 reports the firm-level, wholesale-price, and matched EAT–ENVE evidence. Section 7 explains how the model is calibrated to the Mexican geography. Section 8 quantifies the welfare cost of route crime through the shipping channel. Section 9 concludes.

2 Context: Crime and Road Insecurity in Mexico

Mexico is a natural setting for studying crime as an economic friction because insecurity is large, persistent, and directly relevant for firms. Crime is not only a household welfare concern or a local public-safety outcome; it is a recurrent feature of the environment in which establishments hire workers, buy inputs, move goods, and choose where to operate. Around 27% of firms reported in *Encuesta Nacional de Victimización de Empresas* to be victims of crime in 2023, equivalent to 1.3 million victimized establishments and 2.9 million crimes against firms. Insecurity is therefore not a rare disruption faced by a small set of firms; it is a common business condition.

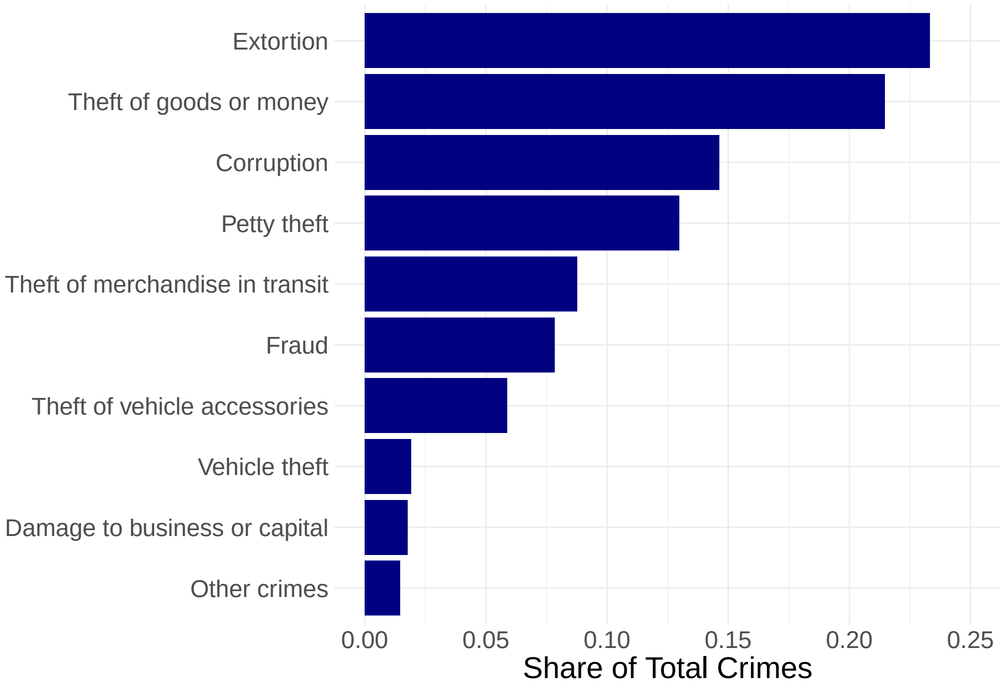

Figure 1: Reported crimes against firms

Notes: The figure reports shares of estimated crimes against firms by crime type. Source: INEGI, *Encuesta Nacional de Victimización de Empresas* (ENVE), 2019 public aggregate tabulations.

The presence of drug cartels and large-scale organized crime makes the Mexican business environment particularly difficult. Figure 1 shows the distribution of reported crimes. Of all reported crimes, extortion and theft of goods or money account for largest shares of reported crimes, while corruption, vehicle theft, and theft of merchandise in transit are also visible categories. These crimes are particularly important because they directly entail costs in the transportation and distribution sectors.

The economic relevance of crime is also visible in its costs. In this same survey, ENVE, it is estimated that the direct cost of insecurity and crime for economic units in 2023 was equivalent to 0.51 percent of GDP. This measure includes direct losses and preventive expenditures.

Further, the incidence of crime affects firms across the size distribution. Figure 2 shows the top three reported crimes by firm size. For micro and small firms, the top crimes are mainly extortion, theft of goods or money, and corruption. For large firms, theft of merchandise in transit appears among the top three crimes, together with vehicle theft and extortion. This pattern is consistent with a simple exposure logic: larger firms operate broader supplier and customer networks, move larger shipment values, and rely more intensively on logistics. Their exposure to crime is therefore not confined to the establishment, but rather it travels with their goods.

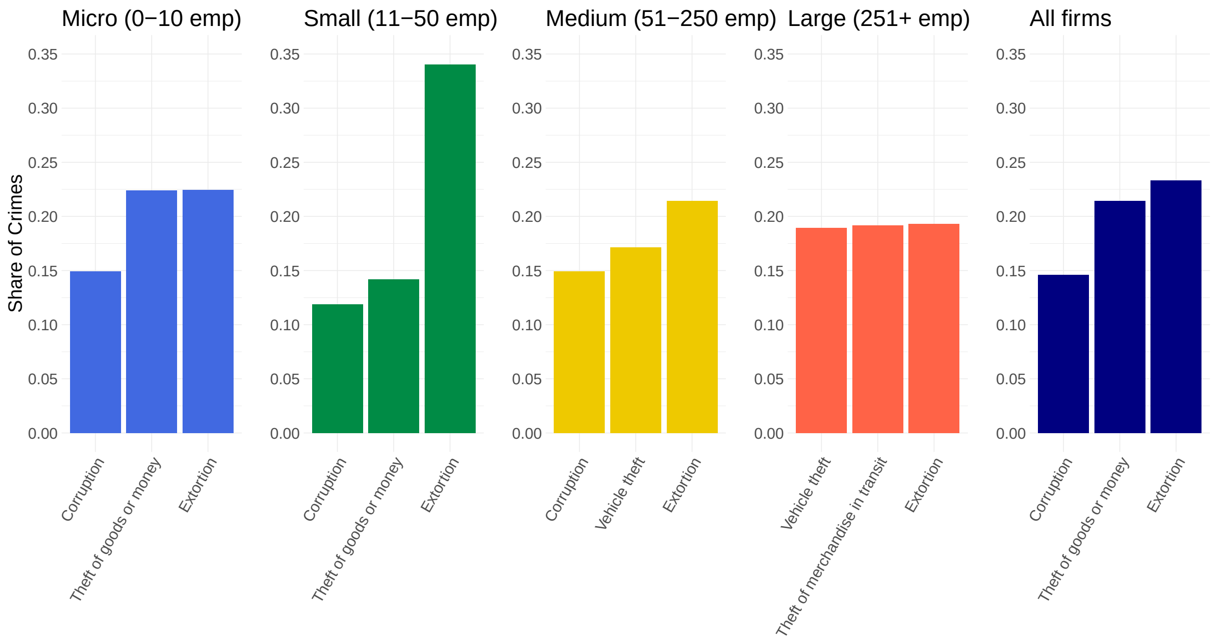

Figure 2: Top reported crimes by firm size, public ENVE 2019

Notes: The figure reports the top three crime categories within each firm-size group. Shares are within firm-size category. Source: INEGI, *Encuesta Nacional de Victimización de Empresas* (ENVE), 2019 public aggregate tabulations.

Mexico is also a good setting to study agricultural prices. Production of fruits and vegetables is highly specialized across space, while consumption is dispersed across large destination markets; as a result, the same crop is often shipped from concentrated producing regions to multiple wholesale markets through different road corridors. This creates spatial variation to study route-level frictions, and connects directly to a literature showing that domestic trade costs shape agricultural prices, market integration, and welfare (Donaldson, 2018; Sotelo, 2020; Atkin and Donaldson, 2022). Mexico is especially useful because SNIIM reports high-frequency wholesale prices by product, origin, and destination market, and SIAP production data make it possible to identify the producing areas behind each reported origin. These features allow us to map price differences across markets to the road corridors over which goods are likely shipped. We will explain in more detail these data in the following sections.

3 Model

This section lays out a spatial model of agricultural trade with transport intermediaries. Goods are produced in locations, moved through a road network with multiple feasible routes, and consumed in destination locations. Crime enters the benchmark model as a route-level shipping friction: insecurity along a corridor raises the effective cost of moving goods through that corridor. The purpose of the model is to connect route-level crime exposure to delivered prices, trade flows, and welfare in a general-equilibrium setting.

3.1 Environment

Goods and space. There are K agricultural goods indexed by k and a set of locations. We use o when a location is the origin of a shipment and d when a location is the destination of a shipment. In the general-equilibrium counterfactual, origins and destinations are therefore two roles played by the same set of inhabited locations. Time is indexed by t. For each origin–destination pair (o, d) there is a set of feasible road routes $\mathcal{R}_{od}$, indexed by r.

The model therefore has three spatial objects:

1. **Origins.** Production takes place in locations o, which differ in agricultural productivity, land endowments, wages, population, and access to the road network.
2. **Destinations.** The same locations appear as destinations d when residents buy delivered agricultural varieties. Destination expenditure is local income, not an exogenous market shifter.
3. **Routes.** Each origin–destination link can be served through several alternative corridors. These routes differ in travel time, road quality, geography, and the crime exposure encountered along the way.

Agents. There are three types of agents. First, competitive agricultural producers in origin o produce good k using labor and land. Second, transport intermediaries buy goods at the farm gate and deliver them to destination locations. Third, residents in each destination location aggregate delivered origin varieties through CES demand. The model also keeps track of local income, cost of living, and population, which are the objects needed to compute real-income welfare.

Within-period sequence. Within each period t, producers choose inputs and supply output at the farm gate, intermediaries organize delivery from each origin to each destination through the road network, residents spend local income on tradables and non-tradables, and local wages, prices, and population adjust to clear markets.

3.2 Production

Technology. Origin o produces good k under constant returns to scale with labor and fixed land:

$$Y_{ot}^{k} = A_{ot}^{k} \left(L_{ot}^{k}\right)^{\mu} \left(\bar{H}_{o}^{k}\right)^{1-\mu}, \tag{1}$$

where A_{ot}^{k} is productivity, L_{ot}^{k} is labor employed in that crop and origin, $\bar{H}_{o}^{k}$ is fixed agricultural land, and $\mu \in (0,1)$ is the labor share.

Factor prices and unit cost. Under perfect competition, producers sell at unit cost. Let w_{ot} denote the local wage and let r_{ot}^{k} denote the rental price of land used in origin o and crop k. Cost minimization implies the unit production cost

$$c_{ot}^{k} = \frac{1}{A_{ot}^{k}} \left(\frac{w_{ot}}{\mu}\right)^{\mu} \left(\frac{r_{ot}^{k}}{1-\mu}\right)^{1-\mu}. \tag{2}$$

The farm-gate cost is increasing in wages and land rents and decreasing in productivity. Because land is fixed, origins with scarcer effective land face higher marginal cost.

3.3 Transport Intermediation and Route-Level Trade Costs

The role of intermediaries. For each origin–destination pair (o, d) and good k, a transport intermediary buys the good at the farm-gate cost c_{ot}^{k} and delivers it to destination d. The intermediary does not take the road network as a black box. Instead, it faces an explicit routing problem: different corridors imply different delivery costs because they differ in physical difficulty and in exposure to crime.

Route-specific generalized shipping cost. Each route $r \in \mathcal{R}_{od}$ is characterized by route-specific crime exposure κ_{odrt}. This scalar object summarizes the insecurity faced along corridor r in period t and is taken as primitive in the model. The intermediary's route-level generalized cost has two parts: a non-crime component that reflects travel time and other predictable shipping difficulty, and a crime component that could capture, for example, extortion, robbery risk, and defensive expenditures. We write

$$g_{odrt} = \delta\, \ell_{odr} + \phi\, \kappa_{odrt}, \tag{3}$$

where ℓ_{odr} captures time-invariant route characteristics such as travel time, distance, or infrastructure quality, and $\delta > 0$ is the elasticity of trade costs with respect to non-crime route difficulty. The parameter $\phi > 0$ is the semi-elasticity of the transport wedge with respect to route crime.

Route-level transport wedge. The generalized shipping cost induces a route-specific cost-equivalent ad valorem wedge

$$\tau_{odrt} = \exp(\delta\, \ell_{odr} + \phi\, \kappa_{odrt})\,. \tag{4}$$

This wedge is the object that enters delivered marginal cost. It should be interpreted as a price-equivalent shipping multiplier: it summarizes all per-unit route-specific delivery frictions in one term, including distance, fuel, labor, maintenance, security costs, delay, and crime risk. When $\ell_{odr} = 0$ and $\kappa_{odrt} = 0$, the route is frictionless and $\tau_{odrt} = 1$.

The intermediary's shipping problem. To deliver one effective unit from origin o to destination d in period t, the intermediary combines route-specific shipping services. Let x_{odrt} denote the amount of route-r shipping service used. The intermediary solves the unit-cost problem

$$\tau_{odt} = \min_{\{x_{odrt}\}_{r\in\mathcal{R}_{od}}} \sum_{r\in\mathcal{R}_{od}} \tau_{odrt} x_{odrt} \quad \text{s.t.} \quad \left[\sum_{r\in\mathcal{R}_{od}} \eta_{odr} \left(x_{odrt}\right)^{\frac{\zeta-1}{\zeta}}\right]^{\frac{\zeta}{\zeta-1}} \geq 1, \tag{5}$$

where $\eta_{odr} > 0$ is a route quality weight and $\zeta > 1$ is the elasticity of substitution across route services. This problem is best interpreted as the aggregate representation of a continuum of intermediaries with heterogeneous operational fit across routes. Some shipments use one corridor, others use another, and the aggregate delivery technology summarizes that allocation in a tractable way.

Bilateral trade cost. The dual of the shipping problem is the origin–destination delivery wedge

$$\tau_{odt} = \left[\sum_{r \in \mathcal{R}_{od}} \eta_{odr} \left(\tau_{odrt} \right)^{1-\zeta} \right]^{\frac{1}{1-\zeta}}. \tag{6}$$

This is the reduced-form bilateral trade cost that enters the rest of the model. Importantly, it is not assumed directly. It is generated by the intermediary's routing problem.

Route shares. Conditional on serving the origin–destination pair, the share of shipping service allocated to route r is

$$\omega_{odrt} = \eta_{odr} \left(\frac{\tau_{odrt}}{\tau_{odt}} \right)^{1-\zeta}. \tag{7}$$

Routes with lower crime and lower physical difficulty attract more traffic. The parameter ζ governs how easily intermediaries can reallocate across corridors. When ζ is high, traffic diverts strongly toward safer or faster routes. When ζ is low, the intermediary is effectively captive to the available network.

Delivered marginal cost and price. Once the shipping problem is solved, the delivered marginal cost of one unit of good k from origin o to destination d in period t is simply

$$mc^k_{odt} = c^k_{ot} \cdot \tau_{odt}. \tag{8}$$

If the intermediary sector is competitive, the delivered price equals delivered marginal cost. More generally, under isoelastic demand the intermediary charges a constant product-specific markup $\mu^k_I \geq 1$, so that

$$p^k_{odt} = \mu^k_I \, c^k_{ot} \cdot \tau_{odt}. \tag{9}$$

This is the key object that links route-level crime to delivered prices. Taking logs gives

$$\ln p^k_{odt} = \ln \mu^k_I + \ln c^k_{ot} + \ln \tau_{odt}. \tag{10}$$

Thus, any empirical delivered-price regression must absorb origin–product–time production costs and isolate the part of the bilateral delivery wedge that varies with route crime.

3.4 Demand and Price Aggregation

Origin aggregation. Destinations combine origin varieties through an Armington CES structure. For each good k, the expenditure share of destination d on origin o is

$$S^k_{odt} = \beta^k_{od} \left(\frac{p^k_{odt}}{P^k_{dt}} \right)^{1-\sigma}, \qquad P^k_{dt} = \left[\sum_o \beta^k_{od} \left(p^k_{odt} \right)^{1-\sigma} \right]^{\frac{1}{1-\sigma}}, \tag{11}$$

where S^k_{odt} is the expenditure share of origin o in destination d for good k, β^k_{od} is an origin preference or quality shifter, and $\sigma > 1$ is the elasticity of substitution across origins.

Goods aggregation. At the top tier, destinations aggregate across agricultural goods:

$$P^T_{dt} = \left[\sum_k \alpha^k_d \left(P^k_{dt} \right)^{1-\rho} \right]^{\frac{1}{1-\rho}}, \tag{12}$$

where α^k_d is the taste weight on good k and $\rho > 1$ is the elasticity of substitution across goods.

Expenditure and trade flows. Let I_{dt} denote total income in destination location d and let α denote the expenditure share on tradables. Tradable expenditure is

$$E^T_{dt} = \alpha I_{dt}. \tag{13}$$

In the shipping-only benchmark used for the first counterfactual, there are no non-tradables or land-rent feedbacks, so local income is labor income and $E^T_{dt} = w_{dt} L_{dt}$. More generally, I_{dt} includes all local factor income, including land rents when the fixed factor is active.

Let s^k_{dt} be the expenditure share of good k in destination d,

$$s^k_{dt} = \alpha^k_d \left(\frac{P^k_{dt}}{P^T_{dt}} \right)^{1-\rho}. \tag{14}$$

Nominal expenditure by destination d on good k from origin o is then

$$X^k_{odt} = S^k_{odt} s^k_{dt} E^T_{dt}. \tag{15}$$

This is the object that closes the goods market and links destination income to origin revenues.

3.5 Additional Equilibrium Blocks

Non-tradable sector. Each location also has a non-tradable sector such as local services, storage, and housing. For an origin o, output is produced with labor and a fixed factor:

$$Y^{NT}_{ot} = A^{NT}_{ot} \left(L^{NT}_{ot} \right)^{1-\xi} \bar{H}^{\xi}_o, \tag{16}$$

where A^{NT}_{ot} is non-tradable productivity, L^{NT}_{ot} is labor in the local sector, and $\bar{H}_o$ is fixed non-tradable land or space. The associated non-tradable price is denoted p^{NT}_{ot}.

Cost of living. Residents consume tradables and non-tradables. Let P^T_{ot} denote the tradable price index relevant for households in origin o, and let $\alpha \in (0, 1)$ be the expenditure share on tradables. The local cost-of-living index is

$$\mathcal{P}_{ot} = \left(P^T_{ot} \right)^{\alpha} \left(p^{NT}_{ot} \right)^{1-\alpha}. \tag{17}$$

Population. Population can respond to real income and local amenities. Let B_{ot} denote the amenity level of origin o, $\mathcal{P}_{ot}$ its cost-of-living index, and $\nu > 0$ the mobility parameter. A convenient reduced-form expression for the population allocation is

$$L_{ot} \propto \bar{T}_o \left(\frac{B_{ot}\, w_{ot}}{\mathcal{P}_{ot}} \right)^{\nu}, \tag{18}$$

where $\bar{T}_o$ is an exogenous location weight. When ν is small, migration is sluggish and population is nearly fixed. When ν is large, people reallocate more strongly toward high-real-wage, high-amenity locations.

3.6 Route Crime

Benchmark channel. In the main specification, crime affects the economy by raising route-specific shipping costs. The crime object is the route exposure κ_{odrt} in Eq. (4). Holding non-crime distance and route quality fixed, higher κ_{odrt} raises the route transport wedge τ_{odrt}, shifts traffic across available corridors through Eq. (7), and raises the bilateral delivery wedge τ_{odt} through Eq. (6).

This channel maps directly into delivered prices. For any good k, route crime increases the delivered marginal cost in Eq. (8) and therefore the delivered price in Eq. (9). Destinations then substitute across origins and goods according to the CES demand system. In general equilibrium, these price changes feed back into local revenues, wages, destination expenditure, and real-income welfare.

Scope. The benchmark focuses on the route-shipping channel because it is the channel most directly connected to the empirical price and firm-cost evidence. Appendix A records model extensions that are not used in the current quantitative exercise.

3.7 Equilibrium and Benchmark Cases

Equilibrium. Given route characteristics, land endowments, demand shifters, and baseline local fundamentals, a competitive equilibrium in each period consists of agricultural labor allocations, wages, non-tradable prices, bilateral trade costs, route shares, delivered prices, destination price indices, income, expenditure, and population levels such that:

1. producers minimize costs and sell at unit cost,
2. intermediaries solve the shipping problem in Eq. (5),
3. residents allocate expenditure according to the CES demand system,
4. destination tradable expenditure equals local tradable income under balanced trade,
5. labor, land, and goods markets clear, and
6. population satisfies the mobility condition in Eq. (18).

The key balanced-trade restriction is that destination expenditure is generated inside the model. In the benchmark with fixed population, no non-tradables, and labor-only agricultural production,

$$E^T_{dt} = w_{dt} L_{dt}, \qquad w_{ot} L_{ot} = \sum_{d,k} X^k_{odt}. \tag{19}$$

With land or non-tradables active, the same logic applies using total local factor income rather than labor income alone. Appendix B derives the closure and the wage fixed point in detail.

A useful benchmark. The quantitative benchmark isolates the route-shipping channel. Population is fixed, non-tradables are inactive, and crime enters only through the intermediary shipping wedge in Eq. (4). This benchmark is the version we calibrate below and use for the first counterfactual.

4 Data

We combine firm surveys, wholesale price records, agricultural production data, municipal crime data, and the road network. The firm data identify how transport intermediaries respond when their shipment portfolios become more exposed to crime. The price data identify whether those route-level shocks also appear in delivered agricultural prices. The remaining data sources locate production, destination markets, and crime along the road corridors that connect them.

4.1 The EAT transport-firm panel

Our firm-level evidence comes from the confidential *Encuesta Anual de Transportes* (EAT), which surveys land-transport firms in SCIAN 48–49. The survey is organized into three series, corresponding to the 2008, 2013, and 2018 questionnaire frames. We use the linked EAT panel for the years in which firms can be merged to our route-level crime measures. For each firm-year, EAT reports a rich set of operating outcomes, including total revenue, total costs, fuel expenditures, insurance expenditures, employment, total trips, total distance, total tons moved, total vehicles, and the total value shipped.

EAT is particularly useful for our setting because firms report their main origin-destination (OD) markets and the revenue share associated with each one. Each firm reports up to four top OD pairs. These reported revenue shares are the outer weights that allow us to aggregate route-level crime exposure into a firm-level measure. Table 1 summarizes the coverage of the three series, while Table 2 reports the distribution of the main firm outcomes used in the paper.

Table 1: EAT Panel Coverage by Survey Series

Series	First year	Last year	Firm-years	Mean firms/year	Max firms/year
Serie 2008	2015	2018	1,513	378	553
Serie 2013	2017	2021	1,696	424	424
Serie 2018	2021	2022	530	265	266

Notes: Coverage is computed from the EAT firm-year sample that can be matched to OD-level crime exposure. Serie 2013 is our preferred subpanel because the number of firms is constant across years.

Table 2: Selected Descriptive Statistics in the EAT Panel

Variable	Observations	Mean	P25	Median	P75
Total revenue	3,739	306,617	44,668	153,743	349,232
Shipping revenue	3,738	296,938	43,120	145,424	342,172
Total costs	3,739	180,855	26,643	92,373	212,618
Fuel expenditure	3,209	75,272	11,692	39,633	91,000
Insurance expenditure	3,721	6,553	309	2,768	7,633
Employment	3,739	211	29	100	257
Total vehicles	3,672	244	20	98	296
Total trips	3,316	31,387	3,646	12,096	32,292
Total distance	3,316	12,654,968	1,589,705	5,263,200	15,000,600
Total tons shipped	3,316	585,627	54,318	207,360	603,654
Value shipped	2,531	425,406	0	6,500	175,864

Notes: This table reports the pooled distribution of the main firm outcomes used in the EAT analysis after constructing the firm-level route-crime exposure measure. Monetary amounts are reported in the survey units used by EAT. All descriptive moments are rounded to whole numbers. In the regression analysis we work with the logarithm of these outcomes.

Our preferred estimating sample is Serie 2013. Relative to the other two frames, it delivers the cleanest panel structure for a fixed-effects design: it covers four years, from 2017 through 2021, and contains the same 424 firms in every year. By contrast, Serie 2008 exhibits substantial changes in the number of observed firms over time, and Serie 2018 covers only two years. We therefore treat Serie 2013 as the main subpanel and use the pooled sample as a complement.

4.2 Crime and road network data

Our crime measures combine municipality-year outcomes from the Secretariado Ejecutivo del Sistema Nacional de Seguridad Pública (SESNSP) with the Mexican road network. The empirical crime bundle contains transporter robbery, extortion, and homicide, measured at the municipality-year level over 2011–2025. These are the crimes most closely connected to route security and the broader violence environment faced by firms and shipments.

Figure 3 plots annual national totals for the bundle and for each individual crime. As it can be seen, the aggregate bundle is relatively stable and at lower levels between 2011 and 2013, and then it increases towards the end of the decade. This is the temporal variation that we will exploit, as we will explain in the following section.

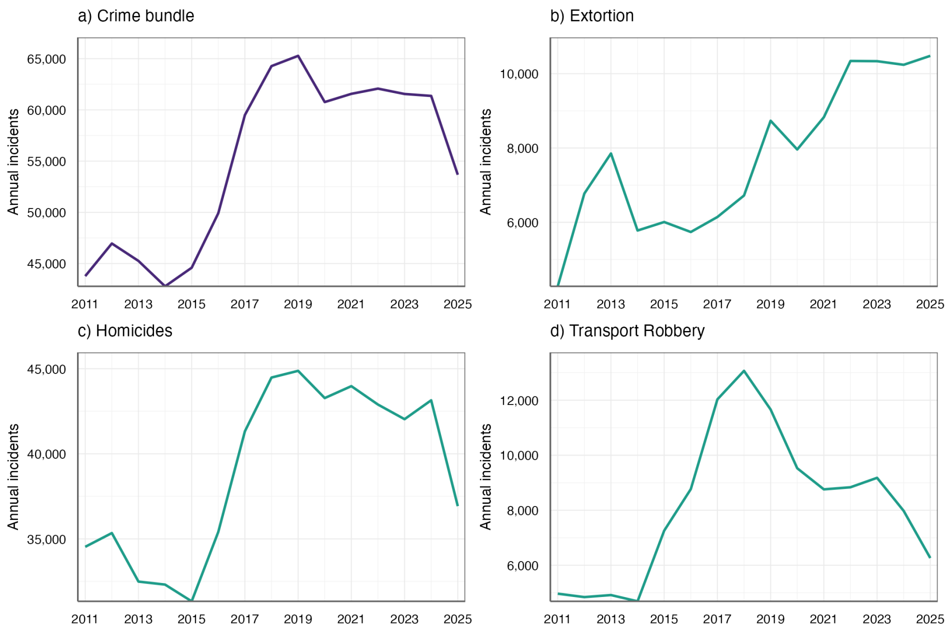

Figure 3: Temporal variation of crime

Notes: The figure reports annual national incident totals for the crime measures used to construct route-level exposure. Panel a reports the three-crime bundle; panel b reports extortion; panel c reports homicide; and panel d reports transporter robbery. Sources: Secretariado Ejecutivo del Sistema Nacional de Seguridad Pública (SESNSP), municipal criminal incidence records.

Crime is far from uniform across space. Figure 4 maps municipalities by total incidents per road kilometer over 2011–2025, separately for the crime bundle and for each component. The distribution is highly uneven across space, with high-exposure municipalities appearing mostly in central Mexico, around major metro areas such as Monterrey and Mexico City, Sinaloa, and in border cities such as Tijuana and Matamoros.

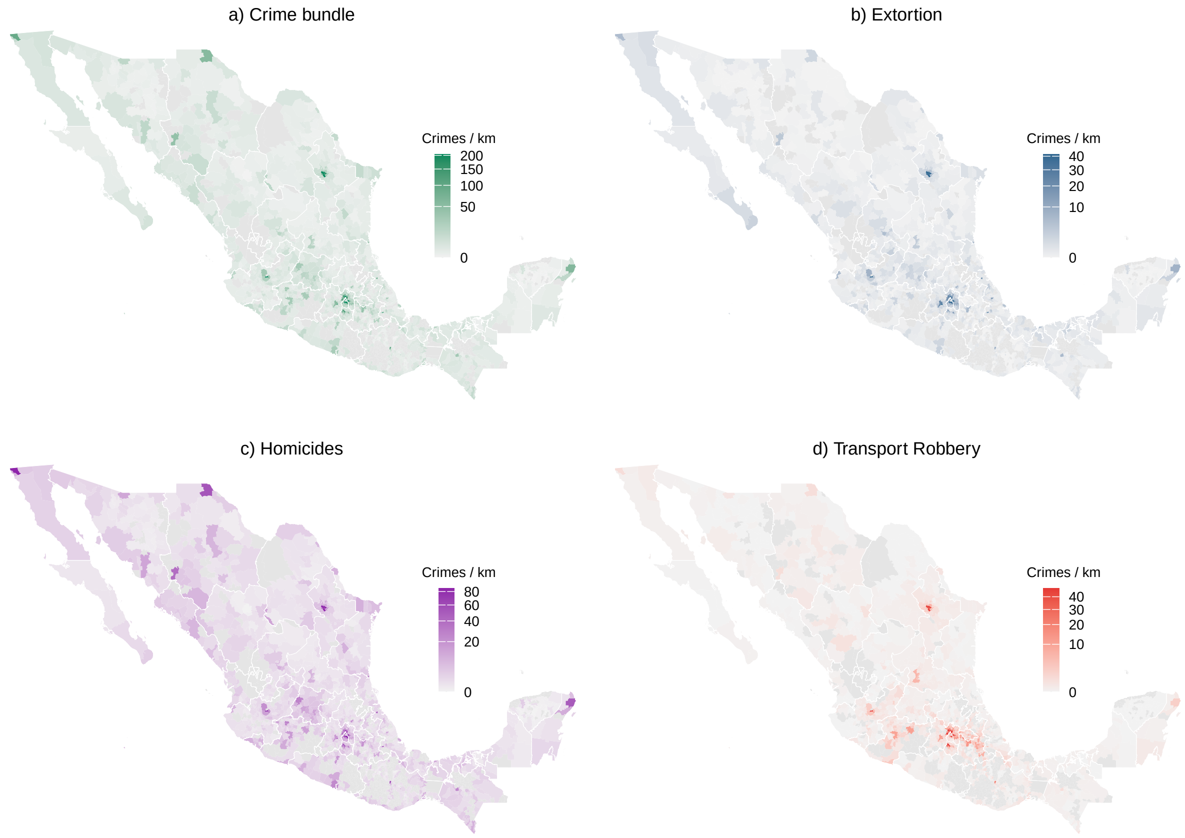

Figure 4: Municipality-level crime exposure per road kilometer, 2011–2025

Notes: Each panel reports total incidents over 2011–2025 divided by municipal road kilometers. The crime bundle is the sum of transporter robbery, extortion, and homicide. Color scales are panel-specific square-root scales capped at the 99th percentile to preserve within-panel variation in the presence of extreme values. Municipalities without valid road-kilometer denominators are shown in gray. Sources: Secretariado Ejecutivo del Sistema Nacional de Seguridad Pública (SESNSP), municipal criminal incidence records; Red Nacional de Caminos; INEGI, 2010 municipal boundaries.

We construct feasible origin-destination routes between the markets reported by EAT firms and between SIAP-weighted agricultural origins and the wholesale markets observed in SNIIM. The route construction uses the Red Nacional de Caminos road network, keeping the geography fixed so that time variation in exposure comes from crime rather than from changes in the road map. We then overlay municipality-level crime outcomes on the municipalities traversed by each route to obtain route-year measures for each crime type. This route-level panel is the raw input for the exposure objects used in the empirical strategy. The next section describes how route-year crime is aggregated into OD-level and firm-level exposure measures for the EAT and SNIIM regressions.

4.3 Enterprise victimization survey (ENVE)

The *Encuesta Nacional de Victimización de Empresas* (ENVE) is INEGI's enterprise victimization survey. It records whether economic units were victims of crime, the types of crimes they experienced, the number of incidents, whether incidents were reported to authorities, direct monetary losses, and spending on prevention. ENVE is therefore useful for measuring the firm-side costs of insecurity that do not appear in administrative crime records: extortion, theft, vehicle theft, theft of merchandise in transit, prevention spending, and the large gap between crimes experienced and crimes reported.

In the paper, ENVE plays a complementary role. Public ENVE aggregates motivate the setting by showing that insecurity is widespread across Mexican firms and that transport-relevant crimes are especially salient for larger firms. We also link ENVE to EAT when firm identifiers allow it, which gives a smaller matched sample used to interpret the EAT cost regressions: the match helps separate direct losses from preventive spending and gives a scale for the crime costs reported by transport firms.

4.4 Wholesale prices and agricultural origins

SNIIM wholesale prices. Our price evidence comes from the *Sistema Nacional de Información e Integración de Mercados* (SNIIM), which records wholesale prices for fruits and vegetables in Mexican wholesale markets. Each quote identifies a product, commercial presentation, origin, destination market, date, and reported price. The cleaned daily file contains about 15 million price observations from January 2000 through February 2026, covering 222 products and 60 destination markets. For the empirical price exercise, we use the frequent wholesale price and aggregate the daily records to product–presentation–origin–destination monthly cells. The yearly panel used in the main price table then averages monthly log prices within product–origin–destination–year cells and averages product varieties within the same crop so that the price outcome has the same crop–origin–destination–year grain as the route-crime exposure.

Figure 5 illustrates the time-series variation in four high-coverage products after removing broad seasonal and annual movements. The residual price series show that there remains substantial within-route variation once common calendar-month and year means are removed. Appendix Figure 9 summarizes the same residual variation across destination markets over time, which is useful for assessing whether price gaps across markets are stable or widen in periods when transport conditions deteriorate.

An interesting fact across these two figures is that dispersion of prices seems to increase over time, and particularly after 2010. This increase in dispersion coincides with the rise in crime during the 2010 decade. This is of course not causal evidence, but a useful fact to recall when identifying more formally the effect of crime on prices in the following section.

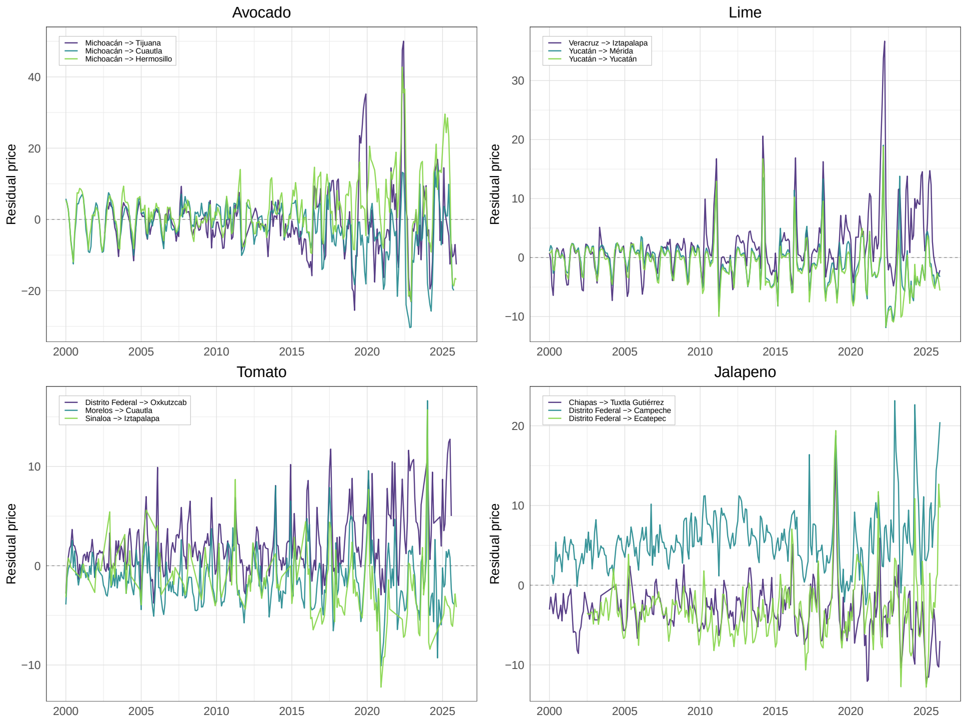

Figure 5: Residual wholesale price series for selected SNIIM products

Notes: Each panel plots monthly residual frequent wholesale prices for one product. Lines correspond to the three origin–destination pairs with the most observations for that product. Residual prices remove product-specific calendar-month and year means and add back the product mean. Prices are nominal Mexican pesos per kilogram.

Two features of SNIIM matter for the design. First, destinations are specific wholesale markets, not just states, so we keep destination markets distinct when building routes and estimating price regressions. Second, SNIIM origins are reported as origin states or origin labels, not as precise farm municipalities. This means that SNIIM can identify the state from which a product is shipped, but an additional production source is needed to locate the relevant origin within that state.

Agricultural production (SIAP). We use the *Cierre de la Producción Agrícola* from SIAP to locate agricultural production within SNIIM origin states. SIAP reports annual municipality-level production by crop for Mexico. We map SNIIM product names to SIAP crop names; the matched set contains 22 SIAP crops, including the main fruits and vegetables used in the price exercise. These data enter the paper in two ways. First, production-weighted crop-state centroids provide geographic origin points for routing SNIIM shipments through the road network. Second, municipality shares of state-crop production provide the weights used to aggregate municipality-level route exposure into state-origin by destination-market by crop exposure.

The production weights are predetermined relative to the crime variation used in the price regressions. We use pre-Calderón production geography to construct static crop-state centroids and use the 2005 SIAP distribution to compute the municipality shares. This is important because the weights should describe where production was located before the rise in violence, rather than allowing the origin weights themselves to adjust to later crime shocks. Figure 6 summarizes this concentration. For most matched crops, production within a state is concentrated enough that a production-weighted origin is informative, and the leading production municipalities are relatively stable over time.

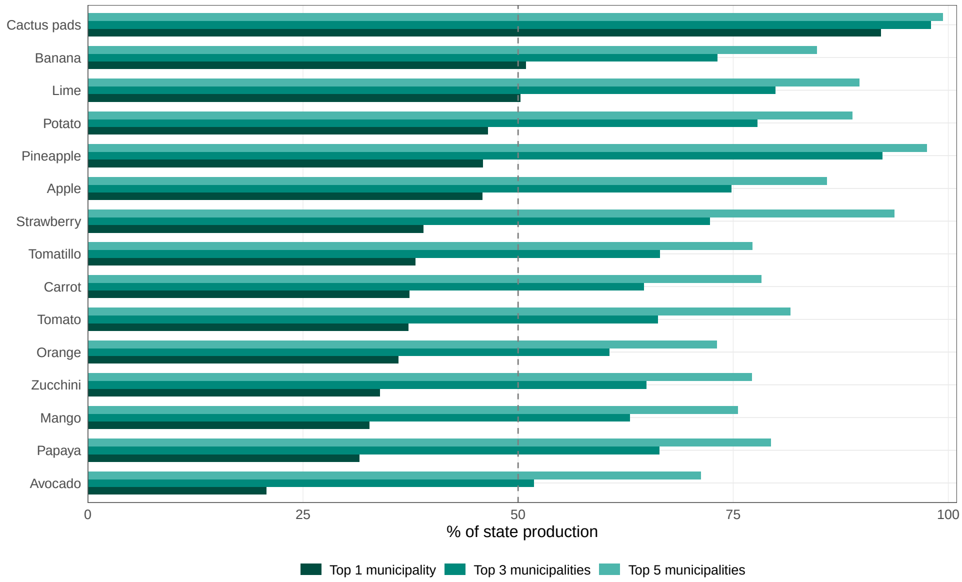

Figure 6: Within-state concentration of SIAP production

Notes: Bars show the production-weighted share of state-crop output produced by the largest one, three, or five municipalities within each producing state for selected matched crops. Shares are computed from 2005–2006 municipal production and averaged across producing states using state-crop production as weights. The dashed vertical line marks 50 percent of state production.

5 Empirical Strategy

The model provides the empirical roadmap for the paper. It says that crime matters because goods are moved through routes, not because they originate or end in a single place. That route-based logic leads to two empirical questions. First, do transport firms whose shipment portfolios are more exposed to crime experience higher costs and adjust their operations? Second, do destination markets reached through more crime-exposed corridors face higher wholesale prices? We organize the empirical strategy around those two questions.

The regressions below should be read as model-guided empirical tests rather than as exact estimating equations for every structural parameter. The model disciplines the exposure measure, the relevant margins, and the sign predictions. We then use the estimated magnitudes as inputs and moments for the quantitative calibration.

5.1 Measuring Exposure to Crime

A route-based notion of exposure. In the model, what matters for trade is not only crime at the endpoint locations, but crime encountered along the corridor that connects the origin and the destination. Our empirical exposure measure follows that logic. Firms ship across space, pass through multiple municipalities, and can often use more than one feasible route between the same origin and destination. Exposure to crime should therefore be measured at the route level and then aggregated up using the geography of the shipment.

Route-level crime intensity. For each route r connecting origin o to destination d, we compute route-level exposure as the average crime intensity across the interior municipalities traversed by that route:

$$\kappa_{odrt} = \frac{1}{|M_{odr}^{int}|} \sum_{m \in M_{odr}^{int}} \rho_{mt}, \tag{20}$$

where ρ_{mt} is the municipality-year crime bundle, standardized by road kilometers, and M_{odr}^{int} denotes the set of interior municipalities on route r. This is the empirical counterpart of the scalar route-level crime object in the model.

Figure 7 illustrates the spatial viaration of crime across different route alternatives from selected origins. For example, focusing on panel c) Monterrey, it can be seen that a firm that ships from Monterrey to Guadalajara or Ciudad de Mexico is subtantially more exposed to crime relative to shipping to Ciudad Juarez in the North. At the same time, route alternatives matter. If a firm ships from Tijuana, as in panel d), it will not have many alternatives to ship to Ciudad Juarez or Monterrey as it has available only single major roads. As a result, this firm will be heavily exposed to changes in crime that occur along its only way.

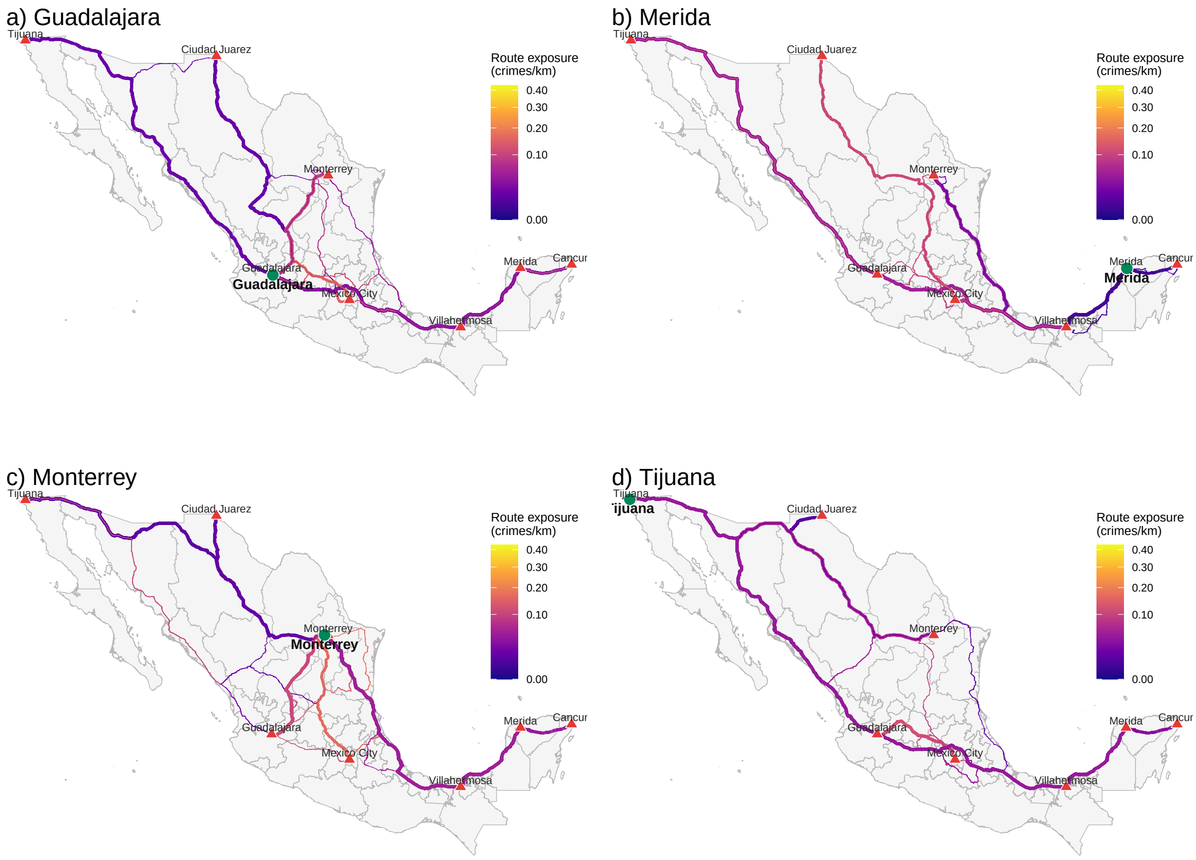

Figure 7: Route alternatives and crime exposure for selected origins

Notes: Each panel displays road-network route alternatives from the origin municipality named in the panel to eight wholesale-market destinations: Tijuana, Ciudad Juarez, Mexico City, Guadalajara, Monterrey, Cancun, Villahermosa, and Merida. Routes are colored by average crime exposure over 2011–2025, measured as the three-crime bundle per kilometer along the route. The color scale is capped at the 95th percentile across these example routes. Line width increases with inverse-time route share; triangles mark destination markets and filled circles mark origins.

From route-level crime to OD exposure. Most origin-destination pairs have several feasible routes. Because realized route choice is not observed directly in all of our datasets, we approximate the relative importance of each feasible corridor using predetermined route weights. In the baseline, ω_{odr} is the inverse-travel-time share of route r within the OD-specific route set. For each route we compute a baseline-relative crime shift,

$$g_{odrt} = \frac{\kappa_{odrt} - \bar{\kappa}_{odr,2011\text{–}2013}}{\bar{\kappa}_{odr,2011\text{–}2013}}, \tag{21}$$

and aggregate across feasible routes:

$$\kappa_{od,t} = \sum_{r \in \mathcal{R}_{od}} \omega_{odr} \cdot g_{odrt}. \tag{22}$$

This OD-level exposure index captures both where crime is located and how relevant each corridor is for moving goods between that origin and destination.

Identification and the shift-share exposure design. Our route-crime measure has the structure of a shift-share exposure design. In the jargon of Goldsmith-Pinkham et al. (2020) and Borusyak et al. (2025), the object in equation (22) is a weighted average of crime shifts, where the shares ω_{odr} are predetermined route weights and the shifts g_{odrt} are changes in crime exposure along each feasible corridor. The shift-share literature, in sum, establishes that in order for the instrument to be valid, we need that either the shares or the shifts are exogenous. We interpret our identifying variation primarily through the share-exogeneity logic emphasized by this literature.

The shares are not chosen as a function of contemporaneous crime, prices, or firm outcomes; they are engineering features of the road network, determined by feasible route paths and inverse driving time. Moreover, they are tailored to the treatment: they measure precisely the extent to which an origin–destination pair is mechanically exposed to crime while goods are in transit, rather than a generic measure of market size, distance, or local economic activity. The identifying comparison is therefore between otherwise similar shipment relationships whose predetermined road geography exposes them differentially to subsequent changes in crime along transport corridors.

The shifts are less naturally viewed as fully as-good-as-random municipality-level shocks. Crime may respond to broader local conditions and to the value of economic activity on some corridors. For this reason, we construct shifts as changes relative to the 2011–2013 baseline, a period in which the relevant road-crime measures were comparatively low and stable. This removes persistent differences in crime levels across corridors and focuses the variation on the sharp post-baseline increase in insecurity. Under our preferred interpretation, the exclusion restriction is that, conditional on the fixed effects and controls in the estimating equation, origin–destination pairs, firms, or destination markets with different predetermined route shares would not have experienced systematically different changes in costs or prices absent the subsequent crime increases along those routes.

5.2 Effects of Crime on Transport Firms

Using the model as a guide. In the model, τ_{odt} is a reduced-form ad valorem shipping wedge. It summarizes the per-unit frictions faced when serving an origin–destination pair, including distance, fuel, labor, maintenance, security costs, delay, and crime risk. Because route crime raises this wedge, the model predicts that firms serving more exposed OD pairs should experience changes in outcomes that move with shipping frictions. We therefore begin with the firm data, where the route-cost mechanism is most direct.

Firm-level exposure. EAT reports each firm's main origin-destination markets together with their revenue shares. We use those shares to aggregate OD-level crime exposure into a firm-level Bartik-style index:

$$B_{f,t} = \frac{\sum\limits_{od \in \mathcal{M}_{f,t}} w^{rev}_{f,od,t} \cdot \kappa_{od,t}}{\sum\limits_{od \in \mathcal{M}_{f,t}} w^{rev}_{f,od,t}}, \tag{23}$$

where $\mathcal{M}_{f,t}$ is the subset of the firm's reported OD pairs that matches to the route-level crime files. This exposure measure asks whether firms whose commercial geography places more weight on crime-exposed corridors experience different outcomes than otherwise similar firms. It is the firm-level analogue of the model's route-crime wedge: the same OD exposure $\kappa_{od,t}$ that shifts τ_{odt} is averaged over the firm's reported shipment portfolio.

Firm-level specification. Our baseline firm-level specification is

$$\ln Y_{f,t} = \alpha_f + \alpha_t + \beta \cdot B_{f,t} + \varepsilon_{f,t}, \tag{24}$$

where α_f are firm fixed effects that control for time-invariant characteristics of the firm, and α_t are year fixed effects that control for economy-wide conditions. We estimate this equation for total operating costs, fuel expenditure, insurance expenditure, total trips, total distance, employment, revenue, and shipped value.

We do not treat the EAT regression as a full structural estimation of the model. Instead, it is the first reduced-form test of the mechanism emphasized by the theory. For cost-side outcomes, β is the empirical counterpart of ϕ: it measures how strongly observed firm costs respond to the same route-crime variation that shifts τ_{odt} in the model. If crime operates through the route-cost channel, then β should be positive for outcomes such as total costs, fuel, and distance. For activity outcomes such as revenue or shipped value, the sign is theoretically more ambiguous because firms may pass costs through, reroute, or shrink activity.

5.3 Effects of Crime on Wholesale Prices

The second implication of the model is that higher shipping costs should not remain confined to the balance sheet of transport firms. In the origin–destination version of the model, the delivered wholesale price of product k shipped from origin o to destination d in year t is

$$P^k_{od,t} = \mu^k_I c^k_{ot} \tau_{od,t}, \tag{25}$$

where c^k_{ot} is the origin–product–time farm-gate cost and $\tau_{od,t}$ is the bilateral transport wedge generated by the route network. Taking logs gives

$$\ln P^k_{od,t} = \ln \mu^k_I + \ln c^k_{ot} + \ln \tau_{od,t}. \tag{26}$$

Thus, the model says that delivered wholesale prices vary because of origin-specific production costs and because of bilateral transport costs. The empirical price regression must therefore control for origin–product–time production costs and isolate the component of the bilateral transport wedge that varies with crime exposure along the routes connecting o and d.

Mapping route crime into the transport wedge. The model's routing problem implies that the OD transport wedge is an aggregate of route-specific wedges. Let κ_{odrt} denote crime exposure along route r between origin o and destination d in year t. A first-order approximation to the route aggregator around baseline route shares gives

$$\Delta \ln \tau_{od,t} \simeq \phi \sum_{r \in \mathcal{R}_{od}} \omega_{odr,0}\, \Delta \kappa_{odrt}, \tag{27}$$

where $\omega_{odr,0}$ is the baseline importance of route r in the OD transport wedge and ϕ is the semi-elasticity of route costs with respect to crime exposure.

Equation (27) is a change equation. Combining it with Eq. (26) gives

$$\Delta \ln P^k_{od,t} \simeq \Delta \ln c^k_{ot} + \phi \sum_{r\in\mathcal{R}_{od}} \omega_{odr,0}\, \Delta\kappa_{odrt}. \tag{28}$$

Define the scalar OD crime exposure bundle as

$$\kappa_{od,t} \equiv \sum_{r\in\mathcal{R}_{od}} \omega_{odr,0}\, \Delta\kappa_{odrt}. \tag{29}$$

Then the model implies that changes in delivered wholesale prices are proportional to changes in route-share-weighted crime exposure, after accounting for changes in origin–product production costs. The same approximation can be written in levels once the exposure measure is normalized relative to a baseline period. In particular,

$$\ln \tau_{od,t} \simeq \lambda_{od} + \phi\kappa_{od,t}, \tag{30}$$

where λ_{od} absorbs the baseline OD transport wedge and all time-invariant route characteristics, including distance, road quality, and geography. Substituting Eq. (30) into Eq. (26) yields the levels counterpart of the model-implied price equation:

$$\ln P^k_{od,t} \simeq \ln c^k_{ot} + \lambda_{od} + \phi\kappa_{od,t} + \text{constant}. \tag{31}$$

Price regression. Guided by Eq. (31), our the model-implied wholesale-price specification is

$$\ln P^k_{od,t} = \theta\kappa_{od,t} + \alpha^k_{ot} + \psi^k_{dt} + \lambda_{od} + \varepsilon^k_{od,t}, \tag{32}$$

where $P^k_{od,t}$ is the observed SNIIM wholesale price for good k shipped from origin o to destination market d, and $\kappa_{od,t}$ is the scalar route-share-weighted crime exposure bundle for that OD pair. The coefficient θ is the price pass-through of route crime exposure into delivered wholesale prices.

The fixed effects in Eq. (32) come directly from the price equation. The origin–product–year fixed effects α^k_{ot} absorb the farm-gate cost term $\ln c^k_{ot}$, including origin-specific productivity, wages, land rents, harvest shocks, and source-side crime that affects production. The destination–product–year fixed effects ψ_{dkt} absorb product-specific wholesale-market shocks at the destination, such as local demand, storage and handling conditions. The origin–destination–product fixed effects λ_{od} absorb permanent bilateral heterogeneity, including distance and road quality. Identification therefore comes from changes over time in route-level crime exposure within the same OD–

product cell, after comparing only variation that is not explained by origin–product–year or destination–product–year shocks.

An alternative specification that we run is

$$\ln P^k_{od,t} = \theta\kappa_{od,t} + \alpha^k_o + \psi^k_d + \lambda^k_t + \varepsilon^k_{od,t}, \tag{33}$$

where variation comes from relative prices from the same good across different destinations, rather than within an origin-destination pair. This specification effectively compares, for example, the price of an avocado from Michoacán delivered at Mexico City relative to Guadalajara at a given point in time, controling for origin-crop and destination-crop time-invariant characteristics.

6 Results

This section presents three sets of evidence. We first use the EAT firm panel to test whether route-crime exposure appears in transport firms' operating costs and activity. We then use SNIIM wholesale prices to ask whether the same exposure is reflected in delivered agricultural prices. Finally, we use the matched EAT–ENVE sample to describe the direct burden of victimization and prevention spending among transport firms.

6.1 Firm-level evidence from EAT

The firm-level regressions estimate Eq. (24): log firm outcomes are regressed on firm-level route-crime exposure, with firm and year fixed effects. The preferred sample is EAT Serie 2013. As discussed in Section 4, this is the cleanest subpanel for a within-firm design because it follows a stable set of firms over 2017–2021. Appendix D reports analogous results using the pooled EAT sample.

Table 3: Firm-level EAT regressions: Series 2013, cost outcomes

	Log total costs				Log fuel expense				Log insurance expense			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Transporter-robbery exposure	-0.00005***				-0.00004***				-0.00000			
	(0.00001)				(0.00001)				(0.00002)			
Extortion exposure		0.00439***				0.00659***				0.00096		
		(0.00062)				(0.00101)				(0.00197)		
Homicide exposure			0.09660**				0.06475				-0.08916	
			(0.04552)				(0.05947)				(0.18326)	
Crime exposure				0.02855**				0.02400*				0.00265
				(0.01178)				(0.01429)				(0.03774)
Observations	775	800	805	805	775	800	805	805	696	715	720	720
Firm FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Firm IDs	303	314	314	314	303	314	314	314	278	287	287	287
Years	4	4	4	4	4	4	4	4	4	4	4	4

Notes: Each column is a separate firm-level regression with firm and year fixed effects; only the row for the exposure included in that column is populated. The displayed regressors are firm-year crime exposure measures built in the INEGI lab by merging each firm's reported top origin-destination revenue shares to OD-year route-crime exposure, then averaging over matched OD pairs and renormalizing over the matched share. OD-year route exposure is based on municipal crimes normalized by road kilometers along the relevant routes and expressed as a percent-change-from-base shift. The single-crime rows report transporter robbery, extortion, and homicide exposures; the bundled row reports the crime-bundle exposure available in the frozen lab return. Standard errors in parentheses are clustered at the firm-series level. All numeric coefficient and standard-error entries are reported to five decimals. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 3 reports the cost-side results. The clearest and most stable pattern comes from extortion exposure. A one-unit increase in extortion exposure–that is, a 100 percent increase relative to the baseline exposure index–is associated with a 0.00439 increase in log total costs and a 0.00659 increase in log fuel expenses. Both coefficients are statistically significant at the one percent level. The implied magnitudes are modest, about 0.4 and 0.7 percent, respectively, but they are precisely estimated and lie on the margins most directly connected to the model's route-shipping channel. By contrast, the insurance coefficient is small and not statistically significant, suggesting that the observed adjustment is not mainly an insurance-premium response.

The other crime measures point in the same broad direction although unevenly. Homicide exposure is positive and statistically significant for total costs, with a coefficient of 0.09660, but it is not significant for fuel or insurance. The bundled crime exposure is positive for total costs and fuel expenses, with coefficients of 0.02855 and 0.02400, significant at the five and ten percent levels, respectively; the insurance estimate is again small and insignificant. Thus, the cost table supports the central claim of the paper: crime exposure raises transport operating costs. The evidence is driven most cleanly by extortion and, secondarily, by the bundled measure, rather than by every crime category.

A separate caveat applies to transporter-robbery exposure. In 2018, transporter robbery was recoded from a local to a federal crime category. In our municipal series, this institutional change appears to generate an artificial decline in municipal transporter-robbery counts after 2018. Because our exposure measure relies on municipality-year crime counts, the transporter-robbery coefficients are contaminated by this reporting break. We therefore report them for transparency, but we do not interpret them as evidence for the main mechanism. This is especially important because the significant cost coefficients are economically tiny and have the opposite sign from the route-cost channel.

Table 4: Firm-level EAT regressions: Series 2013, shipment and activity outcomes

	Log shipped value				Log total distance				Log total trips				Log total tons			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
Transporter-robbery exposure	-0.00001				-0.00015***				-0.00001				-0.00003			
	(0.00002)				(0.00001)				(0.00001)				(0.00002)			
Extortion exposure		-0.01565**				0.00637***				0.00579***				-0.01152***		
		(0.00645)				(0.00193)				(0.00164)				(0.00346)		
Homicide exposure			0.00256				0.09133				-0.07391				0.19021*	
			(0.32725)				(0.08879)				(0.08060)				(0.09923)	
Crime exposure				-0.07496				0.01358				0.03003				0.08424**
				(0.08265)				(0.02772)				(0.02988)				(0.03365)
Observations	505	525	529	529	775	800	805	805	775	800	805	805	775	800	805	805
Firm FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Firm IDs	236	244	245	245	303	314	314	314	303	314	314	314	303	314	314	314
Years	4	4	4	4	4	4	4	4	4	4	4	4	4	4	4	4

Notes: Each column is a separate firm-level regression with firm and year fixed effects; only the row for the exposure included in that column is populated. The displayed regressors are firm-year crime exposure measures built in the INEGI lab by merging each firm's reported top origin-destination revenue shares to OD-year route-crime exposure, then averaging over matched OD pairs and renormalizing over the matched share. OD-year route exposure is based on municipal crimes normalized by road kilometers along the relevant routes and expressed as a percent-change-from-base shift. The single-crime rows report transporter robbery, extortion, and homicide exposures; the bundled row reports the crime-bundle exposure available in the frozen lab return. Standard errors in parentheses are clustered at the firm-series level. All numeric coefficient and standard-error entries are reported to five decimals. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 4 shows how firms adjust on shipment and operating margins. Extortion again provides the most coherent pattern. A 100 percent increase in extortion exposure is associated with lower shipped value, higher total distance, more trips, and lower tons shipped. These estimates are consistent with the interpretation that extortion raises the cost of serving exposed shipment portfolios: firms travel farther and complete more trips, while the value and tonnage moved through those portfolios fall. This combination is consistent with a story of rerouting, fragmentation of shipments, delays, or reduced service on exposed corridors.

The remaining activity estimates are less systematic. Homicide exposure is imprecise for shipped value, distance, and trips, and positive for total tons at the ten percent level. The bundled crime exposure is positive and statistically significant only for total tons; its effects on shipped value, distance, and trips are not statistically significant. Transporter robbery again delivers very small coefficients, with a significant negative estimate only for total distance, and should be interpreted cautiously for the reporting reasons described above. Taken together, Tables 3 and 4 indicate that the firm-level evidence is not a generic "crime increases everything" result. It is a cost-side and operations-side pattern, with extortion as the crime category that most clearly maps into the transport-cost mechanism.

Table 5: Yearly crop-level price regressions: crime exposure

	Log price							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Transporter-robbery exposure	0.00001				0.00016^{***}			
	(0.00001)				(0.00005)			
Extortion exposure		0.00060				0.00278		
		(0.00185)				(0.00357)		
Homicide exposure			0.01127^{**}				0.00703	
			(0.00524)				(0.00889)	
Crime exposure				0.01186^{**}				0.00596
				(0.00490)				(0.00732)
Observations	20,419	21,503	21,708	21,708	20,419	21,503	21,708	21,708
Outcome mean (log price)	2.488	2.483	2.487	2.487	2.488	2.483	2.487	2.487
Crop x year FE	✓	✓	✓	✓				
Origin x crop FE	✓	✓	✓	✓				
Destination x crop FE	✓	✓	✓	✓				
Origin x crop x year FE					✓	✓	✓	✓
Destination x crop x year FE					✓	✓	✓	✓
OD pair FE					✓	✓	✓	✓
Two-way cluster SE	O,D	O,D	O,D	O,D	O,D	O,D	O,D	O,D

Notes: The dependent variable is the yearly mean of monthly log SNIIM modal wholesale prices, collapsed to crop-origin-destination-year cells by first averaging monthly log prices at the product-origin-destination-year level and then averaging across product varieties within crop-origin-destination-year cells. Exposure variables are yearly percent-change flownet route exposures relative to the 2011–2012 annual route-level base. At the OD-year level, monthly SESNSP crimes are summed to municipality-year cells, normalized by municipal road kilometers, averaged over the interior municipalities crossed by each route, and aggregated across alternative routes with inverse travel-time weights. The crop-state-origin bridge uses SIAP production shares to average municipality-origin OD exposure to SNIIM crop-origin-state cells. Each column is a separate OLS regression; only the row for the exposure included in that column is populated. Standard errors in parentheses are clustered two-way by origin state and destination market. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

6.2 Wholesale price evidence from SNIIM

We next ask whether route-level crime exposure is reflected in wholesale agricultural prices. The SNIIM regressions use a crop–origin–destination–year panel constructed by averaging monthly log modal prices and then averaging across product varieties within crop cells. The table reports two fixed-effect designs. Columns (1)–(4) include crop-year, origin-crop, and destination-crop fixed effects. This specification exploits cross-destination price variation in SNIIM: intuitively, it compares the price of the same crop from a given producing origin across different destination markets that are reached through corridors with different crime exposure, while absorbing common crop-year shocks and permanent origin-crop and destination-crop differences.

Columns (5)–(8) are more demanding. They add origin-crop-year fixed effects, destination-crop-year fixed effects, and OD-pair fixed effects, so identification comes from changes in route exposure within OD pairs after absorbing origin-side supply shocks, destination-side demand shocks, and permanent corridor characteristics.

Table 5 shows that the price evidence is positive but noisier than the firm-level evidence. In the cross-destination specification, homicide exposure and the bundled crime exposure are positive and statistically significant at the five percent level, with coefficients of 0.01127 and 0.01186. These magnitudes imply that a 100 percent increase in exposure is associated with roughly 1.1–1.2 percent higher wholesale prices. Extortion exposure is small and statistically insignificant in this price table, and transporter-robbery exposure is essentially zero in the same fixed-effect design.

The saturated specifications in columns (5)–(8) point in the same direction but with less precision. The bundled coefficient falls to 0.00596 and is not statistically significant; homicide remains positive but insignificant; and extortion remains positive but imprecise. Transporter-robbery exposure is positive and statistically significant in column (5), but the coefficient is only 0.00016. Given both the small magnitude and the reporting break in transporter robbery after 2018, this estimate is not central to our interpretation. We therefore read the SNIIM evidence as suggestive pass-through from route crime exposure to delivered prices, strongest in the across-destination variation emphasized by the data. Given the estimates, we conclude that there is a noisy crime-to-price pass-through with an elasticity between the range 0.6% and 1.2%.

6.3 Matched EAT–ENVE evidence

In addition to the reduced-form analysis, we provide complementary descriptive evidence of how crime affects firms. We do so by analyzing a subset of transportation firms that appear in both the transport EAT and victimization EAT surveys. While the EAT reveals how exposure is associated with operating outcomes, ENVE records whether the firm suffered any type of crime and the direct economic burden of victimization and prevention. The matched EAT-ENVE sample contains 237 firm-year observations corresponding to 154 firms. Of those firms, 84 report at least one targeted crime and 70 do not, so 54.5 percent of matched firms are victimized.

Table 6 reports the ENVE descriptive statistics for the matched sample. The available ratio summaries are reported for victimized matched firms, which makes the table especially useful for quantifying the intensive margin of crime costs. The distributions are highly right-skewed. Among victimized firms, average targeted losses are 2.82 million, average prevention spending is 1.03 million, and average total crime cost reaches 3.90 million. Medians are much smaller, which indicates that the burden is concentrated in the upper tail rather than spread evenly across firms.

Table 6: ENVE matched sample: victimization, losses, and prevention

Metric	Firms	Mean	Median	Min	Max	P25	P75
Panel A. Victimization and incidents							
Victim extortion (%)	82	16	0	0	100	0	0
Victim goods in transit (%)	84	53	100	0	100	0	100
Victim vehicle theft (%)	82	84	100	0	100	100	100
Number of extortion incidents	14	3	1	1	12	1	3
Number of goods-in-transit incidents	49	2	1	1	19	1	2
Number of vehicle theft incidents	68	3	2	1	24	1	3
Panel B. Economic losses							
Conservative targeted losses (MXN)	60	3,216,695	1,800,000	10,000	22,000,000	600,000	3,000,000
Total targeted losses (MXN)	69	2,816,202	1,000,000	0	22,000,000	277,500	2,850,000
Goods-transit proxy losses (MXN)	47	8,155	0	0	295,000	0	0
Vehicle theft losses (MXN)	60	3,216,695	1,800,000	10,000	22,000,000	600,000	3,000,000
Losses / total costs (%)	69	113	1	0	9,231	0	2
Losses / total revenue (%)	69	58	1	0	4,800	0	1
Panel C. Prevention spending							
Prevention spending (MXN)	67	1,028,642	250,000	10,000	30,000,000	80,000	500,001
Prevention / total costs (%)	67	34	0	0	2,769	0	1
Prevention / total revenue (%)	67	18	0	0	1,440	0	0
Panel D. Total crime burden							
Total crime cost (losses + prevention, MXN)	56	3,901,903	2,000,000	18,000	30,085,000	591,250	4,196,250
Total crime cost / total costs (%)	56	174	1	0	12,000	0	3
Total crime cost / total revenue (%)	56	90	1	0	6,240	0	3

Notes: The table summarizes firms observed in both the EAT transport survey and the ENVE victimization survey, separately for firms reporting a crime and firms not reporting a crime. Rows are grouped into victimization and incident counts, economic losses, prevention spending, and total crime burden. Columns report the number of firms represented in each row and the mean, median, minimum, maximum, 25th percentile, and 75th percentile of the row variable. Monetary variables are in Mexican pesos. Victimization indicators and ratios to total costs or total revenue are reported as percentages. Total crime burden is the sum of economic losses and prevention spending.

From table 6 we can establish four main descriptive facts. In a given year, the typical firm: i) suffers extortion, cargo robbery, and vehicle theft with 0.16, 0.53, and 0.84 probability, respectively; ii) suffers between 2 and 3 incidents of extortion, cargo robbery, and vehicle theft— with the maximum counts being 12, 19, and 24 incidents, respectively; iii) that yearly direct economic losses represent 113% of total costs and 58% of total revenue; and that iv) crime-prevention spending represents 34% of total costs and 18% of total revenue. This evidence shows that the burden of crime for a typical firm can be substantial, potentially an order of magnitude larger than regular taxes.

7 Calibration

This section describes how we bring the model in Section 3 to the data. The exercise is intentionally static. The time subscripts in the model are useful for motivating the empirical specifications, but the counterfactuals compare two steady environments: the observed economy with crime as a route-level friction and a counterfactual economy in which that friction is removed.

7.1 Geography and Route Choice

Origins, destinations, and goods. The quantitative geography is built from Mexican municipalities. Origins are municipalities that produce agricultural goods. Destinations are the market municipalities that appear in SNIIM price records, which give the abstract destination index d its empirical interpretation in the quantification. Destination expenditure is then disciplined by local income in the model rather than treated as an exogenous market-size shifter. Goods are the agricultural products for which we can connect SNIIM prices to production shares.

Routes. For every origin–destination pair, we construct a set of feasible driving routes on the road network. In the current implementation we keep the four lowest-cost alternatives for each pair. These alternatives are meant to capture the fact that shipments are not tied to a single shortest path. When one corridor becomes more expensive because of crime, intermediaries may shift some traffic toward other feasible corridors.

The non-crime component of route costs is measured by the distance of the route in kilometers and enters in logs. Thus the route-level transport wedge used in the quantitative exercise is

$$\tau^r_{odr} = \exp(\delta \, \log(\mathrm{km}_{odr}) + \phi \, \kappa_{odr}) \, , \tag{34}$$

where km_{odr} is route distance in kilometers and κ_{odr} is the normalized route crime index described below. We set $\delta = 0.25$. This value treats distance as a trade-cost elasticity rather than as a raw gravity coefficient. It lies above the clean route-based estimate in Donaldson (2018), who recovers trade costs from price gaps and route-level effective distance, and close to the freight-cost distance elasticities around 0.3 summarized by Anderson and van Wincoop (2004). In robustness exercises we will vary this parameter over the range suggested by these benchmarks. The important normalization is that δ multiplies log kilometers. If instead distance entered in kilometers, the comparable coefficient would be a small semi-elasticity per kilometer rather than 0.25.

We currently set the raw route-quality shifter equal to one for every active route. The CES weights η_{odr} are therefore equal-normalized within each origin–destination pair, so that the route aggregator does not mechanically make pairs with more alternatives cheaper simply because more routes are available.

7.2 Crime Exposure

Municipality crime intensity. The crime object entering Eq. (34) is designed to keep the same route-based logic as the empirical Bartik construction, without importing the full dynamic Bartik object into the static model. The model treats crime as a scalar route friction. On the data side, we construct this scalar as total annual municipal crime counts across the transport-relevant categories used in the paper. Because municipalities differ sharply in the amount of road infrastructure they contain, we normalize annual counts by municipal road kilometers. This gives a crime intensity measured as incidents per road kilometer.

Route aggregation. For each route, we identify the municipalities crossed by the route and exclude the origin and destination municipalities. The exclusion keeps the route measure focused on crime encountered along the corridor rather than crime at the shipment endpoints. The raw route crime measure is the mean of average annual crime intensity across the interior municipalities crossed by the route:

$$\widetilde{\kappa}_{odr} = \frac{1}{|\mathcal{M}^I_{odr}|} \sum_{m \in \mathcal{M}^I_{odr}} \overline{\left(\frac{\text{crime}_m}{\text{road km}_m}\right)}, \tag{35}$$

where $\mathcal{M}^I_{odr}$ is the set of interior municipalities crossed by route r, and the overbar denotes the average over available years.

The model-facing crime object is a unit-free index. We divide the raw route measure by its mean among routes with positive exposure:

$$\kappa_{odr} = \frac{\widetilde{\kappa}_{odr}}{\overline{\widetilde{\kappa}}_+}, \qquad \overline{\widetilde{\kappa}}_+ = \frac{1}{|\mathcal{R}_+|} \sum_{(o,d,r) \in \mathcal{R}_+} \widetilde{\kappa}_{odr}. \tag{36}$$

Thus $\kappa_{odr} = 1$ corresponds to an average positively exposed route, while $\kappa_{odr} = 0$ corresponds to no measured interior-route exposure. Removing crime in the counterfactual means setting this route crime index to zero, leaving the non-crime distance component and the route set fixed.

Calibrating the route-crime slope. The EAT cost regressions give the current calibration target for ϕ. Their crime regressors are percent-change Bartik exposures in proportional units: a one-unit increase is a 100 percent increase relative to the base exposure, not a one percentage point increase. For the draft benchmark we set $\phi = 0.02855$, the coefficient on the total-crime exposure in the total-cost specification. This choice maps the firm-side evidence most directly to the transport-cost channel in Eq. (34). It also keeps the quantitative model agnostic about the composition of crime: the counterfactual removes the scalar route-crime friction, not a separate set of category-level frictions.

7.3 Economic Fundamentals

Labor and population weights. Municipality labor size comes from the 2024 Economic Census. We use total employed workers to construct the fixed population weights $\bar{T}_o$. In the shipping-only benchmark the mobility parameter is set to zero, so these weights pin down the spatial distribution of labor. They are not merely starting values for an iterative migration process; with fixed population, they are the population allocation used in the equilibrium.

Productivity and land. The agricultural productivity terms $\bar{A}_o^k$ are calibrated from municipality-level census output together with crop production shares. The Economic Census gives total gross output and employment at the municipality level, while the agricultural data give crop-specific production shares. We allocate municipality output and labor across crops using those crop shares and then invert the Cobb–Douglas production function in Eq. (1). This inversion uses a calibration labor share of 0.75. In the current equilibrium computation, however, the production block sets $\mu = 1$ so that land does not introduce an additional fixed-factor loop before the draft counterfactual is disciplined. The same crop shares also provide the fixed agricultural support terms $\bar{H}_o^k$ and the baseline origin weights in destination demand.

This inversion is a practical first pass. It lets high-output, high-labor municipalities enter the model with stronger fundamentals, while preserving crop-specific heterogeneity from the agricultural production data. It does not yet use crop-specific census labor or crop-specific value added, which are not available in the current calibration inputs.

Demand and expenditure. Goods weights are based on aggregate crop production, and origin weights within each good are based on crop production shares. Destination expenditure is closed inside the model. The empirical destination geography has fewer nodes than the full municipality set, so we assign each municipality's income to a destination catchment. Let χ_{do} denote the share of municipality o's income assigned to destination d, with $\sum_d \chi_{do} = 1$. In the shipping-only benchmark, there are no non-tradables or land-rent feedbacks, so destination tradable expenditure is assigned local labor income:

$$E_d^T = \sum_o \chi_{do} w_o L_o. \tag{37}$$

If the destination set equals the set of municipalities, χ_{do} is the identity matrix and this reduces to $E_d^T = w_d L_d$. With land and non-tradables active, this object should be interpreted as the tradable share of total local factor income. This is the balanced-trade closure derived in Appendix B. In the current quantification, χ_{do} assigns each municipality to the nearest SNIIM destination coordinate. Equal SNIIM-market expenditure weights are therefore not a structural calibration target; they are only useful as a diagnostic or partial-equilibrium comparison.

Table 7: Current structural parameters

Parameter	Value	Role and current calibration
$R_{od}^{\max}$	4	Geography: Route alternatives kept for each origin–destination pair. Current route-choice build keeps four.
σ	5.0	Demand: Origin-variety elasticity within an agricultural good. Placeholder Armington value.
ρ	3.0	Demand: Elasticity across agricultural goods. Placeholder value.
ζ	4.0	Intermediation: Elasticity across alternative route services. Placeholder value for rerouting.
α	1.0	Preferences: Tradable expenditure share. Set to one in the route-shipping benchmark.
μ	1.0	Production: Agricultural labor share. Set to one; productivity inversion uses a 0.75 calibration share.
ξ	0.3	Production: Fixed-factor share in non-tradables. Inactive when $\alpha = 1$.
δ	0.25	Transport costs: Route-cost elasticity with respect to log kilometers. Set to 0.25 from the distance-elasticity literature.
ϕ	0.02855	Transport costs: Route-cost semi-elasticity with respect to the unit-free route crime index. Set to 0.02855 from the EAT total-cost coefficient on the total-crime percent-change Bartik.
ν	0.0	Population: Population mobility. Set to zero; population is fixed at Censo labor weights.

Notes: Defaults used in the current quantitative implementation. The benchmark lets route crime affect transport costs and holds population fixed. Placeholder values are not final estimates.

7.4 Parameters

Table 7 reports the current parameter values used in the quantitative implementation. The benchmark exercise isolates the route-shipping channel: crime affects route-level transport costs through ϕ, distance affects route costs through δ, and intermediaries can substitute across feasible routes according to ζ. Population is fixed by setting $\nu = 0$. This benchmark is the channel most directly disciplined by the firm-side cost evidence and the route-exposure design.

8 Counterfactuals

Welfare cost of crime. We use the calibrated model to ask how much route crime matters through the shipping channel alone. The benchmark economy is the calibrated Mexican equilibrium with observed route crime entering transport costs through

$$\tau_{odr} = \exp(\delta \log(\mathrm{km}_{odr}) + \phi\kappa_{odr})\,. \tag{38}$$

Table 8: Shipping-crime counterfactual

	Change (percent)
Panel A. Welfare	
Population-weighted welfare, $V_o = w_o/P_o$	+1.35
P10 municipality welfare	+0.72
P90 municipality welfare	+5.02
Panel B. Trade costs and prices	
Mean bilateral trade cost, τ_{od}	-2.91
Mean route-level trade cost, τ_{odr}	-2.90
Mean destination tradable price index, P_d^T	-6.46
Panel C. Nominal OD expenditure	
Mean nominal OD expenditure change	-2.94

Notes: The table reports percent changes from the calibrated benchmark to the no-crime counterfactual. The benchmark is the shipping-only model with observed route crime entering route costs through $\tau_{odr} = \exp\{\delta \log(\mathrm{km}_{odr}) + \phi\kappa_{odr}\}$. The counterfactual sets $\kappa_{odr} = 0$ for every route and leaves routes, distance, fundamentals, and non-shipping crime channels fixed. Welfare is real income $V_o = w_o/P_o$ because amenity crime is shut down in the benchmark. Population-weighted welfare uses baseline municipality labor weights; the P10 and P90 rows are percentiles across municipality-level welfare changes. Mean route-level trade costs use active route alternatives. The nominal OD expenditure row averages percentage changes across origin-destination pairs with positive baseline expenditure; flows are nominal expenditure in the model numeraire.

The no-crime counterfactual sets $\kappa_{odr} = 0$ for every route. We keep the route network, distance, route-quality weights, productivities, demand weights, and destination catchments fixed. In this exercise, crime does not directly affect farm productivity, non-tradable productivity, amenities, or demand shifters. Those channels are part of the general model, but we leave them inactive because we do not yet have enough empirical discipline to calibrate their crime semi-elasticities. The counterfactual should therefore be read as the welfare cost of route crime operating through shipping costs.

Table 8 reports the resulting equilibrium changes. Welfare is the real-income object from the model, $V_o = w_o/P_o$ in the shipping-only benchmark. Population-weighted

Figure 8: Municipality Welfare Changes from Removing Route Crime

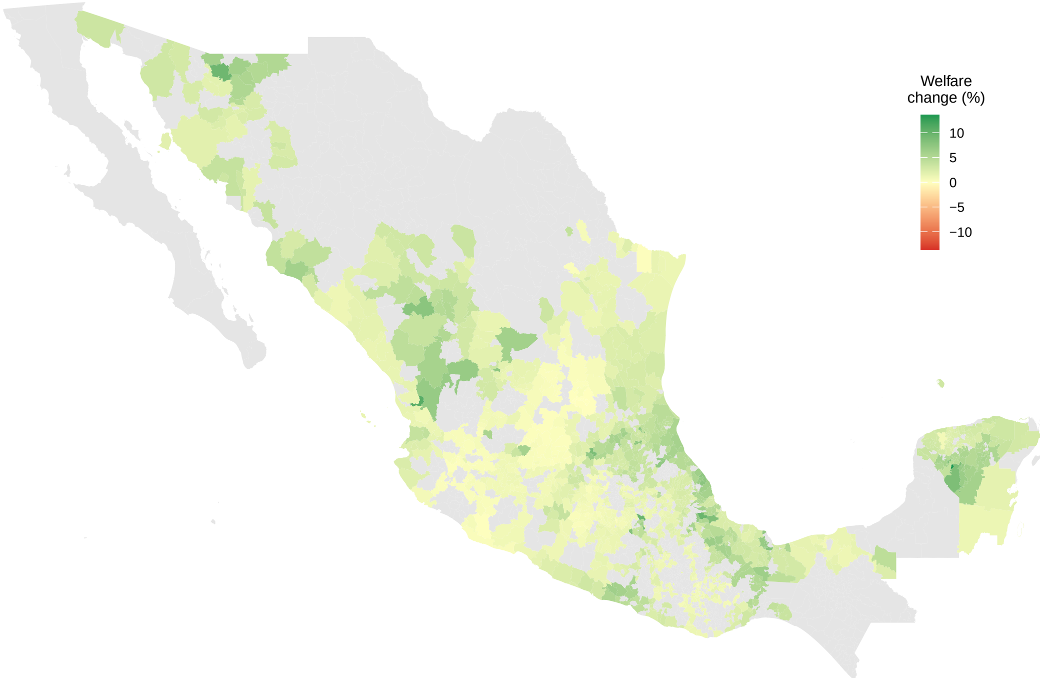

Notes: The map reports percent changes in welfare, $V_o = w_o/P_o$, for agricultural producing origin municipalities in the calibrated shipping-only benchmark. Municipalities not included in the current structural origin universe are shown in grey.

welfare rises by 1.35 percent when route crime is removed. The gains are not concentrated only in the upper tail: the 10th percentile municipality gains 0.72 percent, while the 90th percentile gains 5.02 percent. These numbers reflect both the direct fall in delivered prices and the equilibrium response of wages and sourcing shares.

The mechanism is visible in the price and trade-cost rows. Removing route crime lowers mean bilateral trade costs by 2.91 percent and mean active route-level costs by 2.90 percent. Destination tradable price indices fall by 6.46 percent, which is larger than the average trade-cost change because cheaper routes change the relative attractiveness of origins and goods in the CES demand system.

The nominal OD expenditure row should be interpreted as a reallocation statistic rather than as a real aggregate-trade quantity. It averages percentage changes in nominal origin–destination expenditure across pairs with positive baseline expenditure, and those flows are measured in the model numeraire. A decline in this statistic can coexist with higher welfare because the counterfactual lowers delivered prices, changes sourcing shares, and lets equilibrium wages and local expenditure adjust. The welfare object is real income, while this row summarizes how nominal spending is redistributed across OD links.

Spatial effects. Figure 8 maps the municipality-level welfare changes underlying Panel A. This figure is defined over agricultural producing origins in the current structural origin universe. Most modeled producing municipalities gain from removing the route-crime component, while municipalities outside this origin universe are shown in grey. The gains are not spatially uniform. The largest visible gains among producing municipalities appear in parts of Veracruz, Durango, Sonora, Sinaloa, Hidalgo, and Yucatán, consistent with the idea that route-crime reductions matter most where exposed corridors account for an important part of delivered costs.

9 Conclusion

This paper studies shows how crime—extortion, robbery, homicides—across shipping routes affects firm outcomes and agricultural goods prices in the context of Mexico. Exploiting detailed confidential transportation-firm microdata, origin–destination wholesale agricultural prices, and administrative crime records we show that doubling our crime exposure i) increases transportation firms' operational costs, with estimates ranging from 0.4% to 9.6%, depending on the type of crime; and ii) increases agricultural prices between 0.6% and 1.1%, although estimates are noisy.

The contribution of the paper is to demonstrate empirically a crime-to-cost and crime-to-price channel, and to quantify the aggregate welfare effects using a spatial trade model with transport intermediaries. We show that counterfactually removing the route-crime friction lowers mean bilateral trade costs by about 2.9%, lowers destination tradable price indices by about 6.5%, and raises average welfare by about 1.4%.

The immediate policy implication of the paper is that road security should be treated as a form of trade facilitation. Policy could prioritize economically central shipping routes, particularly corridors with high exposure to crime and limited route substitutes, rather than focusing exclusively on crime at origins or destinations. In this sense, reducing predation on transport networks can operate like an improvement in domestic infrastructure by lowering the effective cost of connecting markets.

A natural next step is to discipline the broader channels that are present in the model but inactive in the benchmark counterfactual. Confidential CPI microdata, together with land and housing price data, would allow us to study how crime affects non-tradable and service prices, local costs of living, and the spatial allocation of economic activity. This would move the quantification from the current shipping-only exercise toward a fuller measure of the real aggregate effect of crime, taking into account trade, local productivity, and prices across tradables, non-tradables, and services.

Future work should also use the route structure to evaluate targeted policy experiments. Natural extensions are to remove crime only on the highest-exposure corridors, or compare route-based crime reductions to origin-based reductions. These exercises would help distinguish policies that reduce the measured friction where it is most severe from policies that generate the largest general-equilibrium gains.

A Model extensions

The main text studies crime as a route-level shipping friction. This appendix records additional margins that can be embedded in the same equilibrium environment but are not used in the current quantitative exercise. These extensions are useful for organizing future versions of the model, where crime may affect local production conditions, amenities, or sourcing preferences in addition to transport costs.

Farm-gate productivity. Crime in producing origins may reduce agricultural productivity by disrupting planting, harvesting, storage, and local input markets. One way to introduce this channel is

$$A^k_{ot} = \bar{A}^k_{ot} \exp(-\phi_A \kappa_{ot}) \, , \tag{39}$$

where κ_{ot} is crime at the source and $\phi_A > 0$ is the productivity semi-elasticity with respect to local crime.

Non-tradable productivity. Crime may also lower productivity in local services, warehousing, distribution, and other non-tradables. This can be written as

$$A^{NT}_{ot} = \bar{A}^{NT}_{ot} \exp(-\phi_{NT} \kappa_{ot}) \, , \tag{40}$$

with $\phi_{NT} > 0$. This channel would raise local non-tradable prices and feed back into the cost of living.

Amenities. Crime may make locations less desirable to workers and firms through fear, harassment, or a deterioration in public space. A reduced-form amenity channel is

$$B_{ot} = \bar{B}_{ot} \exp(-\phi_B \kappa_{ot}) \, , \tag{41}$$

where $\phi_B > 0$. In a model with mobile labor, lower amenities would affect population and wages through the mobility condition in Eq. (18).

Demand avoidance. Buyers may avoid sourcing from crime-exposed supply chains beyond the direct cost increase, for example because reliability falls or because certain origins become reputationally risky. This can be captured through the Armington demand shifter:

$$\beta^k_{od} = \bar{\beta}^k_{od} \exp(-\psi\, \kappa_{odt}) \,, \tag{42}$$

where κ_{odt} is a bilateral supply-chain crime exposure index and $\psi > 0$ governs the strength of buyer avoidance.

B Balanced-Trade Closure and Wage Determination

This appendix derives the expenditure closure used in the quantitative model. The key point is that the destination expenditure object is not an independent primitive in a closed general-equilibrium trade model. It is local income earned by residents of the destination location.

B.1 Expenditure Shares

For each destination location d and good k, the CES price index across origins is

$$P^k_{dt} = \left[\sum_o \beta^k_{od} \left(p^k_{odt} \right)^{1-\sigma} \right]^{\frac{1}{1-\sigma}} . \tag{43}$$

The expenditure share of destination d on origin o for good k is therefore

$$S^k_{odt} = \beta^k_{od} \left(\frac{p^k_{odt}}{P^k_{dt}} \right)^{1-\sigma} . \tag{44}$$

Across goods, the tradable price index is

$$P^T_{dt} = \left[\sum_k \alpha^k_d \left(P^k_{dt} \right)^{1-\rho} \right]^{\frac{1}{1-\rho}} , \tag{45}$$

and the expenditure share of destination d on good k is

$$s^k_{dt} = \alpha^k_d \left(\frac{P^k_{dt}}{P^T_{dt}} \right)^{1-\rho} . \tag{46}$$

If total tradable expenditure in destination d is E_{dt}^T, then nominal expenditure on origin o and good k is

$$X_{odt}^k = S_{odt}^k s_{dt}^k E_{dt}^T. \tag{47}$$

This is the gravity block. It determines expenditure shares conditional on prices and total destination expenditure. It does not by itself determine E_{dt}^T.

B.2 Local Income and Balanced Trade

Let I_{dt} denote total income earned by residents of destination location d. With Cobb–Douglas preferences over tradables and non-tradables, tradable expenditure is

$$E_{dt}^T = \alpha I_{dt}. \tag{48}$$

Balanced trade closes the model by setting destination expenditure equal to local income rather than treating it as an external demand shifter.

In the shipping-only benchmark, the production block is labor-only and non-tradables are inactive. Income is therefore labor income:

$$I_{dt} = w_{dt} L_{dt}, \qquad E_{dt}^T = \alpha w_{dt} L_{dt}. \tag{49}$$

When $\alpha = 1$, this reduces to $E_{dt}^T = w_{dt} L_{dt}$.

When land is active in agricultural production, total income also includes payments to the fixed factor. Under constant returns and perfect competition, crop revenue in origin o and good k is

$$R_{ot}^k = \sum_d X_{odt}^k. \tag{50}$$

Labor receives share μ of crop revenue and land receives share $1 - \mu$:

$$w_{ot} L_{ot}^k = \mu R_{ot}^k, \qquad r_{ot}^k \bar{H}_o^k = (1 - \mu) R_{ot}^k. \tag{51}$$

If local residents own the fixed factor locally, total local agricultural income is the sum of labor and land income:

$$I_{ot}^A = \sum_k R_{ot}^k. \tag{52}$$

If land rents are rebated nationally or owned outside the location, the income closure must specify that ownership rule. The shipping-only benchmark avoids this additional incidence choice by setting $\mu = 1$.

B.3 Wage Fixed Point

The wage fixed point follows from combining expenditure shares with labor-market clearing. In the labor-only benchmark, origin revenue is

$$R_{ot} = \sum_{d,k} S^k_{odt} s^k_{dt} w_{dt} L_{dt}. \tag{53}$$

Labor-market clearing requires

$$w_{ot} L_{ot} = R_{ot}. \tag{54}$$

Equations (53)–(54) define the wage system:

$$w_{ot} = \frac{1}{L_{ot}} \sum_{d,k} S^k_{odt}(\mathbf{w}) s^k_{dt}(\mathbf{w}) w_{dt} L_{dt}, \tag{55}$$

where the shares depend on wages through farm-gate costs and delivered prices. Because only relative wages matter, one wage is chosen as the numeraire.

With land active, labor-market clearing instead uses labor's revenue share:

$$w_{ot} L^T_{ot} = \mu \sum_{d,k} X^k_{odt}, \tag{56}$$

and the farm-gate cost depends on the equilibrium labor-to-land ratio,

$$c^k_{ot} = \frac{w_{ot}}{\mu A^k_{ot}} \left(\frac{L^k_{ot}}{\bar{H}^k_o} \right)^{1-\mu}. \tag{57}$$

This is why the richer model requires a subsidiary fixed point over farm-gate costs. The benchmark counterfactual sets $\mu = 1$, so this loop is shut down and the balanced-trade closure reduces to Eq. (55).

C Additional figures

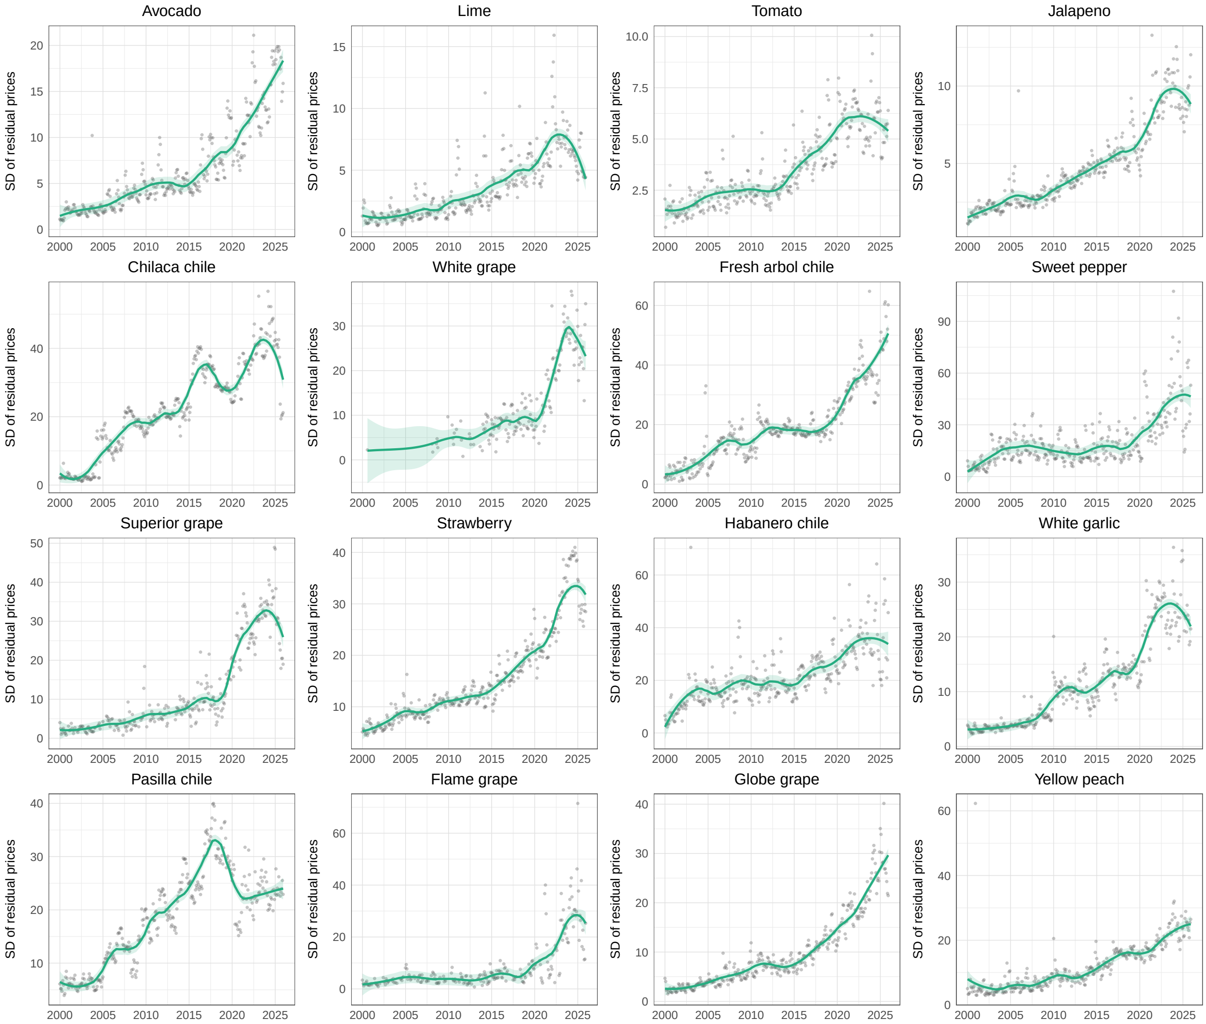

Figure 9: Cross-destination residual price dispersion over time

Notes: Each point is the product-month standard deviation of residual frequent wholesale prices across destination markets for selected products with rising dispersion over time. Residual prices remove product-specific calendar-month and year means and add back the product mean. Product-months require at least three destination markets and ten price observations; lines are local-polynomial smooths with pointwise confidence bands. Prices are nominal Mexican pesos per kilogram.

D Additional EAT Regression Tables

This appendix reports the firm-level EAT regression tables that are not shown in the main text. The first table keeps Serie 2013 but switches to additional firm-scale outcomes. The remaining tables use the pooled series sample, combining series 2008, 2013, and 2018, for the cost, activity, and additional outcome blocks.

Table 9: Firm-level EAT regressions: Series 2013, additional outcomes

	Log employment				Log total vehicles			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Transporter-robbery exposure	0.00000				-0.00004***			
	(0.00001)				(0.00001)			
Extortion exposure		0.00189				0.00658**		
		(0.00236)				(0.00270)		
Homicide exposure			-0.00478				0.00782	
			(0.05804)				(0.04413)	
Crime exposure				0.03041				0.01268
				(0.03809)				(0.02371)
Observations	735	759	764	764	774	799	804	804
Firm FE	✓	✓	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓	✓	✓
Firm IDs	294	305	305	305	303	314	314	314
Years	4	4	4	4	3	3	3	3

Notes: Each column is a separate firm-level regression with firm and year fixed effects; only the row for the exposure included in that column is populated. The displayed regressors are firm-year crime exposure measures built in the INEGI lab by merging each firm's reported top origin-destination revenue shares to OD-year route-crime exposure, then averaging over matched OD pairs and renormalizing over the matched share. OD-year route exposure is based on municipal crimes normalized by road kilometers along the relevant routes and expressed as a percent-change-from-base shift. The single-crime rows report transporter robbery, extortion, and homicide exposures; the bundled row reports the crime-bundle exposure available in the frozen lab return. Standard errors in parentheses are clustered at the firm-series level. All numeric coefficient and standard-error entries are reported to five decimals. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 10: Firm-level EAT regressions: Pooled series 2008, 2013, and 2018, cost outcomes

	Log total costs				Log fuel expense				Log insurance expense			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Transporter-robbery exposure	-0.00003***				-0.00003***				-0.00003***			
	(0.00001)				(0.00001)				(0.00001)			
Extortion exposure		0.00309**				0.00609***				-0.00099		
		(0.00147)				(0.00101)				(0.00378)		
Homicide exposure			0.03537				0.07164*				-0.00158	
			(0.03060)				(0.03700)				(0.09432)	
Crime exposure				0.01099*				0.01034				-0.02625
				(0.00665)				(0.00864)				(0.01750)
Observations	2,272	2,360	2,368	2,368	1,985	2,059	2,066	2,066	2,048	2,122	2,130	2,130
Firm FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Firm IDs	988	1,024	1,025	1,025	801	829	830	830	917	949	950	950
Years	8	8	8	8	7	7	7	7	8	8	8	8

Notes: Each column is a separate firm-level regression with firm and year fixed effects; only the row for the exposure included in that column is populated. The displayed regressors are firm-year crime exposure measures built in the INEGI lab by merging each firm's reported top origin-destination revenue shares to OD-year route-crime exposure, then averaging over matched OD pairs and renormalizing over the matched share. OD-year route exposure is based on municipal crimes normalized by road kilometers along the relevant routes and expressed as a percent-change-from-base shift. The single-crime rows report transporter robbery, extortion, and homicide exposures; the bundled row reports the crime-bundle exposure available in the frozen lab return. Standard errors in parentheses are clustered at the firm-series level. All numeric coefficient and standard-error entries are reported to five decimals. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 11: Firm-level EAT regressions: Pooled series 2008, 2013, and 2018, shipment and activity outcomes

	Log shipped value				Log total distance				Log total trips				Log total tons			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
Transporter-robbery exposure	-0.00004				-0.00006**				-0.00002				-0.00001			
	(0.00017)				(0.00003)				(0.00003)				(0.00004)			
Extortion exposure		-0.01643**				0.00647***				0.00427				-0.01770***		
		(0.00824)				(0.00246)				(0.00267)				(0.00538)		
Homicide exposure			-0.41002				0.06225				-0.00043				0.03074	
			(0.26857)				(0.05800)				(0.06206)				(0.09736)	
Crime exposure				-0.07878				0.01374				0.02361				0.04172**
				(0.05062)				(0.01249)				(0.01555)				(0.01918)
Observations	1,314	1,363	1,369	1,369	2,272	2,360	2,368	2,368	2,272	2,360	2,368	2,368	2,272	2,360	2,368	2,368
Firm FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Firm IDs	631	652	654	654	988	1,024	1,025	1,025	988	1,024	1,025	1,025	988	1,024	1,025	1,025
Years	7	7	7	7	8	8	8	8	8	8	8	8	8	8	8	8

Notes: Each column is a separate firm-level regression with firm and year fixed effects; only the row for the exposure included in that column is populated. The displayed regressors are firm-year crime exposure measures built in the INEGI lab by merging each firm's reported top origin-destination revenue shares to OD-year route-crime exposure, then averaging over matched OD pairs and renormalizing over the matched share. OD-year route exposure is based on municipal crimes normalized by road kilometers along the relevant routes and expressed as a percent-change-from-base shift. The single-crime rows report transporter robbery, extortion, and homicide exposures; the bundled row reports the crime-bundle exposure available in the frozen lab return. Standard errors in parentheses are clustered at the firm-series level. All numeric coefficient and standard-error entries are reported to five decimals. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 12: Firm-level EAT regressions: Pooled series 2008, 2013, and 2018, additional outcomes

	Log employment				Log total vehicles			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Transporter-robbery exposure	-0.00001				0.00006***			
	(0.00000)				(0.00002)			
Extortion exposure		0.00093				0.00468**		
		(0.00193)				(0.00216)		
Homicide exposure			-0.02872				-0.04355	
			(0.04942)				(0.03959)	
Crime exposure				0.01120				-0.00319
				(0.01327)				(0.01008)
Observations	2,138	2,222	2,230	2,230	2,271	2,359	2,367	2,367
Firm FE	✓	✓	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓	✓	✓
Firm IDs	939	974	975	975	988	1,024	1,025	1,025
Years	8	8	8	8	8	8	8	8

Notes: Each column is a separate firm-level regression with firm and year fixed effects; only the row for the exposure included in that column is populated. The displayed regressors are firm-year crime exposure measures built in the INEGI lab by merging each firm's reported top origin-destination revenue shares to OD-year route-crime exposure, then averaging over matched OD pairs and renormalizing over the matched share. OD-year route exposure is based on municipal crimes normalized by road kilometers along the relevant routes and expressed as a percent-change-from-base shift. The single-crime rows report transporter robbery, extortion, and homicide exposures; the bundled row reports the crime-bundle exposure available in the frozen lab return. Standard errors in parentheses are clustered at the firm-series level. All numeric coefficient and standard-error entries are reported to five decimals. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

References

Treb Allen, David Atkin, Santiago Cantillo Cleves, and Carlos Eduardo Hernandez. The traveling trucker problem. *AEA Papers and Proceedings*, 114:334–339, 2024.

James E. Anderson and Eric van Wincoop. Trade costs. *Journal of Economic Literature*, 42(3):691–751, 2004.

David Atkin and Dave Donaldson. The role of trade in economic development. In Gita Gopinath, Elhanan Helpman, and Kenneth Rogoff, editors, *Handbook of International Economics*, volume 5, pages 1–59. Elsevier, 2022.

David Atkin and Amit K. Khandelwal. How distortions alter the impacts of international trade in developing countries. *Annual Review of Economics*, 12(1):213–238, 2020.

Kirill Borusyak, Peter Hull, and Xavier Jaravel. A practical guide to shift-share instruments. *Journal of Economic Perspectives*, 39(1):181–204, 2025.

Giulia Brancaccio, Myrto Kalouptsidi, and Theodore Papageorgiou. Geography, transportation, and endogenous trade costs. *Econometrica*, 88(2):657–691, 2020.

Zach Y. Brown, Eduardo Montero, Carlos Schmidt-Padilla, and Maria Micaela Sviatschi. Market structure and extortion: Evidence from 50,000 extortion payments. *Review of Economic Studies*, 92(3):1595–1624, 2025.

Melissa Dell. Trafficking networks and the mexican drug war. *American Economic Review*, 105(6):1738–1779, 2015.

Melissa Dell, Benjamin Feigenberg, and Kensuke Teshima. The violent consequences of trade-induced worker displacement in mexico. *American Economic Review: Insights*, 1(1):43–58, 2019.

Rafael Dix-Carneiro, Rodrigo R. Soares, and Gabriel Ulyssea. Economic shocks and crime: Evidence from the brazilian trade liberalization. *American Economic Journal: Applied Economics*, 10(4):158–195, 2018.

Dave Donaldson. Railroads of the raj: Estimating the impact of transportation infrastructure. *American Economic Review*, 108(4-5):899–934, 2018.

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

Douglas Gollin and Richard Rogerson. Productivity, transport costs and subsistence agriculture. *Journal of Development Economics*, 107:38–48, 2014.

Gaurav Khanna, Carlos Alberto Medina-Durango, Anant Nyshadham, Daniel Ramos-Menchelli, Jorge Andres Tamayo-Castano, and Audrey Tiew. Spatial mobility, economic opportunity, and crime. Working paper, 2025.

Benjamin A. Olken and Patrick Barron. The simple economics of extortion: Evidence from trucking in aceh. *Journal of Political Economy*, 117(3):417–452, 2009.

Sandra Sequeira and Simeon Djankov. Corruption and firm behavior: Evidence from african ports. *Journal of International Economics*, 94(2):277–294, 2014.

Sebastian Sotelo. Domestic trade frictions and agriculture. *Journal of Political Economy*, 128(7):2690–2738, 2020.