Upstream Market Power and the Price of Production Flexibility *

Qinshu Xue

This paper studies how upstream market power shapes downstream firms' production flexibility. Outsourcing is valuable when demand shifts move in-house production away from efficient capacity utilization, but suppliers can charge for this flexibility. I estimate a structural model of upstream pricing, downstream make-or-buy sourcing, and vehicle pricing using a new dataset linking vehicle models to transmission suppliers from 2009 to 2018. In a counterfactual economic bust, sourcing flexibility raises strategic-core downstream profit by $1.10 billion, but upstream suppliers reprice this flexibility, raising input prices by 40 to 52 percent of baseline markups and capturing $62 million. Under a stylized Aisin-monopoly counterfactual motivated by USMCA regional-content rules, monopoly repricing transfers $10 million under normal demand and $64 million during the bust, showing that upstream concentration makes the price of production flexibility more state-dependent. Supplier market power can make production flexibility more expensive precisely when downstream firms need it most.

Keywords: production flexibility, outsourcing, vertical relations, upstream market power

---

*Department of Economics, Shanghai Jiao Tong University. Email: qinshu.xue@sjtu.edu.cn. I am grateful to Panle Jia Barwick, Julieta Caunedo, Giulia Brancaccio, and Jörg Stoye for their invaluable advice, guidance, and support. I would like to thank Adam Deming, Ben Leyden, Shanjun Li, Seth Sanders, Luming Chen, Daniel Yi Xu, Nahim Bin Zahur, Hester Zhang, Bin Zhao, and participants at Cornell, Fudan, CUHK-Shenzhen, Renmin U, EARIE, AMES, NASMES, APIOC, Shufe-IO workshop for helpful comments. I have also greatly benefited from the conversations with industry experts, especially Alex Kearns and David Andrea. Financial support from the Sage Fellowship is gratefully acknowledged. All errors are my own.

1 Introduction

Many industries face substantial demand fluctuations, and these fluctuations shape how firms organize production across the supply chain. A downstream firm can produce inputs internally, preserving control over production but bearing the utilization consequences of its own capacity, or it can outsource production to specialized suppliers, gaining flexibility at the cost of paying upstream markups. This tradeoff is central to the economics of vertical relations.[1]

This paper studies how upstream market power shapes the price of downstream production flexibility when demand fluctuates. The analysis separates three channels. First, upstream market power creates a standard double marginalization channel, raising downstream marginal costs through input markups. Second, outsourcing provides production flexibility by allowing downstream firms to use an outside production option when demand movements shift in-house plants away from efficient capacity utilization. Third, demand shocks alter the value of outsourcing and shift the input-demand curve faced by upstream suppliers; upstream suppliers then adjust input prices, transmitting the shock back through the vertical chain. These channels reveal a role for upstream market power that is obscured when input prices are treated as exogenous: upstream suppliers can make downstream production flexibility more expensive precisely when demand shocks make that flexibility valuable.

The automobile industry lends itself to the analysis. It is highly volatile and is affected by the macroeconomic environment, ongoing technological disruption, and within-industry uncertainty about consumer tastes.[2] In addition, auto production involves more than ten thousand parts, and supply-chain management is central to automakers' business models. Automotive supply contracts make outsourcing a meaningful adjustment margin: unit prices are often specified for an initial delivery period, while exact quantities are less rigidly fixed. These contractual features allow downstream firms to adjust outsourced quantities without

[1] Lafontaine and Slade (2007) survey the theoretical and reduced-form empirical literature on vertical relations.

[2] According to Ramey and Vine (2006), the auto industry represents almost 5% of aggregate GDP in the United States and accounted for almost 25% of its variance over the past 40 years. Their paper shows that volatility declined in the early 2000s. Following their computation method, I find that the auto industry still accounts for more than 20% of the variation in aggregate GDP growth from 2004 through 2019.

frequently renegotiating supplier relationships.

For the empirical analysis, I focus on the outsourcing decision for transmissions, a core component of the powertrain system that accounts for about 7% of the cost of a car. The transmission market provides a suitable empirical setting for two reasons. First, unlike many auto parts that are now fully outsourced, transmissions are still produced both in-house and by independent suppliers, generating meaningful make-or-buy variation. Second, technological barriers and large engineering costs make the U.S. transmission industry highly concentrated. Aisin, ZF, and JATCO dominate the outsourced transmission market, serving more than 90% of outsourced transmissions. These features are increasingly relevant because mergers, acquisitions, and protective trade policies can further alter upstream market structure. In particular, regional-content policies such as the United States–Mexico–Canada Agreement may raise the value of qualifying local suppliers and increase the effective concentration of the upstream market. Despite technical differences across components and industries, the core mechanism can also apply to other settings in which upstream consolidation or policy-induced entry barriers change the price of downstream production flexibility.

To model the interaction between upstream transmission suppliers and downstream vehicle manufacturers, I consider a static three-stage game. In the first stage, upstream suppliers simultaneously post input prices, anticipating downstream sourcing responses and the distribution of demand and cost shocks. In the second stage, downstream firms choose sourcing intensities at the firm–transmission level, deciding how much of a transmission program to produce in-house and how much to outsource. In the third stage, product-level demand and marginal-cost shocks are realized, and downstream firms set final vehicle prices in a differentiated-products Bertrand game. The model delivers equilibrium sourcing probabilities, downstream vehicle prices, and upstream input prices, which are often treated as exogenous in previous work.

I estimate the model using a novel dataset that links transmission suppliers to individual vehicle products sold in the U.S. market from 2009 to 2018. I combine these supplier links with data on vehicle prices, sales, and characteristics. The linked supplier–vehicle data are central for identification because they reveal which products are produced in-house, which products are outsourced, and which upstream supplier serves each outsourced product.

The empirical strategy proceeds in four steps. First, I estimate a random coefficients demand system for differentiated vehicles. Demand estimation uses standard BLP-style instruments, differentiation instruments, and cost shifters, including local manufacturing wages at assembly sites. Second, I recover downstream marginal costs from the downstream vehicle-pricing first-order conditions. Third, I estimate the nonlinear in-house transmission cost function using variation in realized output, sourcing status, and equilibrium vehicle prices. The make-or-buy variation is informative about the relative cost of in-house production versus outsourcing, while output variation identifies how in-house costs change with capacity utilization. Fourth, I recover upstream supplier marginal costs by inverting the upstream pricing first-order conditions. This inversion uses the model-implied elasticity of expected outsourced demand with respect to upstream input prices.

A key empirical contribution is to construct demand fluctuations in a way that is consistent with the structural demand system. I recover product-level demand shocks from the random-coefficients demand model and use an empirical-Bayes residual decomposition to separate them into a common firm–transmission component and a product-specific component. The common component captures shocks relevant for the sourcing decision unit, while the product-specific component captures residual variation across products within a program. By shrinking noisy product-level residual moments toward a common distribution, this procedure produces a demand-shock process disciplined by the estimated demand system rather than by raw sales volatility.

The estimated in-house cost function is U-shaped. Average in-house transmission costs are high at low output because of poor capacity utilization, fall over an intermediate range because of scale economies, and rise again when output approaches capacity limits. There is also substantial heterogeneity in in-house production costs across downstream firms, consistent with observed sourcing patterns. Firms such as Daimler, which produce nearly all transmissions in-house, have among the lowest estimated in-house production costs. These estimates provide the structural basis for measuring the production-flexibility value of outsourcing.

In the counterfactual analysis, I first quantify the value of sourcing flexibility during a large negative demand shock. I simulate an economic bust by reducing the 2018 market size by one third. When strategic downstream firm–transmission pairs are forced to use the full in-house action, total welfare falls by $29.8 billion relative to the baseline, and strategic-core profit falls by 41 percent. Allowing these firms to adjust sourcing decisions substantially mitigates the shock. Holding upstream prices fixed, sourcing flexibility raises strategic-core profit by $1.10 billion, recovering 17.8 percent of the bust-induced strategic-core loss. The upstream sector therefore provides an economically important production-adjustment margin during downturns.

Upstream suppliers respond to the increased value of outsourcing. When suppliers reoptimize input prices in the bust, per-transmission markups rise by $16.4 to $22.8, which equals 40 to 52 percent of baseline markups. This confirms the shock-propagation channel. The downstream demand shock is transmitted upstream through the sourcing margin and partly transformed into higher input prices. In the four-supplier baseline market, however, competitive discipline limits the share of flexibility gains extracted from downstream firms. Upstream repricing transfers an additional $62 million from downstream to upstream firms, which equals 5.6 percent of the strategic-core flexibility gain. Upstream market power therefore alters the incidence of flexibility gains without eliminating most of the downstream adjustment value in the observed market structure.

I next examine how an increase in upstream concentration changes this transfer. This counterfactual is motivated by the United States–Mexico–Canada Agreement, which raises the regional value-content requirement for transmissions and may raise the effective entry barrier faced by non-North-American upstream suppliers. I consider a stylized monopoly case in which Aisin is the only available independent upstream supplier. Under normal demand, monopoly repricing transfers only $10 million from downstream to upstream firms. During the bust, the same pricing margin transfers $64 million, with the additional $54 million induced specifically by the demand contraction. Almost all of the downstream loss falls on firms active on the make-or-buy margin. Upstream concentration does not necessarily generate larger total extraction than four-supplier competition, but it makes the price of production flexibility substantially more state-dependent, concentrating extraction in precisely the downturns when downstream firms most need flexibility.

Related Literature: My work relates to three strands of literature: (i) vertical relations, the make-or-buy decision, and upstream market power; (ii) auto industry vertical integration and trade policy; and (iii) supply chain adjustments to economic shocks.

A long theoretical literature shows how technology, uncertainty, and upstream market structure shape firm boundaries (Carlton, 1979; Green, 1986; Lafontaine and Slade, 2007; Loertscher and Riordan, 2019). Related work studies input sourcing by multinational firms and global value chains (Garetto, 2013; Alfaro et al., 2019), and recent work by Bilal and Lhuillier (2021) studies outsourcing as a response to decreasing returns in in-house production. I contribute to this literature by structurally estimating sourcing decisions as an endogenous firm-level choice and by modeling upstream suppliers as strategic price-setters who anticipate downstream sourcing responses. This allows me to quantify a welfare channel of upstream market power that is absent when input prices or sourcing decisions are treated as fixed.

The paper also relates to the empirical literature on demand fluctuations and industry organization. Collard-Wexler (2013) studies how demand fluctuations affect entry, exit, investment, and market structure in the ready-mix concrete industry. My paper shares the idea that nonlinear payoffs make demand fluctuations matter for industry structure. The mechanism differs: in my setting, demand fluctuations affect the vertical organization of production by changing the value of outsourcing along nonlinear in-house cost curves and by inducing an upstream pricing response.

My paper also contributes to the literature on vertical integration and trade policy in the automobile industry. Classic work studies transaction costs, supplier switching costs, and subcontracting systems in auto production (Klein et al., 1978; Klein, 1988, 2000; Monteverde and Teece, 1982; Langlois and Robertson, 1989; Masten et al., 1989; Taylor and Wiggins, 1997). A common theme in this literature is that the organization of production depends not only on firm boundaries, but also on the market structure of the supplier base. This point is closely related to Stigler (1951), who emphasizes that markets supporting disintegrated production are themselves endogenous.

Recent work studies how trade policy and market power affect the automobile sector. Grossman et al. (2024) analyze tariff propagation through firm-to-firm supply relationships under bilateral negotiation. Heeney et al. (2026) study tariff pass-through from imported parts to vehicle marginal costs, and Duarte et al.

(2025) estimate downstream conduct and scale economies in the U.S. automobile market. My paper adds a component-level sourcing margin and endogenous upstream pricing. Rather than taking input prices or sourcing allocations as fixed, I model how upstream suppliers price the downstream make-or-buy margin. This is important for policies such as USMCA regional-content requirements, which can change the effective set of qualifying suppliers and thereby alter upstream concentration, outsourcing prices, and downstream production flexibility.

Finally, the paper relates to work on supply-chain adjustment and shock propagation. Production-network models show how microeconomic shocks propagate through input-output linkages (Long and Plosser, 1983; Acemoglu et al., 2012, 2017), and empirical studies document propagation through firm-to-firm supply chains after natural disasters and other disruptions (Barrot and Sauvagnat, 2016; Carvalho et al., 2020). Recent work also documents large sourcing reallocations in response to trade tensions and COVID-era disruptions (Alfaro and Chor, 2023). I contribute by providing a micro-founded mechanism through which upstream market power shapes firms' ability to reallocate sourcing after demand shocks, using structurally recovered demand shocks rather than sales-volatility proxies.[3]

This paper proceeds as follows. Section 2 describes the industry background and the data used. Section 3 presents the model. Section 4 discusses identification and estimation procedure. Section 5 reports structural estimation results and Section 6 addresses the economic questions of interest via counterfactual analysis. Section 7 concludes.

2 Industry Background and Data Description

2.1 The Transmission Industry

My research focuses on passenger vehicles in the U.S. market and the upstream market for automobile transmissions. A transmission is a core component of the powertrain: it transmits power from the engine to the wheels and affects driving performance, fuel economy, and vehicle quality. Transmission products can be

[3]Earlier literature uses sales volatility or self-reported uncertainty perceptions to measure demand shocks (Walker and Weber, 1984; Acemoglu et al., 2003; di Giovanni and Levchenko, 2009); Bloom (2014) surveys these proxies.

broadly classified by type and speed. The main types are manual transmission (MT), automatic transmission (AT), continuously variable transmission (CVT), and automated manual transmission (AMT). Except for CVT, each type is offered with several speed options; higher-speed transmissions generally allow smoother gear changes and improved fuel efficiency. AT is the dominant transmission type in the U.S. market, while CVT has gained market share because of its fuel-efficiency advantages.

The transmission market is a useful setting for studying how upstream market power affects downstream production flexibility. Although many automobile input markets are often viewed as competitive buyer markets, transmissions are an exception. Technological barriers, large engineering costs, and the need for vehicle-specific integration leave only a small number of independent suppliers. Table 1 reports summary statistics for the industry from 2009 to 2018. Alongside a substantial amount of in-house production, only six independent suppliers serve the U.S. passenger-vehicle market. Aisin, ZF, and JATCO dominate the AT segment, which accounts for more than 70% of total transmission market share. CVT, the second most important transmission type, is primarily supplied by Aisin and JATCO. Competition is more dispersed in manual transmissions, a more mature technology. For the rest of the paper, I group the three smaller suppliers, GETRAG, Eaton, and TREMEC, into an other-supplier category.

The transmission industry is also directly relevant to recent trade policy. Under the United States–Mexico–Canada Agreement (USMCA), transmissions are classified as super-core components.[4] To qualify for preferential tariff treatment, vehicles must satisfy higher Regional Value Content requirements, and super-core components such as transmissions must satisfy regional-content and labor-value-content rules. Major transmission suppliers such as Aisin, ZF, and JATCO are foreign-headquartered firms. To continue serving the U.S. market under the new rules, they may need to locate production in North America or help downstream manufacturers avoid substantial penalties. These requirements increase entry barriers and can change the effective concentration of the upstream transmission market. This makes the setting well-suited for studying how upstream concentration affects the price of downstream sourcing flexibility.

[4]More information can be found at https://ustr.gov/trade-agreements/free-trade-agreements/united-states-mexico-canada-agreement.

2.2 The Automobile Industry

Figure 1 shows fluctuations in U.S. passenger-vehicle sales over the past four decades. These fluctuations reflect macroeconomic cycles, gasoline prices, trade policy, changes in consumer composition, and substitution toward other mobility options. Such demand movements affect the scale at which automakers operate their in-house production capacity. Historically, firms could buffer demand fluctuations by maintaining excess capacity to ensure final-product delivery. After the 2008 financial crisis, however, many automakers reduced excess capacity to improve cost efficiency. This makes the make-or-buy margin more important: when realized demand moves output away from efficient utilization of in-house capacity, firms may rely more heavily on outside suppliers.

Transmissions provide a particularly clean margin for studying this problem. Compared with downstream vehicle manufacturers, specialized transmission suppliers can serve multiple automakers and multiple vehicle programs. An upstream supplier designs a standardized transmission product and works with downstream manufacturers to integrate vehicle-specific software and calibration. Since much of the customization is managed through software, transmissions sold to different downstream firms can be physically similar while still being tailored to specific vehicle models.[5] This feature makes outside suppliers natural providers of production flexibility.

Industry contracts reinforce this adjustment margin. Manufacturing contracts often last for the life cycle of a product program. Figure 2a shows that downstream firms rarely change transmission providers in my sample. According to the survey evidence in Mueller et al. (2008), automotive supply contracts rarely specify exact quantities, and minimum quantities to be absorbed by downstream firms are also uncommon. Unit prices, however, are typically specified for the initial delivery period, and prices in later periods are either pre-specified through stepwise reduction schedules or renegotiated annually. These contract features imply that downstream firms can adjust the quantity outsourced to a supplier without frequently changing supplier identity.

Unlike many automobile components that are almost completely outsourced,

[5]Udo Rügheimer, automotive spokesman at Bosch, stated: "We are in the ideal position for scale effects. We sell similar components, not identical, but similar components to the Tata Nano and also the Mercedes S-Class. In principle, the parts can be produced on the same machine."

transmissions are still often produced in-house, similar to engines and motors. Figure 2b documents the evolution of in-house transmission production from 2009 to 2018. The figure shows substantial movement in both output-weighted and product-weighted in-house shares, although the two measures do not always move in the same direction. This variation indicates that firms actively adjust the intensive margin of sourcing even though supplier switches are rare. The total variation in the in-house indicator is 0.478, and 30% of this variation comes from within-product changes. These patterns motivate a model in which supplier identity is predetermined, while downstream firms choose the intensity with which each firm–transmission program relies on in-house production versus outsourcing.

2.3 Data

One of the main obstacles in studying vertical relationships is the lack of production-network data, since business-to-business transactions are often private. Producers of automobile parts typically mark their names on the parts they produce, so who-supplies-whom links can in principle be tracked through vehicle tear-down reports. In practice, tracking each of the thousands of parts in each vehicle model is infeasible. Transmissions are different because they must be matched precisely to the vehicle model to ensure compatibility. Replacement-parts platforms such as Transend therefore collect detailed transmission information at the trim level, including transmission product codes and the producing firm.[6] I use this information to construct a dataset linking vehicle models to transmission suppliers for all passenger vehicles sold in the U.S. market from 2009 to 2018.

I combine the transmission-supplier links with data on light-vehicle sales, prices, and characteristics from WardsAuto. WardsAuto reports Manufacturer Suggested Retail Price (MSRP), weight, engine displacement, horsepower, length, width, wheelbase, EPA fuel-economy ratings, drive type, and transmission type. I define a product as a make–model–transmission combination, such as Honda Accord AT6. Because characteristics are reported at the trim level, I construct a baseline version of each product using median characteristics across trims within a make–model–transmission cell. WardsAuto also reports sales by model and transmission. Although MSRP may differ from transaction prices, discounts tend to be largely

[6] https://transend.us/.

manufacturer-specific (Nurski and Verboven, 2016); I absorb these differences with manufacturer fixed effects in the demand estimation.

Three additional data sources complete the dataset. First, I collect local manufacturing wages at assembly locations from the Bureau of Labor Statistics to use as cost shifters.[7] Second, following Petrin (2002), I use the Consumer Expenditure Survey (CEX) automobile supplement to measure new-vehicle purchase probabilities by income group. Third, I sample from the Current Population Survey (CPS) from 2009 to 2018 to approximate the distribution of household demographics.

Table 2 summarizes the key variables. Each year is treated as a separate market, and the sample contains 3,848 make–model–transmission-year products. Following Berry et al. (1995), I use the total number of U.S. households as market size. As in Berry et al. (1995), only a small share of households purchase a new vehicle in a given year, so product-level market shares are small. Apart from a few cases, each product either uses an in-house transmission or sources from one outside supplier. On average, 65% of products use in-house transmissions, and 12% of products use outsourced transmissions from Aisin.

3 A Model of Sourcing Flexibility and Upstream Pricing

This section develops a three-stage model of sourcing flexibility and upstream pricing. Upstream suppliers first set annual input prices. Downstream automakers then choose sourcing intensities for each firm–transmission program. Downstream firms subsequently set final vehicle prices after observing product-level demand and cost states. Firms are risk-neutral. Demand fluctuations matter because they move realized output along nonlinear in-house cost curves. Outsourcing is therefore valuable as a production-flexibility option, while upstream market power determines how much of this flexibility value is extracted through input prices.

The model separates three effects of upstream market power. First, even holding sourcing fixed, upstream suppliers charge markups over marginal cost; this is the standard static double-marginalization channel. Second, because downstream firms can reallocate production between in-house and outsourced transmissions,

[7] Local wages are also used as cost shifters in Wollmann (2018) and Grieco et al. (2021).

upstream suppliers price the downstream production-flexibility option; I call this the flexibility-pricing channel. Third, when demand or cost conditions shift downstream firms' incentive to outsource, upstream suppliers adjust input prices in response. This endogenous price response changes how shocks propagate through the vertical chain; I call this the shock-propagation channel.

I first use a simple linear-demand example to illustrate the mechanism; the full derivations are in Appendix A. The example shows that outsourcing is valuable when realized output falls in regions where in-house production is relatively inefficient, and that suppliers with market power raise input prices when downstream firms' incentive to outsource increases. The full empirical model then embeds this mechanism in random-coefficients demand, multiproduct vehicle pricing, and a finite-action sourcing game.

3.1 Mechanism in a Simple Model

Consider a downstream firm that can either produce an input in-house or buy it from an upstream supplier at input price τ. Let final-good demand be $q(\xi, p)$, where ξ is a demand shock realized after the sourcing decision. The firm's in-house input production cost is $c(q)$, and the outsourced input has unit price τ. Holding final-good pricing fixed for exposition, the relative value of outsourcing in state ξ is governed by

$$\Delta(\xi;\tau) = \pi^{O}(\xi;\tau) - \pi^{I}(\xi) = q(\xi)\left[c^{AC}(q(\xi)) - \tau\right],$$

where $c^{AC}(q) = c(q)/q$ is the average in-house input cost. Outsourcing is valuable in states where the in-house cost of producing the input is high relative to the upstream price. If $c(\cdot)$ is nonlinear, changes in the distribution of ξ affect the expected value of outsourcing by changing the probability that output lies in such regions:

$$E_{\xi}[\Delta(\xi;\tau)] = E_{\xi}\left[q(\xi)\left(c^{AC}(q(\xi)) - \tau\right)\right].$$

Demand fluctuations affect the value of outsourcing by changing where realized output falls on the in-house cost curve. Outsourcing becomes more valuable when demand realizations place downstream firms in output regions where in-house production is relatively inefficient. If demand realizations mostly lie in regions where in-house production is cheaper than outsourcing, the value of outsourcing need not increase.

The upstream supplier chooses τ, anticipating this sourcing margin. Let $Q^O(\tau)$ be the downstream firm's expected outsourced input demand. The supplier's expected profit is

$$\Pi^U(\tau) = (\tau - mc^U)Q^O(\tau),$$

and the first-order condition is

$$Q^O(\tau) + (\tau - mc^U)\frac{\partial Q^O(\tau)}{\partial \tau} = 0. \tag{1}$$

Equation (1) highlights the distinction between the static markup and the flexibility-pricing channel. The term $Q^O(\tau)$ is the direct revenue effect of raising the input price, while $(\tau - mc^U)\partial Q^O(\tau)/\partial\tau$ captures how higher input prices reduce outsourced demand. When the downstream value of outsourcing increases, the entire demand curve $Q^O(\tau)$ shifts outward. A supplier with market power responds by raising τ, so the downstream shock is partly transformed into a higher input price. This endogenous price response is the shock-propagation channel.

3.2 Environment and Timing

I index consumer households by i, products by j, downstream firms by f, transmission types by h, upstream suppliers by s, and years by t. A product is a make–model–transmission combination. In each year t, there is a set of downstream vehicle manufacturers $\mathcal{F}_t$, a set of upstream transmission suppliers $\mathcal{S}_t$, a set of transmission types $\mathcal{H}_t$, and a set of vehicle products $\mathcal{J}_t$.

Each product $j \in \mathcal{J}_t$ belongs to a downstream firm $f(j)$ and uses a transmission type $h(j)$. If product j is outsourced, it is linked to a predetermined upstream supplier $s(j)$. I model the sourcing decision at the firm–transmission level. For a firm–transmission pair (f, h), let $\mathcal{J}_{fht} \subseteq \mathcal{J}_t$ denote the set of products offered by firm f that use transmission type h in year t.

The model has three stages.

1. **Upstream pricing.** Upstream suppliers simultaneously set input prices. In the general model, supplier s sets a unit price τ_{sht} for transmission type h in

year t. This price is set before the realization of product-level demand and marginal-cost shocks.

2. **Downstream sourcing.** After observing the upstream prices τ_t, downstream firms choose sourcing intensities. The decision unit is a firm–transmission pair (f, h). Its action $a_{fht} \in \mathcal{A}_{fht}$ determines how many products within $\mathcal{J}_{fht}$ are produced in-house. Conditional on the action, the specific products assigned to in-house production are randomly drawn from $\mathcal{J}_{fht}$. Products not assigned to in-house production are outsourced to the pair's predetermined upstream supplier.

3. **Downstream vehicle pricing.** After sourcing decisions are made, product-level demand shocks ξ_{jt} and marginal-cost shocks ω_{jt} are realized and observed by downstream firms. Downstream firms then set final vehicle prices in a differentiated-products Bertrand game.

The timing captures the idea that sourcing choices are made before the realization of product-level shocks, while vehicle prices respond to realized demand and cost conditions. Firms are risk neutral, so demand and cost fluctuations affect sourcing decisions through expected profits rather than through direct risk preferences. The production-flexibility value of outsourcing arises because realized output determines where a firm operates on its nonlinear in-house cost curve.

I impose three assumptions that are motivated by the industry background and the data.

Assumption 1: Predetermined supplier relationships. The identity of the outside transmission supplier for each firm–transmission pair is predetermined. Downstream firms choose the intensive margin of how much to outsource to this supplier, rather than choosing among suppliers each year.[8]

Assumption 2: Predetermined product line. The downstream product line is taken as fixed within a year.[9]

[8] This assumption is consistent with the data: supplier switches are rare at the firm–transmission level. Transmission programs require large upfront engineering, calibration, and software-integration costs, so supplier relationships are typically formed before annual production decisions are made. The model can accommodate annual supplier choice by expanding the action set, but doing so substantially increases the computational burden and is not the primary margin observed in the data.

[9] Vehicle model entry and exit are slow relative to annual sourcing and pricing decisions because

Assumption 3: Linear upstream prices. Upstream transmission contracts are summarized by linear unit-price components that do not vary with realized downstream demand shocks.[10]

3.3 Stage 3: Downstream Vehicle Pricing

Consumer demand. I model consumer demand using a random-coefficients logit model. Consumer i's indirect utility from purchasing product j in year t is

$$u_{ijt} = \underbrace{X_{jt}\beta - \alpha p_{jt} + \xi_{jt}}_{\delta_{jt}} + \underbrace{\nu_{i0}\beta_\nu^0 + \log(Y_i)\beta_d^p p_{jt}}_{\mu_{ijt}} + \epsilon_{ijt}.$$

The vector X_{jt} includes vehicle characteristics, transmission characteristics, brand fixed effects, and year fixed effects. The term p_{jt} is the vehicle price. The shock ξ_{jt} is a product-level demand shock observed by firms when final-good prices are set but unobserved by the econometrician. The outside option has mean utility normalized to zero.

The market share of product j is

$$s_{jt}(p_t, \xi_t) = \int \frac{\exp\left(\delta_{jt} + \nu_{i0}\beta_\nu^0 + \log(Y_i)\beta_d^p p_{jt}\right)}{1 + \sum_{m \in \mathcal{J}_t} \exp\left(\delta_{mt} + \nu_{i0}\beta_\nu^0 + \log(Y_i)\beta_d^p p_{mt}\right)} dF_\nu(\nu_{i0}) dF_Y(Y_i). \quad (2)$$

Market demand is $D_{jt} = M_t s_{jt}$, where M_t is market size.

Downstream marginal cost. Let $I_{jt} \in \{0, 1\}$ indicate whether product j's transmission is produced in-house. The downstream marginal cost of product j is

$$mc_{jt} = \widetilde{mc}_{jt} + \omega_{jt} + (1 - I_{jt})\tau_{sht} + I_{jt} c'_{f(j)h(j)t}(D_{jt}), \quad (3)$$

where $\widetilde{mc}_{jt} = X_{jt}^c \gamma$ is the vehicle marginal-cost component excluding the transmission sourcing margin, ω_{jt} is a product-level marginal-cost shock, and τ_{sht} is the

product launches require platform design, engineering, safety testing, and supplier integration. I therefore focus on sourcing and pricing decisions for the observed product set.

[10] Industry evidence and conversations with experts suggest that supply contracts specify unit prices for an initial delivery period and often include pre-specified or renegotiated price paths, while exact annual quantities and minimum purchase commitments are limited. The estimated upstream price component should be interpreted as an average linear unit-price component rather than a literal invoice price for each product.

input price paid if product j's transmission is outsourced. In-house marginal cost $c'_{f(j)h(j)t}(D_{jt})$ can be high at low output because of low utilization, decrease over an intermediate range because of scale economies, and increase again when output approaches capacity constraints. This nonlinearity is what gives outsourcing production-flexibility value.

Pricing equilibrium. Given upstream prices τ_t, sourcing assignments I_t, and realized shocks $e_t = (\xi_t, \omega_t)$, downstream firms set vehicle prices simultaneously. Product-level profits are

$$\pi_{jt} = D_{jt}(p_t)\left[p_{jt} - \widetilde{mc}_{jt} - \omega_{jt} - (1 - I_{jt})\tau_{sht}\right] - I_{jt}c_{f(j)h(j)t}(D_{jt}(p_t)).$$

For a multiproduct firm, the first-order condition for product j is

$$s_{jt}(p_t) + \sum_{k \in \mathcal{J}_t} \Gamma_{jk,t}\left[p_{kt} - mc_{kt}(p_t, \tau_t, I_t, e_t)\right] \frac{\partial s_{kt}(p_t)}{\partial p_{jt}} = 0. \tag{4}$$

Here $\Gamma_{jk,t} = 1$ if products j and k are priced by the same downstream pricing unit and 0 otherwise. In the empirical implementation, Γ_t groups products by firm–transmission pair. Let $p_t^*(\tau_t, I_t, e_t)$ denote the downstream Bertrand equilibrium price vector, and let $D_t^*(\tau_t, I_t, e_t)$ and $\pi_t^*(\tau_t, I_t, e_t)$ denote the associated demand and product profits.

3.4 Stage 2: Downstream Sourcing Decisions

Before product-level demand and marginal-cost shocks are realized, downstream firms choose sourcing intensities. A Stage-2 player is a firm–transmission pair (f, h). The action

$$a_{fht} \in \mathcal{A}_{fht}$$

determines the number or fraction of products in $\mathcal{J}_{fht}$ whose transmissions are produced in-house. The action set is finite and may differ across firm–transmission pairs because product counts differ.

Conditional on an action profile $a_t = (a_{fht})_{(f,h)}$, the model draws product-level sourcing assignments $I_t^n(a_t)$, $n = 1, \ldots, N_A$. For each draw, the number of in-house products in each firm–transmission pair equals the chosen action. Products not selected for in-house production are outsourced to the pair's predetermined supplier. This random-allocation device reduces the dimensionality of the combinatorial problem while preserving the within-program nature of the sourcing decision.[11]

For a given action profile, the expected product profit is

$$E\pi^*_{jt}(a_t, \tau_t) = \frac{1}{N_A} \sum_{n=1}^{N_A} E_{e_t} \left[\pi^*_{jt}\left(\tau_t, I^n_t(a_t), e_t\right) \right].$$

The expected payoff of firm–transmission pair (f, h) is

$$E\pi_{fht}(a_{fht}, a_{-fht}, \tau_t) = \sum_{j \in \mathcal{J}_{fht}} E\pi^*_{jt}(a_t, \tau_t).$$

The sourcing game is modeled as a private-information discrete game. Each firm–transmission pair receives an action-specific payoff shock $\epsilon_{fht}(a)$, observed by the decision maker but not by rivals or the econometrician.[12] I assume that these shocks are i.i.d. Type-I extreme value across actions and firm–transmission pairs.

The private information structure serves two purposes. First, it is economically natural in this setting because sourcing decisions depend on program-specific information that is difficult for rivals and the econometrician to observe. Second, it regularizes a high-dimensional discrete game. Complete-information discrete games with many players and multiple sourcing intensities can have multiple equilibria and discontinuous best responses. Private information yields smooth logit best responses and represents equilibrium as a fixed point in choice probabilities, which is useful for computing comparative statics with respect to upstream prices.[13]

[11] For example, if a firm–transmission pair has ten products and chooses action $a_{fht} = 1/2$, then five products are assigned to in-house production in each assignment draw. The identity of the five products is randomly drawn. This avoids modeling every product-level allocation as a separate combinatorial choice. Yang (2020) also adopt similar method for tractability concerns.

[12] These shocks capture unobserved factors that affect the desirability of a sourcing intensity, such as engineering constraints, managerial capacity, plant-level frictions, or unobserved cost differences. Similar private-information payoff shocks are widely used in empirical discrete games; see, for example, Rust (1994) and Seim (2006)

[13] Private-information entry and discrete games are often used to address equilibrium-selection and computational issues. Espín-Sánchez et al. (2023) provide sufficient conditions for equilib-

Let $\sigma_{fht}(a)$ denote the probability that firm–transmission pair (f, h) chooses action a. Given beliefs about rivals' strategies, the deterministic component of the expected value of action a is

$$E\Pi_{fht}(a, \tau_t; \sigma_{-fht}) = \sum_{a_{-fht}} E\pi_{fht}(a, a_{-fht}, \tau_t) \prod_{(f',h') \neq (f,h)} \sigma_{f'h't}(a_{f'h'}).$$

The logit best response is

$$\sigma_{fht}(a) = \frac{\exp\left(E\Pi_{fht}(a, \tau_t; \sigma_{-fht})/\lambda_\epsilon\right)}{\sum_{a' \in \mathcal{A}_{fht}} \exp\left(E\Pi_{fht}(a', \tau_t; \sigma_{-fht})/\lambda_\epsilon\right)},$$

where λ_ϵ is the scale parameter of the private-information shocks. A Stage-2 Bayesian Nash equilibrium is a set of choice probabilities $\sigma_t^*(\tau_t)$ satisfying

$$\sigma_t^*(\tau_t) = \Psi(\sigma_t^*(\tau_t), \tau_t). \tag{5}$$

3.5 Stage 1: Upstream Pricing

Upstream suppliers set prices to maximize expected profit, internalizing how downstream firms' sourcing probabilities and vehicle prices respond to upstream prices. In the general model, supplier s's expected profit is

$$\Pi_{st}^U(\tau_t) = \sum_{h \in \mathcal{H}_{st}} \left(\tau_{sht} - mc_{sht}^U\right) ED_{sht}(\tau_t), \tag{6}$$

where mc_{sht}^U is supplier s's marginal cost of producing transmission type h, and $ED_{sht}(\tau_t)$ is expected outsourced demand for supplier s's transmission type h.

For action profile a_t, define

$$Q_{sht}(a_t, \tau_t) = \sum_{j \in \mathcal{J}_t} \mathbf{1}\{j \text{ is outsourced to supplier } s \text{ for transmission } h\} E\left[D_{jt}^*\left(\tau_t, I_t(a_t), e_t\right)\right],$$

where the expectation integrates over product-level assignment draws and de-

---

rium uniqueness in a class of private-information entry games. My sourcing game differs from their binary-entry environment because actions are multi-valued and payoffs are generated by a downstream pricing equilibrium; I therefore use their results as motivation rather than as a direct uniqueness theorem.

mand/cost shocks. The equilibrium expected demand faced by supplier s for transmission type h is

$$ED_{sht}(\tau_t) = \sum_{a_t} Pr_t^*(a_t|\tau_t) Q_{sht}(a_t, \tau_t),$$

where $Pr_t^*(a_t|\tau_t) = \prod_{(f,h)} \sigma_{fht}^*(a_{fht}|\tau_t)$.

The first-order condition for τ_{sht} is

$$0 = ED_{sht}(\tau_t) + \left(\tau_{sht} - mc_{sht}^U\right) \frac{\partial ED_{sht}(\tau_t)}{\partial \tau_{sht}}. \tag{7}$$

The derivative of expected supplier demand can be decomposed as

$$\frac{\partial ED_{sht}(\tau_t)}{\partial \tau_{\ell rt}} = \sum_{a_t} \frac{\partial Pr_t^*(a_t|\tau_t)}{\partial \tau_{\ell rt}} Q_{sht}(a_t, \tau_t) + \sum_{a_t} Pr_t^*(a_t|\tau_t) \frac{\partial Q_{sht}(a_t, \tau_t)}{\partial \tau_{\ell rt}}.$$

The first term is the sourcing-probability response. The second term is the downstream demand response induced by the pass-through of upstream prices into vehicle marginal costs and final-good prices.

The static markup channel is captured by the margin $\tau_{sht} - mc_{sht}^U$: holding downstream sourcing fixed, upstream suppliers charge a markup over their marginal cost. The flexibility-pricing channel operates through $ED_{sht}(\tau_t)$: downstream firms' ability to shift production from in-house plants to outside suppliers creates an input-demand curve faced by the upstream supplier, and the supplier prices this sourcing margin. The shock-propagation channel operates through the dependence of both $Q_{sht}(a_t, \tau_t)$ and $Pr_t^*(a_t|\tau_t)$ on the upstream price vector and on the distribution of downstream demand and cost states. When shocks increase downstream firms' incentive to outsource, $ED_{sht}(\tau_t)$ shifts outward; a supplier with market power responds by raising τ_{sht}, partly transforming downstream production shocks into higher input prices.

Equation (7) can be inverted to recover the supplier's upstream marginal cost:

$$mc_{sht}^U = \tau_{sht} + \frac{ED_{sht}(\tau_t)}{\partial ED_{sht}(\tau_t)/\partial \tau_{sht}}.$$

In the empirical implementation, upstream transaction prices are not observed. I therefore use a parsimonious supplier-year price index τ_{st}, which should be interpreted as an average linear unit-price component for supplier s in year t. This restriction reduces the dimensionality of the upstream pricing problem and makes the counterfactual equilibrium computationally feasible. Under this restriction, the summation over h in (6) is absorbed into the supplier-year expected demand $ED_{st}(\tau_t)$, and the same inversion logic recovers a supplier-year effective marginal cost mc^U_{st}. I examine robustness to calibrated upstream cost curvature and scaled upstream marginal costs in the counterfactual analysis.

Why posted upstream prices rather than bilateral bargaining. I model upstream prices as supplier-set linear unit-price components rather than as pair-specific Nash-in-Nash bargaining outcomes.[14] This choice is disciplined by both the economic question and the data. The paper studies how upstream market power prices the downstream sourcing margin and thereby affects shock propagation, rather than how pair-specific surplus is divided across buyers. These channels operate through the aggregate input-demand elasticity faced by upstream suppliers and through downstream firms' sourcing responses. A bilateral bargaining model would instead put most of the structure on pair-specific surplus division and price discrimination across downstream buyers, which is not the primary margin in this paper.

The institutional environment also makes a common-price-index approximation natural. Transmission programs involve large upfront engineering and software integration costs, and upstream and downstream firms interact repeatedly over many years. For a given program, downstream firms typically work with one outside supplier rather than simultaneously negotiating with a set of complementary suppliers. In such an environment, repeated interaction and information about peer purchasing limit the scope for large unexplained price dispersion across comparable buyers(Grennan and Swanson, 2020). I therefore summarize upstream contracts by linear unit-price components and focus on how those prices respond to downstream sourcing incentives.

This restriction is also necessary empirically. I do not observe transaction prices between automakers and transmission suppliers, and marginal-cost data are not

[14] Empirical applications of bilateral bargaining often study settings such as cable television, hospitals, insurers, and employers, where negotiated prices or margins can be disciplined by richer contract or claims data (Chipty and Snyder, 1999; Crawford and Yurukoglu, 2012; Crawford et al., 2018; Gowrisankaran et al., 2015; Ho and Lee, 2017).

available for all upstream suppliers. I therefore cannot discipline a bargaining model with observed pair-specific margins or estimate bargaining parameters as in settings with richer contract data. In the general model, supplier s charges a unit price τ_{sht} for transmission type h in year t. In the empirical implementation, I use a parsimonious supplier-year price index τ_{st}, which should be interpreted as an average linear unit-price component rather than a literal invoice price for every product. This restriction keeps the counterfactual problem tractable while preserving the core channels of interest: upstream markups, the price of downstream sourcing flexibility, and shock propagation through the vertical chain.[15]

3.6 Equilibrium

An equilibrium in year t consists of upstream prices τ_t^*, downstream sourcing choice probabilities σ_t^*, and downstream vehicle prices p_t^* such that:

1. Given upstream prices τ_t^*, product-level sourcing assignments I_t, and realized demand and cost shocks e_t, downstream vehicle prices $p_t^*(\tau_t^*, I_t, e_t)$ solve the differentiated-products Bertrand pricing first-order conditions in (4).

2. Given upstream prices τ_t^*, downstream sourcing probabilities $\sigma_t^*(\tau_t^*)$ form a Bayesian Nash equilibrium of the Stage-2 sourcing game and satisfy (5).

3. Given downstream sourcing equilibrium $\sigma_t^*(\tau_t)$ and downstream pricing equilibrium $p_t^*(\tau_t, I_t, e_t)$, upstream prices τ_t^* satisfy the upstream suppliers' first-order conditions in (7).

4 Identification and Estimation

This section describes the estimation of the demand system, the construction of demand and cost shock draws used in the sourcing game, the estimation of downstream marginal costs and in-house transmission costs, and the recovery of upstream marginal costs. The key objects are the demand parameters $\theta^d =$

[15] Villas-Boas (2007) shows how detailed retail and wholesale marginal-cost data can be used to compare alternative vertical-pricing models. Such transaction-level input-price data are not available in my setting.

$(\beta, \alpha, \beta_\nu^0, \beta_d^p)$, the downstream marginal-cost parameters, the in-house cost function $c(\cdot)$, the upstream price index τ_{st}, and the upstream marginal costs mc_{st}^U.

4.1 Estimating Downstream Demand-side Parameters θ^d

I estimate vehicle demand using the random-coefficients logit model introduced in Section 3. For a candidate vector of nonlinear parameters (β_ν^0, β_d^p), the model-implied market share of product j in year t is

$$s_{jt}(\delta_t, p_t; \beta_\nu^0, \beta_d^p) = \int \frac{\exp\left(\delta_{jt} + \nu_{i0}\beta_\nu^0 + \log(Y_i)\beta_d^p p_{jt}\right)}{1 + \sum_{m \in J_t} \exp\left(\delta_{mt} + \nu_{i0}\beta_\nu^0 + \log(Y_i)\beta_d^p p_{mt}\right)} dF_\nu(\nu_{i0}) dF_Y(Y_i).$$

Market demand is $D_{jt} = N_t s_{jt}$, where N_t is market size. Following Berry (1994), for each candidate value of the nonlinear parameters I invert observed market shares to recover the vector of mean utilities

$$\delta_{jt} = f(s_t, p_t, \beta_\nu^0, \beta_d^p).$$

The recovered demand residual is therefore

$$\xi_{jt} = f(s_t, p_t, \beta_\nu^0, \beta_d^p) - \left(X_{jt}\beta - \alpha p_{jt}\right).$$

The main demand-side endogeneity concern is that vehicle prices are correlated with the unobserved product quality component ξ_{jt}. I use excluded instruments that shift prices and substitution patterns but are plausibly orthogonal to current product-level demand shocks. The first group is the standard BLP-style instruments based on exogenous characteristics of rival products. These instruments are valid under the maintained assumption that vehicle characteristics are chosen before the realization of the demand shock and affect product demand through the demand system rather than through ξ_{jt}.

Because standard BLP instruments can be weak when products are close substitutes along only a few dimensions, I also use differentiation instruments following Gandhi and Houde (2019). These instruments measure how isolated a product is in characteristics space and in predicted-price space. The instrument set is

$$z_{jt} = \left\{ x_{jt}, \sum_{j' \neq j} (d^k_{jt,j'})^2, \sum_{j' \neq j} (d^{\hat{p}}_{jt,j'})^2 \right\}.$$

Here $d^k_{jt,j'}$ is the distance between product j and product j' along characteristic k. I use length, horsepower, fuel efficiency, and transmission speed. The term $d^{\hat{p}}_{jt,j'}$ is the distance between predicted prices, where predicted prices are obtained from exogenous product characteristics and cost shifters. These differentiation instruments help identify substitution patterns and the nonlinear parameters by exploiting variation in the closeness of competing products.

I also use local manufacturing wages at the vehicle assembly location as cost shifters. These wages shift marginal costs and therefore equilibrium prices, but are unlikely to be correlated with contemporaneous product-level demand shocks after controlling for observed characteristics and fixed effects. Firms rarely reallocate products across assembly sites within the sample period, which supports the exclusion restriction.[16] Local wages are also used as cost shifters in Wollmann (2018) and related automobile demand applications.

4.2 Constructing Demand Shock Draws

The sourcing game requires expectations over product-level demand and marginal-cost states. I construct these states from the residuals recovered from the demand and supply systems. The goal is to generate shock draws that are disciplined by the structural demand estimates rather than by raw sales volatility.

Let $\hat{\xi}_{jt}$ denote the demand shock recovered from the demand system. Figure 3 reports diagnostics for the recovered product-level demand shocks. Panel (a) shows that the distribution is centered close to zero and is approximately single-peaked, with some left-tail observations. Panel (b) plots the distribution by year. The yearly medians remain close to zero, while the within-year dispersion is sizable. These patterns indicate that the recovered shocks are not primarily aggregate

[16] For vehicles assembled outside the United States, I use comparable wage measures when available. The wage data are more detailed for Canada, Mexico, and Japan. For other countries, I use country-level wage data. All foreign wage measures are converted into U.S. dollars using purchasing power parity.

year effects; instead, they contain product-level residual variation after controlling for observed characteristics, brand fixed effects, and year fixed effects.

I decompose the recovered demand residual into a common firm–transmission component and a product-specific component:

$$\hat{\xi}_{jt} = \zeta_{f(j)h(j)t} + u_{jt}.$$

The common component $\zeta_{f(j)h(j)t}$ captures shocks that affect products within the same firm–transmission program and therefore line up with the Stage-2 sourcing decision unit. The product-specific component u_{jt} captures residual demand variation across products within a program.

This decomposition has a natural empirical-Bayes interpretation. The observed demand residuals are noisy realizations of latent product and firm–transmission demand states. Estimating variance components and shrinking noisy product-level moments toward a common distribution reduces the influence of short product histories and extreme observations. This is useful because many products are observed for only a limited number of years, while sourcing decisions are made at a coarser firm–transmission level. The empirical-Bayes logic is standard in normal-means and compound-decision problems: information is pooled across related units to recover latent heterogeneity more stably than using raw unit-level moments alone (Efron, 2016).

I estimate two variance components. The first, $\sigma^2_{\zeta,fh}$, governs the common firm–transmission component. The second, $\sigma^2_{u,j}$, governs product-specific residual variation. The simulation shock is then generated as

$$\xi^{draw}_{jt} = \zeta^{draw}_{f(j)h(j)t} + u^{draw}_{jt}.$$

This structure preserves the distinction between shocks relevant for the sourcing unit and shocks idiosyncratic to individual vehicle products.

I construct two demand-shock panels. In the iid specification, ζ_{fht} and u_{jt} are drawn independently across years:

$$\zeta^{draw}_{fht} \sim N(0, \sigma^2_{\zeta,fh}),$$
$$u^{draw}_{jt} \sim N(0, \sigma^2_{u,j}).$$

In the persistent specification, both components follow an AR(1) process. For consecutive observed years t_{k-1} and t_k,

$$\zeta^{draw}_{fht_k} = \rho^{\Delta}\zeta^{draw}_{fht_{k-1}} + \nu^{\zeta}_{fht_k}, \qquad \nu^{\zeta}_{fht_k} \sim N\left(0, \left(1-\rho^{2\Delta}\right)\sigma^2_{\zeta,fh}\right),$$

where $\Delta = t_k - t_{k-1}$. The same recursion is used for u_{jt}. This innovation variance preserves the unconditional variance of each component. In the baseline persistent specification, I set $\rho = 0.46$.[17]

For marginal-cost shocks, I use the residuals recovered from the downstream supply equation. I draw cost shocks by bootstrapping from the empirical distribution of $\hat{\omega}_{jt}$. Before bootstrapping, I winsorize the cost residuals at the tails in the baseline production run to avoid allowing a small number of extreme cost residuals to dominate the fixed-point price solver. I report robustness to the winsorization rule.

The demand and cost shock panels are pre-generated and aligned with the product-year simulation panel. Each column is one Monte Carlo draw of the full panel, and the same pre-generated shocks are reused across counterfactuals and equilibrium iterations. This implements common random numbers and reduces Monte Carlo noise when comparing equilibria. For counterfactuals that scale demand volatility, I multiply the loaded demand-shock matrix by the volatility factor at simulation time rather than regenerating the panel. Because both Gaussian and AR(1) draws are linear in the standard deviation, this scaling is equivalent to drawing from the scaled shock distribution.

I use the AR(1) factor-structure shock panel as the baseline and report iid shocks as a robustness check. The pre-generated panels include N independent draws; antithetic pairing is not imposed in the baseline because the stored shock matrices are intended to be a flexible simulation input. I use larger independent panels to verify Monte Carlo convergence.

[17] The value $\rho = 0.46$ is estimated from the recovered demand shocks using a pooled within-product AR(1) regression on consecutive-year observations. The generated AR(1) shock panel reproduces this persistence: the empirical AR(1) coefficient is 0.464, while the iid panel has coefficient 0.004.

4.3 Estimating Downstream Supply-side Parameters θ^s

I estimate downstream marginal costs using the vehicle-pricing first-order conditions and the sourcing structure observed in the supplier–vehicle data. The marginal cost of product j in year t is

$$mc_{jt} = X_{jt}^c \gamma + (1 - I_{jt})\tau_{s(j)t} + I_{jt} c'_{f(j)h(j)t}(D_{jt}) + \omega_{jt},$$

where $I_{jt} = 1$ if the transmission is produced in-house and $I_{jt} = 0$ if it is outsourced. The term $X_{jt}^c \gamma$ captures the vehicle marginal-cost component excluding the transmission sourcing margin, $\tau_{s(j)t}$ is the upstream price index for the supplier serving product j, and $c'_{f(j)h(j)t}(D_{jt})$ is the marginal cost of producing the transmission in-house.

I parametrize the in-house transmission cost function flexibly as a third-order polynomial in product demand:

$$c(D_{jt}) = c_1 D_{jt} + c_2 D_{jt}^2 + c_3 D_{jt}^3.$$

This specification allows average in-house costs to be high at low output because of poor capacity utilization, decline over an intermediate range because of scale economies, and rise again when output approaches capacity constraints. I also allow selected downstream firms to have different in-house cost levels, reflecting persistent heterogeneity in production technology and plant organization.

The upstream transaction prices are not observed. I therefore estimate a supplier-year price index τ_{st} and interpret it as an average linear unit-price component for supplier s in year t.[18] To reduce dimensionality, I approximate each supplier's price path with a quadratic time trend,

$$\tau_{st} = \tau_s + \tau_s^{trend} t + \tau_s^{trend_2} t^2,$$

where t is the number of years since 2009.

[18] To reduce the number of parameters and make counterfactual upstream pricing computationally feasible, I assume that supplier s charges the same average price component τ_{st} across products in year t. The estimate should be interpreted as an average unit-price component rather than a literal invoice price for each transmission product. Council (2015) reports direct manufacturing-cost differences across transmission types; for example, the incremental cost of an eight-speed transmission relative to a ZF six-speed transmission is reported to be about $61.84 in EPA/FEV estimates.

The key supply-side endogeneity concern is that D_{jt} is an equilibrium quantity and therefore depends on the unobserved marginal-cost shock ω_{jt}. I instrument for D_{jt} using predicted demand constructed from exogenous product characteristics and the differentiation instruments used in the demand estimation. Specifically, I use Lasso to project D_{jt} on the high-dimensional set of excluded demand shifters and product characteristics, and then use the predicted value as an instrument for realized demand in the cost equation. Belloni et al. (2012) show that IV estimators based on Lasso or Post-Lasso first stages are root-n consistent and asymptotically normal under approximate sparsity, so standard inference procedures can be applied.

The supply-side parameter vector is $\theta^s = (\gamma, \tau_s, \tau_s^{trend}, \tau_s^{trend_2}, c_1, c_2, c_3)$, together with the firm-specific in-house cost shifters. Because a product is either produced in-house or outsourced, the level of the in-house cost function and the level of the upstream price index are not separately identified without a normalization. I normalize the 2009 price of the other-supplier group, τ_O, to zero. The in-house cost level and the supplier price indices should therefore be interpreted relative to this normalized base group.

4.4 Recovering Upstream Marginal Costs

With the estimated demand system, downstream marginal-cost parameters, in-house cost function, and upstream price index τ_{st}, I recover upstream suppliers' marginal costs from the upstream pricing first-order conditions in Equation 7. For each year, I solve the sourcing game at the estimated upstream prices, compute each supplier's expected outsourced demand, and invert the supplier first-order condition to recover mc_{st}^U. The inversion uses the fact that, in equilibrium, an upstream supplier's price reflects both its markup and the responsiveness of expected outsourced demand to that price.

The recovered mc_{st}^U is a supplier-year effective marginal cost implied by the model, the estimated upstream price index, and the maintained linear-contract restriction. Because the data contain four upstream supplier groups over ten years, I do not estimate a flexible upstream cost curve in the baseline. Instead, I recover one marginal-cost term for each supplier-year. In the counterfactual analysis, I assess robustness by scaling the recovered marginal costs and by allowing calibrated upstream cost curvature.

The main computational challenge is solving the downstream sourcing game. The full game is high dimensional because each active firm–transmission pair can choose among multiple sourcing intensities, and each sourcing intensity corresponds to many possible product-level allocations. To keep the problem computationally feasible, I solve the private-information sourcing game for a strategic core of firm–transmission pairs that are active on the make-or-buy margin and account for large supplier-level demand. In the baseline, I select the largest active firm–transmission pair associated with each major upstream supplier. Products outside the strategic core keep their observed sourcing assignment, but still enter supplier demand and respond to upstream prices through the downstream vehicle-pricing equilibrium.

For each strategic action profile, I simulate product-level sourcing assignments and demand/cost shock draws, solve the downstream vehicle-pricing equilibrium, and compute expected profits and supplier-level outsourced demand. I then solve the Stage-2 private-information equilibrium and use the resulting action-profile probabilities to compute expected outsourced demand for each supplier. The derivative of expected supplier demand with respect to upstream prices is computed analytically. It combines two components: the response of downstream sourcing probabilities and the response of downstream vehicle demand through price reoptimization. Details of the derivative calculation, the strategic-core approximation, and sensitivity tests for the number of players, action grid, assignment draws, and shock draws are reported in Appendix B.

5 Estimation Results

I discuss the estimated demand and marginal cost parameters of downstream firms, and the cost parameters of upstream firms.

5.1 Estimation Results for θ^d and θ^s

Table 3 reports the demand estimates. Column (1) reports a standard logit specification that ignores consumer heterogeneity. Column (2) reports the random-coefficients demand estimates. Relative to the logit specification, the random-coefficients model implies a larger price coefficient and more sensible estimates for product characteristics. Consumers value larger vehicles, higher horsepower, and higher fuel efficiency on average. For example, a one-mile-per-gallon increase in fuel efficiency is equivalent to a price decrease of $418, and a one-meter increase in vehicle size is equivalent to a price decrease of $380. Vehicle type also matters: SUVs have a premium of $5,217 relative to sedans.

Transmission characteristics are also important for demand. Vehicles equipped with ZF transmissions have a premium of $10,961. This estimate should be interpreted as a composite premium: it reflects the value of ZF transmissions as well as potential complementarities between ZF transmissions and the vehicle models that use them. The estimated heterogeneity in the outside option is also economically meaningful. Its standard deviation is about 60% of the mean, capturing substantial dispersion in consumers' outside-option values. Accounting for this heterogeneity generates more flexible substitution patterns and more sensible markups.

The demand system implies reasonable own-price elasticities. I allow price sensitivity to vary with income, and the estimates imply that higher-income consumers are less price-sensitive. Consistent with profit maximization in oligopoly, all own-price elasticities are greater than one in absolute value. Table 4 reports summary statistics for price elasticities, marginal costs, and margins. The average gross margin is 41.3%, broadly in line with Berry et al. (1995), Goldberg and Verboven (2001), and Nurski and Verboven (2016). More expensive vehicle models also have higher marginal costs, consistent with vertical quality differentiation.

Table 5 reports the downstream supply estimates. Vehicle characteristics enter marginal costs with the expected signs: horsepower, size, and fuel efficiency all increase production costs. Foreign vehicles are more costly to build, and SUVs have higher marginal costs than sedans. The estimated in-house transmission cost function is nonlinear. Average in-house production costs first decline with output, reflecting scale economies and improved capacity utilization, and then increase when output becomes sufficiently high, reflecting capacity constraints and adjustment frictions.

The estimates also reveal substantial heterogeneity in in-house production costs across downstream firms. This heterogeneity is consistent with observed sourcing patterns. Daimler, Honda, Hyundai, and Volkswagen have sizable in-house transmission production in the data, suggesting that these firms differ in their internal production capabilities. The estimates imply lower average in-house transmission costs for Daimler and Volkswagen, and Daimler's in-house cost decreases further over time. The corresponding time trends are not statistically significant for Honda, Hyundai, and Volkswagen.

The estimated coefficient on local labor costs is negative, which is initially counterintuitive. One possible interpretation is that local manufacturing wages partly proxy for the historical organization of auto production in the United States. The Big Three have traditionally operated in areas with higher unionization and higher labor costs, while some premium foreign manufacturers locate assembly plants in lower-unionization states. Because the model also controls for product characteristics and firm fixed effects, the labor-cost coefficient should be interpreted as a residual cost shifter rather than as a structural labor demand elasticity.

The output variable D_{jt} enters the in-house cost function and is endogenous because equilibrium demand depends on the unobserved marginal-cost shock. I instrument for D_{jt} using a Lasso prediction based on exogenous vehicle characteristics and differentiation instruments from the demand estimation. The first stage is strong: the Cragg–Donald F statistic is 47.05, above the 5% critical value of 20.93. Without the fitted value $\hat{D}_{jt}$, the Cragg–Donald F statistic falls to 6.96 and the cost-function coefficients become implausibly large, indicating that weak instruments materially distort the estimated in-house cost curve.

Figure 4a reports estimated upstream transmission price differences relative to the normalized baseline group τ_O. The three major upstream suppliers have higher estimated prices than the other-supplier group. This pattern is consistent with differences in supplier quality and product mix: automatic transmissions and continuously variable transmissions, which are mainly produced by Aisin, ZF, and JATCO, are more expensive than manual transmissions. The estimated price differences are statistically significant at the 10% level.

Figure 4b illustrates the production-flexibility mechanism using the estimated in-house cost curve and the 2018 ZF transmission price. The in-house production cost curve is U-shaped relative to the approximately constant upstream input price. At low realized output, poor capacity utilization raises the cost of producing transmissions internally. At intermediate output levels, scale economies lower in-house costs. At sufficiently high output, capacity constraints increase in-house costs again. Economic downturns are therefore especially relevant in this setting:

when shrinking demand moves downstream firms toward the low-utilization region of the cost curve, outsourcing becomes a valuable production-adjustment margin. This is the structural source of the production-flexibility value measured in the counterfactual analysis.

5.2 Estimation Results for mc_{st}

Figure 5 reports the recovered upstream marginal costs from 2009 to 2018, expressed relative to the other-supplier group in 2009. The recovered marginal costs are below the corresponding estimated input prices, suggesting that the implied upstream markups are sensible. Consistent with the price estimates, Aisin, ZF, and JATCO have higher recovered marginal costs than the other-supplier group, reflecting differences in transmission types, product quality, and supplier composition.

The recovered marginal cost for ZF varies substantially over time. This pattern is consistent with product-cycle and technology changes, including the introduction of the AT9 transmission in 2013 and subsequent improvements in production efficiency. The other-supplier category combines GETRAG, Eaton, and TREMEC, so movements in its recovered marginal cost may also reflect changes in the composition of suppliers and transmission products within that residual group.

The recovered upstream cost estimates provide the primitives needed for the counterfactual exercises. Holding these marginal costs fixed, I can compare equilibria in which upstream prices remain fixed to equilibria in which upstream suppliers reoptimize prices in response to changes in downstream sourcing incentives. This comparison isolates the price of downstream production flexibility and the vertical shock-propagation channel emphasized in the model.

6 Counterfactuals

This section uses the estimated model to quantify the value of sourcing flexibility and the price that upstream suppliers charge for it. I first decompose the effects of a large demand contraction in the 2018 market. I then examine how a more concentrated upstream structure changes the state dependence of upstream repricing.

6.1 The Impact of an Economic Bust

The recent pandemic has had a drastic impact on the US automobile industry. Especially in the first few months of the pandemic, travel was discouraged, dealers' showrooms were closed, and the demand for new vehicles collapsed.[19] From Figure 1, one can see that sales dropped by almost 50% in early 2020. Many manufacturers significantly decreased the number of shifts or even temporarily shut down some plants. This shock is useful for studying the vertical mechanism in my model because a decline in vehicle demand changes more than final-good sales. It moves in-house transmission production along nonlinear cost curves, changes the relative attractiveness of outsourcing, and shifts the input-demand curve faced by upstream suppliers. Therefore, an economic bust affects the vertical chain through both downstream sourcing decisions and upstream pricing responses. To mimic such a bust, I shrink the size of the downstream market by one-third and recompute the downstream sourcing equilibrium and upstream prices.

This counterfactual maps directly to the three channels in the model. The first is the direct effect of the demand contraction, holding fixed the ability of strategic downstream firms to reallocate production. The second is the production-flexibility channel: downstream firms can respond to the bust by shifting production between in-house and outsourced transmissions. The third is the shock-propagation channel: upstream suppliers internalize the increased demand for outsourced transmissions and adjust input prices in response.

I evaluate the following three counterfactual equilibria. In Scenario 1, strategic firm–transmission pairs are forced to use the full in-house action while non-strategic products keep their observed sourcing assignments. In Scenario 2, I allow the Stage-2 sourcing equilibrium to adjust while holding upstream prices fixed at baseline values. In Scenario 3, upstream suppliers reoptimize prices using the Stage-1 first-order conditions and the recovered 2018 upstream marginal costs. Table 6 reports the decomposition.

I evaluate three counterfactual equilibria. In Scenario 1, strategic firm–transmission pairs are forced to use the full in-house action while non-strategic products keep their observed sourcing assignments. In Scenario 2, the Stage-2 sourcing equilib-

[19] According to the report from Mckinsey, "The effects began in China, where sales plunged 71 percent in February 2020; by April, sales had dropped 47 percent in the United States and dived 80 percent in Europe."

rium adjusts while upstream prices are held fixed at baseline values. In Scenario 3, upstream suppliers reoptimize prices using the Stage-1 first-order conditions and the recovered 2018 upstream marginal costs. Table 6 reports the decomposition.

The demand contraction is large. Relative to the baseline, total welfare in Scenario 1 falls by $29.82 billion, almost entirely on the producer side. Strategic-core profit falls by 41 percent, while consumer surplus falls only marginally because the bust is modeled as a contraction in market size rather than a shift in tastes over the inside goods. Upstream profit also falls, reflecting the smaller volume of outsourced transmissions.

Allowing strategic downstream firms to adjust sourcing decisions partially offsets these losses. Holding upstream prices fixed (Scenario 2 versus Scenario 1), strategic-core profit rises by $1.10 billion, recovering 17.8 percent of the bust-induced strategic-core loss. The make-or-buy margin therefore allows strategic firms to recover roughly 18 cents on every dollar of bust-induced profit loss. When in-house production becomes less attractive because realized output moves into the high-average-cost region of the cost curve, downstream firms reallocate production toward outside suppliers and reduce the cost of absorbing the shock.

Upstream suppliers respond to the increased value of outsourcing. Table 7 shows that reoptimized upstream prices rise by 40 to 52 percent of baseline upstream markups, with suppliers nearly doubling their per-transmission margin in the bust. This confirms that the shock-propagation channel is active: the bust raises the value of sourcing flexibility, and upstream suppliers raise the price of that flexibility.

Comparing reoptimized upstream prices (Scenario 3) to fixed upstream prices (Scenario 2) isolates the extraction channel. Upstream profit rises by an additional $62 million when suppliers reoptimize. Relative to the $1.10 billion strategic-core flexibility gain at fixed upstream prices, suppliers extract roughly 5.6 percent through repricing in the four-supplier baseline market. Two features of the four-supplier market explain why extraction is limited despite large per-unit markup increases. First, the bust contracts the market by one-third, so the larger per-transmission markup is applied to a smaller outsourced volume. Second, four-supplier competition disciplines unilateral price increases, since each supplier internalizes that raising its own price diverts outsourced demand to competitors and induces some downstream firms to return to in-house production.

6.2 Upstream Concentration and State-Dependent Repricing

The previous counterfactual shows that upstream suppliers raise input prices when downstream firms have stronger incentives to outsource, and that four-supplier competition limits how much of the resulting flexibility value is transferred to upstream firms. I next examine how an increase in upstream concentration changes this transfer. This exercise is motivated by the United States–Mexico–Canada Agreement (USMCA), which entered into force on July 1, 2020 and replaced NAFTA. The agreement tightened automotive rules of origin by raising the regional value-content requirement for passenger vehicles and light trucks and by imposing corresponding requirements on core auto parts, including transmissions. These rules increase the value of North American sourcing and may raise the effective entry barrier faced by non-North-American upstream suppliers.

Because the demand and supply estimates are based on data ending before the full implementation of the USMCA, I interpret the exercise as a prospective policy counterfactual rather than an ex-post evaluation. To motivate the relevant sourcing margin, Table 8 reports descriptive evidence from post-USMCA AALA filings. The table focuses on ICE, non-electrified cars and crossovers, excluding pickups, EVs, and PHEVs. I also exclude Korean OEMs because Hyundai, Kia, and Genesis face a separate preferential-trade margin under KORUS. The table is further restricted to firms whose North American transmission share changed by at least five percentage points between the pre-period and the post-period.

The descriptive pattern is consistent with a relocation of transmission sourcing toward North America. In the total responder set, the sales-weighted North American transmission share rises from 58.4 percent in MY 2017–2019 to 76.6 percent in MY 2022–2025, an increase of 18.2 percentage points. The largest positive changes occur at Infiniti, Buick, and Cadillac. Infiniti's North American transmission share increases by 47.9 percentage points, driven by QX50 and QX55 sourcing from Mexico. Buick's share increases by 47.6 percentage points, largely reflecting the Encore GX. Cadillac's share increases by 43.1 percentage points as the model mix shifts away from ATS and toward SUVs whose transmissions are sourced from Tennessee or Mexico. Lexus and Honda also display sizable increases, associated with the RX and HR-V, respectively. Volkswagen is a later mover. Its North American transmission share remains zero through MY 2024 and increases in MY 2025, when Jetta, Taos, and Atlas move to North American transmission sourcing.[20]

These patterns are descriptive rather than causal. They do not by themselves identify the effect of the USMCA, because transmission sourcing also responds to model redesigns, plant allocation decisions, product-cycle timing, and supplier-specific capacity constraints. Nevertheless, they show that the USMCA-relevant sourcing margin is empirically active. I therefore use the structural model to quantify how an increase in upstream concentration changes equilibrium outcomes, particularly when downstream firms face a demand contraction.

To isolate the role of upstream concentration, I consider a stylized monopoly case in which Aisin is the only available independent upstream supplier. Downstream firms may still produce transmissions in-house, but all outsourced transmissions are purchased from Aisin. This assumption avoids the assignment problem that would arise if multiple upstream suppliers were retained, because supplier choice depends on relationship-specific compatibility, engineering constraints, capacity, and contractual frictions that are not modeled. The monopoly-upstream case should therefore be interpreted as an upper-bound exercise that isolates the effect of increased upstream concentration.

Direct extraction-rate comparisons across the four-supplier baseline and the Aisin-monopoly counterfactual are not straightforward. The four-supplier counterfactual aggregates repricing responses across suppliers with different baseline prices and different recovered marginal costs, and the strategic core is distributed across multiple supplier relationships. Reassigning the strategic core to a single supplier under monopoly therefore changes both the level of upstream prices faced by each firm and the marginal-cost primitives used in the upstream first-order condition. To avoid conflating these compositional changes with the effect of concentration, I structure the comparison around a single market structure. Under Aisin monopoly, I separate the level effect of upstream pricing under normal demand from the state-dependent effect that arises when a demand contraction increases the value of outsourcing. Table 9 reports the results.

Under normal demand, allowing Aisin to reoptimize its upstream price has only a small effect on aggregate outcomes. Relative to the fixed-price Aisin-monopoly

[20] On March 26, 2025, the White House announced a 25% tariff on imported passenger vehicles and light trucks, with special treatment for USMCA vehicles. Importers could certify U.S. content, and the tariff would apply only to the non-U.S. content of USMCA vehicles.

case, reoptimization increases upstream profit by $10 million, with a corresponding small reduction in downstream profit and consumer surplus. Monopoly repricing therefore reallocates a small amount of surplus without materially affecting equilibrium quantities or consumer welfare.

The repricing channel becomes substantially more active in the bust. Holding Aisin's price fixed, the bust reduces total welfare by roughly $28.7 billion relative to the normal-demand Aisin-monopoly equilibrium, and the magnitude is similar with reoptimized prices. The interesting margin is therefore not the level of welfare loss but the change in how upstream pricing responds to the bust.

Under the bust, Aisin's repricing raises upstream profit by $64 million, with the offsetting loss falling on downstream firms and consumers. The contrast with normal demand is sharp: the same monopoly structure that transferred only $10 million from downstream to upstream under normal demand transfers $64 million during the bust. The additional repricing effect induced by the bust is therefore approximately $54 million in upstream extraction. This is the state-dependent component of the price of production flexibility under upstream concentration. When demand is normal, downstream firms have weak incentives to outsource and Aisin's profit-maximizing price is close to its baseline. When the bust raises the value of outsourcing, the monopoly supplier captures a much larger share of the resulting flexibility value.

The downstream loss from bust-induced repricing falls almost entirely on the strategic firm–transmission pairs whose sourcing decisions are solved in the Stage-2 game. The firms whose make-or-buy margins are active are precisely the firms that bear the higher price of production flexibility, while firms with predetermined sourcing assignments are largely insulated from the repricing margin.

The contrast with the four-supplier counterfactual is informative even though direct extraction-rate comparison is not clean. In the four-supplier baseline, the bust generated $62 million of upstream extraction, which represented 5.6 percent of the strategic-core flexibility gain. Under Aisin monopoly, the bust generates $64 million of upstream extraction, but $54 million of this is specifically induced by the bust rather than reflecting the level of monopoly pricing under normal demand. Concentration does not necessarily produce larger total extraction, but it concentrates the extraction in the state of the world where downstream firms are least able to absorb it. In this sense, the price of production flexibility becomes more state-dependent as the upstream market becomes more concentrated.

The Aisin-monopoly counterfactual does not generate a large additional total-welfare loss through repricing alone. Instead, it shows that upstream concentration makes the price of production flexibility more state-dependent. The policy implication is that regional-content rules can affect more than the location of production. By changing the effective concentration of upstream supplier markets, they can also change who pays for production flexibility when demand conditions deteriorate, with the cost concentrated in precisely the downturns when downstream firms most need flexibility.

7 Conclusion

This paper studies how upstream market power prices downstream production flexibility. In many supply chains, downstream firms can either produce key inputs internally or source them from specialized suppliers. Outsourcing is valuable not because firms directly value risk reduction, but because it provides an adjustment margin when demand and cost conditions move in-house production away from efficient capacity utilization. When the suppliers providing this flexibility have market power, they can price the downstream make-or-buy margin and change how shocks propagate through the vertical chain.

I quantify this mechanism in the U.S. automobile transmission market using a product-level dataset linking vehicle models to transmission suppliers from 2009 to 2018. The estimated in-house transmission cost curve is U-shaped: average costs are high at low output, fall over an intermediate range, and rise again near capacity constraints. This nonlinearity creates the production-flexibility value of outsourcing. I embed it in a three-stage model in which upstream suppliers set input prices, downstream automakers choose sourcing intensities, and vehicle prices reoptimize after product-level demand and cost shocks.

The counterfactuals show that sourcing flexibility is valuable, but its value is partly priced by upstream suppliers. In an economic bust that reduces the 2018 market size by one third, sourcing flexibility raises total welfare by $0.64 billion and increases strategic-core downstream profit by $1.10 billion when upstream prices are fixed. Once suppliers reoptimize, input prices rise by $16–$23 per transmission, equal to 40–52 percent of recovered baseline markups, shifting surplus upstream. A stylized Aisin-monopoly counterfactual shows that this pricing margin becomes more state-dependent when the upstream market is more concentrated: repricing is small under normal demand but much larger during the bust, with the burden falling mainly on firms active on the make-or-buy margin.

The broader lesson is that supplier market power matters not only through standard double marginalization, but also through the pricing of adjustment margins. Treating input prices as fixed can miss how upstream suppliers respond when downstream firms' need for external production changes. This mechanism extends beyond automobile transmissions to other critical inputs, such as batteries, semiconductors, power electronics, and technology-intensive components, whenever downstream firms retain some internal production capability and specialized upstream suppliers are concentrated.

References

Acemoglu, D., V. M. Carvalho, A. Ozdaglar, and A. Tahbaz-Salehi (2012). The network origins of aggregate fluctuations. *Econometrica 80*(5), 1977–2016.

Acemoglu, D., S. Johnson, J. Robinson, and Y. Thaicharoen (2003). Institutional causes, macroeconomic symptoms: volatility, crises and growth. *Journal of Monetary Economics 50*(1), 49–123.

Acemoglu, D., A. Ozdaglar, and A. Tahbaz-Salehi (2017). Microeconomic origins of macroeconomic tail risks. *American Economic Review 107*(1), 54–108.

Alfaro, L. and D. Chor (2023, September). Global supply chains: The looming "great reallocation". Working Paper 31661, National Bureau of Economic Research.

Alfaro, L., D. Chor, P. Antras, and P. Conconi (2019). Internalizing global value chains: A firm-level analysis. *Journal of political economy 127*(2), 508–559.

Barrot, J.-N. and J. Sauvagnat (2016). Input Specificity and the Propagation of Idiosyncratic Shocks in Production Networks. *The Quarterly Journal of Economics 131*(3), 1543–1592.

Belloni, A., D. Chen, V. Chernozhukov, and C. Hansen (2012). Sparse models and methods for optimal instruments with an application to eminent domain. *Econometrica 80*(6), 2369–2429.

Berry, S., J. Levinsohn, and A. Pakes (1995). Automobile prices in market equilibrium. *Econometrica 63*(4), 841–890.

Berry, S. T. (1994). Estimating discrete-choice models of product differentiation. *The RAND Journal of Economics 25*(2), 242–262.

Bilal, A. and H. Lhuillier (2021). Outsourcing, inequality and aggregate output. Technical report, National Bureau of Economic Research.

Bloom, N. (2014). Fluctuations in uncertainty. *Journal of Economic Perspectives 28*(2), 153–76.

Carlton, D. W. (1979). Vertical integration in competitive markets under uncertainty. *The Journal of Industrial Economics 27*(3), 189–209.

Carvalho, V. M., M. Nirei, Y. U. Saito, and A. Tahbaz-Salehi (2020). Supply Chain Disruptions: Evidence from the Great East Japan Earthquake*. *The Quarterly Journal of Economics 136*(2), 1255–1321.

Chipty, T. and C. Snyder (1999). The role of firm size in bilateral bargaining: A study of the cable television industry. *The Review of Economics and Statistics 81*(2), 326–340.

Collard-Wexler, A. (2013). Demand fluctuations in the ready-mix concrete industry. *Econometrica 81*(3), 1003–1037.

Council, N. R. (2015). *Cost, Effectiveness, and Deployment of Fuel Economy Technologies for Light-Duty Vehicles*. Washington, DC: The National Academies Press.

Crawford, G. S., R. S. Lee, M. D. Whinston, and A. Yurukoglu (2018). The welfare effects of vertical integration in multichannel television markets. *Econometrica 86*(3), 891–954.

Crawford, G. S. and A. Yurukoglu (2012). The welfare effects of bundling in multichannel television markets. *The American Economic Review 102*(2), 643–685.

di Giovanni, J. and A. Levchenko (2009). Trade openness and volatility. *The Review of Economics and Statistics 91*(3), 558–585.

Duarte, M., L. Magnolfi, D. Quint, M. Sølvsten, and C. Sullivan (2025). Conduct and scale economies: Evaluating tariffs in the us automobile market.

Dubé, J.-P., J. T. Fox, and C.-L. Su (2012). Improving the numerical performance of static and dynamic aggregate discrete choice random coefficients demand estimation. *Econometrica 80*(5), 2231–2267.

Efron, B. (2016). Empirical bayes deconvolution estimates. *Biometrika 103*(1), 1–20.

Espín-Sánchez, J.-A., Á. Parra, and Y. Wang (2023). Equilibrium uniqueness in entry games with private information. *The RAND Journal of Economics 54*(3), 512–540.

Gandhi, A. and J.-F. Houde (2019). Measuring substitution patterns in differentiated-products industries. Working Paper 26375, National Bureau of Economic Research.

Garetto, S. (2013). Input sourcing and multinational production. *American Economic Journal: Macroeconomics 5*(2), 118–51.

Goldberg, P. K. and F. Verboven (2001). The Evolution of Price Dispersion in the European Car Market. *The Review of Economic Studies 68*(4), 811–848.

Gowrisankaran, G., A. Nevo, and R. Town (2015). Mergers when prices are negotiated: Evidence from the hospital industry. *American Economic Review 105*(1), 172–203.

Green, J. R. (1986). *Vertical Integration and Assurance of Markets*. MIT Press.

Grennan, M. and A. Swanson (2020). Transparency and negotiated prices: The value of information in hospital-supplier bargaining. *Journal of Political Economy 128*(4), 1234–1268.

Grieco, P. L. E., C. Murry, and A. Yurukoglu (2021). The evolution of market power in the us auto industry. Working Paper 29013, National Bureau of Economic Research.

Grossman, G. M., E. Helpman, and S. J. Redding (2024). When tariffs disrupt global supply chains. *American Economic Review 114*(4), 988–1029.

Heeney, L., K. Christopher R., and M. Jasdeep (2026). Tariffs, global value chains, and the incidence of protection: Evidence from us automobiles. Technical report, National Bureau of Economic Research.

Ho, K. and R. S. Lee (2017). Insurer competition in health care markets. *Econometrica 85*(2), 379–417.

Klein, B. (1988). Vertical integration as organizational ownership: The fisher body-general motors relationship revisited. *Journal of Law, Economics, Organization 4*(1), 199–213.

Klein, B. (2000). Fisher—general motors and the nature of the firm. *The Journal of Law and Economics 43*(1), 105–142.

Klein, B., R. G. Crawford, and A. A. Alchian (1978). Vertical integration, appropriable rents, and the competitive contracting process. *The Journal of Law Economics 21*(2), 297–326.

Lafontaine, F. and M. Slade (2007). Vertical integration and firm boundaries: The evidence. *Journal of Economic Literature 45*(3), 629–685.

Langlois, R. N. and P. L. Robertson (1989). Explaining vertical integration: Lessons from the american automobile industry. *The Journal of Economic History 49*(2), 361–375.

Loertscher, S. and M. H. Riordan (2019). Make and buy: Outsourcing, vertical integration, and cost reduction. *American Economic Journal: Microeconomics 11*(1), 105–23.

Long, J. B. and C. I. Plosser (1983). Real business cycles. *Journal of Political Economy 91*(1), 39–69.

Masten, S. E., J. W. Meehan, and E. A. Snyder (1989). Vertical integration in the u.s. auto industry: A note on the influence of transaction specific assets. *Journal of Economic Behavior Organization 12*(2), 265–273.

Monteverde, K. and D. J. Teece (1982). Supplier switching costs and vertical integration in the automobile industry. *The Bell Journal of Economics 13*(1), 206–213.

Mueller, F., K. O. Stahl, and F. Wachtler (2008). Upstream relationships in the automotive industry: A contractual perspective. *University of Mannheim*.

Nurski, L. and F. Verboven (2016). Exclusive Dealing as a Barrier to Entry? Evidence from Automobiles. *The Review of Economic Studies 83*(3), 1156–1188.

Petrin, A. (2002). Quantifying the benefits of new products: The case of the minivan. *Journal of Political Economy 110*(4), 705–729.

Ramey, V. A. and D. J. Vine (2006, December). Declining volatility in the u.s. automobile industry. *American Economic Review 96*(5), 1876–1889.

Rust, J. (1994). Structural estimation of markov decision processes. *Handbook of econometrics 4*, 3081–3143.

Seim, K. (2006). An empirical model of firm entry with endogenous product-type choices. *The RAND Journal of Economics 37*(3), 619–640.

Stigler, G. J. (1951). The division of labor is limited by the extent of the market. *Journal of Political Economy 59*(3), 185–193.

Taylor, C. R. and S. N. Wiggins (1997). Competition or compensation: Supplier incentives under the american and japanese subcontracting systems. *The American Economic Review 87*(4), 598–618.

Villas-Boas, S. B. (2007). Vertical relationships between manufacturers and retailers: Inference with limited data. *The Review of Economic Studies 74*(2), 625–652.

Walker, G. and D. Weber (1984). A transaction cost approach to make-or-buy decisions. *Administrative Science Quarterly 29*(3), 373–391.

Wollmann, T. G. (2018). Trucks without bailouts: Equilibrium product characteristics for commercial vehicles. *American Economic Review 108*(6), 1364–1406.

Yang, C. (2020). Vertical structure and innovation: A study of the soc and smartphone industries. *The RAND Journal of Economics 51*(3), 739–785.

Figure 1: Light Vehicle Sales in the U.S., 1975–2020

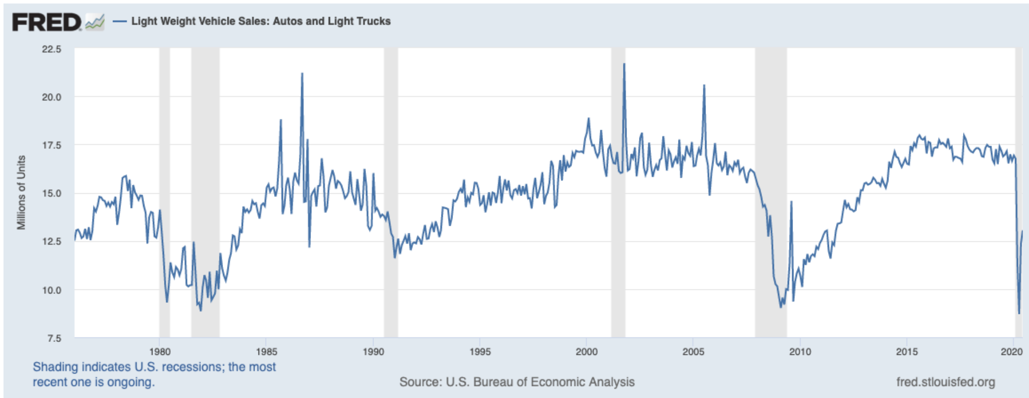

Notes: This figure is downloaded from the St. Louis Fed and is based on statistics from the U.S. Bureau of Economic Analysis. It shows passenger vehicle sales in the U.S. from 1975 to 2020.

Figure 2: Supplier Switching and In-House Transmission Production

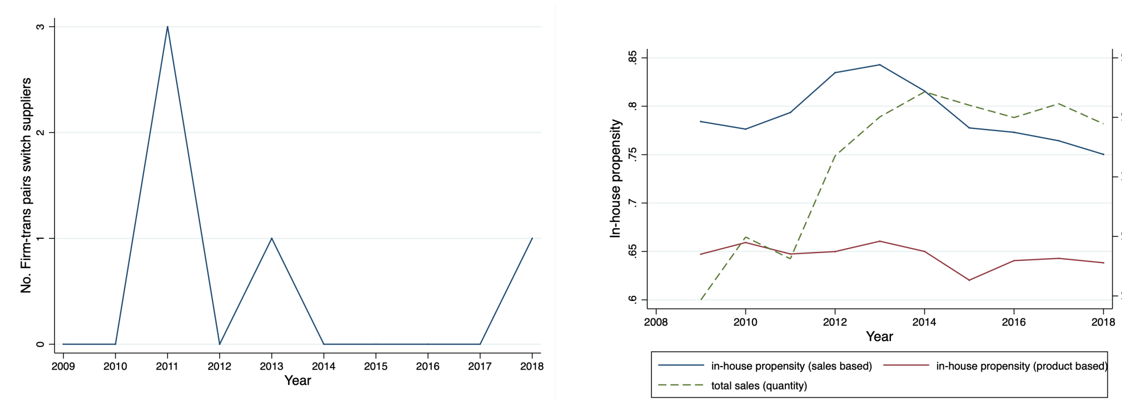

(a) Supplier switching

(b) In-house production share

Notes: Panel (a) shows the number of downstream firm–transmission pairs that switch to a different upstream supplier in each year. The maximum switching fraction is less than 4%. Panel (b) shows the evolution of in-house transmission production. The dashed line is annual output, and the solid lines report the fraction of cars or products using in-house transmissions.

Figure 3: Diagnostics for Recovered Demand Shocks

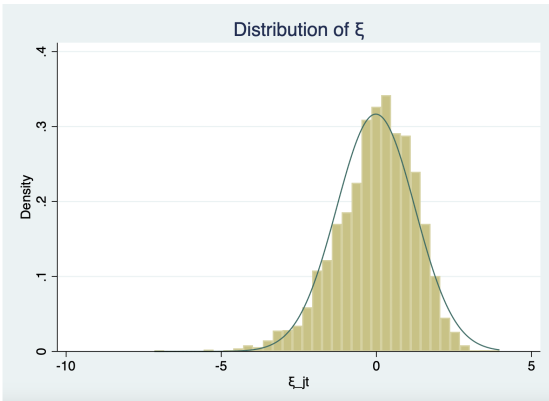

(a) Overall distribution

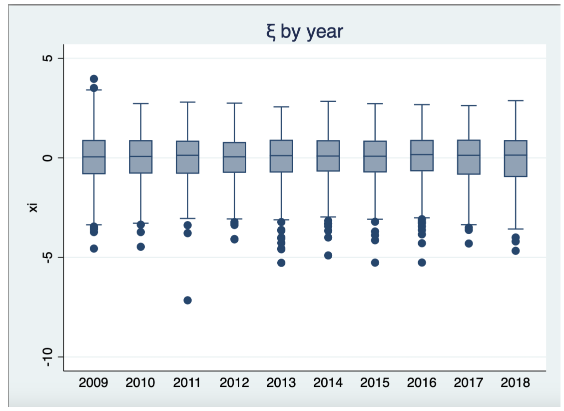

(b) Distribution by year

Notes: Panel (a) plots the density of recovered product-level demand shocks $\hat{\xi}_{jt}$. Panel (b) plots the distribution of $\hat{\xi}_{jt}$ by year. Demand shocks are recovered from the random-coefficients demand system after controlling for observed vehicle characteristics, brand fixed effects, and year fixed effects.

Figure 4: Estimated Transmission Prices and In-House Cost Differences

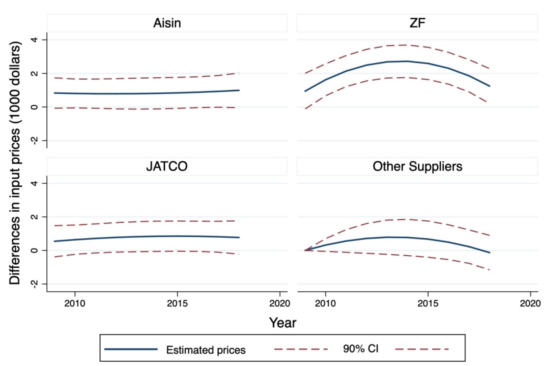

(a) Estimated transmission price differences

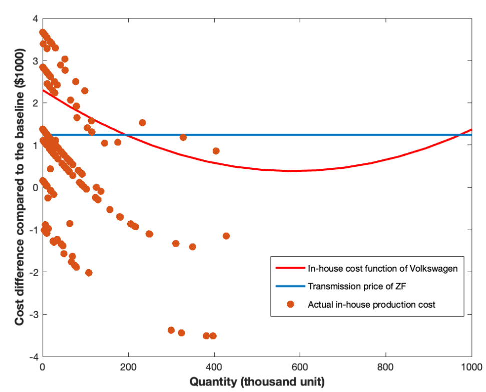

(b) In-house versus outsourcing cost

Notes: Panel (a) shows estimated transmission price differences relative to the other-supplier group in 2009, τ_O. Estimated transmission prices are computed as $\tau_{st} = \hat{\tau}_s + \hat{\tau}_s^{trend} t + \hat{\tau}_s^{trend2} t^2$, where $t = \text{year} - 2009$. Panel (b) plots the cost difference between in-house and outsourced transmission production for selected downstream manufacturers. The red curve is drawn from Volkswagen's estimated in-house cost function. Actual in-house production costs are computed using the estimates and realized equilibrium demand. The blue horizontal line is the equilibrium price of ZF in 2018.

Figure 5: Upstream Suppliers' Marginal Cost Differences

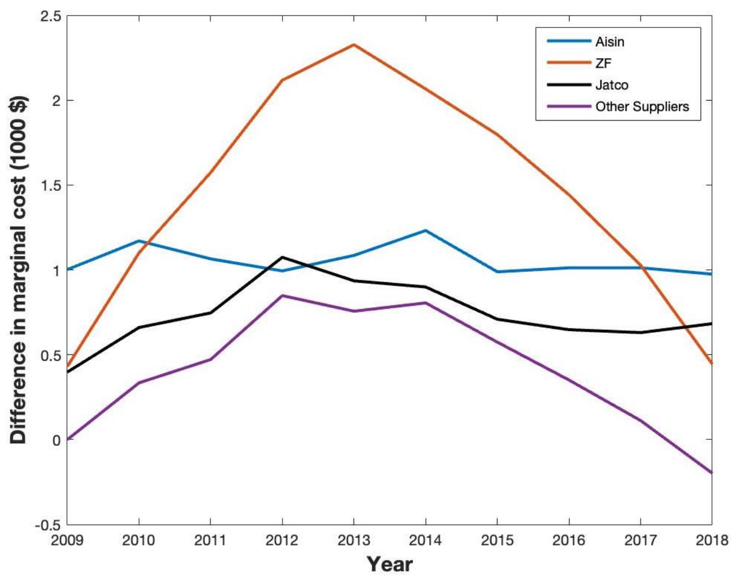

Notes: This figure plots relative upstream marginal costs because the transmission price index uses the other-supplier group's price in 2009 as the normalized base group. Marginal costs are inverted from Equation 7 using the estimated equilibrium upstream prices and parameter primitives in each year.

Table 1: Summary Statistics for Transmission Suppliers

Transmission type	Speed	Transmission share	Supplier	Conditional supplier share
AT	A4	0.168	Aisin	0.45
			JATCO	0.55
	A5	0.183	Aisin	0.66
			JATCO	0.34
	A6	0.465	Aisin	0.86
			ZF	0.14
			JATCO	0.004
	A7	0.025	JATCO	1.00
	A8	0.108	Aisin	0.24
			ZF	0.76
	A9	0.058	ZF	1.00
	A10	0.037	Aisin	1.00
CVT	CVT	0.158	Aisin	0.16
			JATCO	0.84
DCT	DCT6	0.024	Getrag	1.00
	DCT7	0.013	ZF	0.78
			Getrag	0.22
	DCT8	0.001	ZF	1.00
	DCT9	0.00003		
MT	M5	0.024	Aisin	0.34
			Getrag	0.55
			Tremec	0.11
	M6	0.031	Aisin	0.09
			ZF	0.11
			Eaton	0.01
			Getrag	0.46
			Tremec	0.34
	M7	0.001	ZF	0.46
			Tremec	0.54

Notes: This table reports transmission types and supplier concentration in the transmission market. Shares are calculated using the quantity sold. Conditional supplier share is computed among outsourced transmissions within each transmission category.

Table 2: Summary Statistics for Main Variables

Variable	Observations	Mean	Std. Dev.	Min	Max
Market share	3,848	0.0003	0.0006	2.31×10^{-8}	0.0075
Products per year	3,848	384.80	17.42	370	428
Product characteristics					
Price (10^3)	3,848	16.70	9.46	5.14	63.60
Horsepower (10)	3,848	24.73	9.58	7	65
Fuel efficiency (10)	3,848	3.15	1.08	1.31	15.76
Length (10 cm)	3,848	18.54	1.66	10.61	25.45
Foreign	3,848	0.47	0.50	0	1
Pickup	3,848	0.06	0.24	0	1
SUV	3,848	0.31	0.46	0	1
Van	3,848	0.04	0.20	0	1
Transmission characteristics					
CVT	3,848	0.10	0.30	0	1
DCT	3,848	0.07	0.25	0	1
MT	3,848	0.25	0.43	0	1
Transmission low speed	3,848	0.23	0.42	0	1
Transmission high speed	3,848	0.22	0.41	0	1
Transmission supplier					
Aisin	3,848	0.12	0.33	0	1
ZF	3,848	0.09	0.29	0	1
JATCO	3,848	0.08	0.27	0	1
Other suppliers	3,848	0.06	0.24	0	1
In-house	3,848	0.65	0.48	0	1
Micro moment: real income					
New car purchase (10^4)	20,751	4.18			
No new car purchase (10^4)	302,788	2.85			

Notes: This table reports summary statistics for make–model–transmission-year observations in the sample. Each product is defined as a make–model–transmission-year combination. I follow Berry et al. (1995) and use total U.S. households as market size. Prices are adjusted for inflation. Fuel efficiency is defined as miles per dollar following Berry et al. (1995). The demand estimation uses model years 2009–2018 and includes 3,848 products. I exclude vehicles without transmissions, which are mainly electric vehicles.

Table 3: Demand Estimation Results

Variable	IV	BLP	Variable	IV	BLP
Vehicle			*Transmission*		
Constant	-9.118	-46.634	Low speed	-0.486	-0.421
	(1.036)	(0.411)		(0.123)	(0.119)
Constant$\times v$		27.414	High speed	-0.116	-0.112
		(0.551)		(0.147)	(0.158)
Price	-0.049	-0.268	CVT	-0.502	-0.271
	(0.020)	(0.070)		(0.314)	(0.273)
Price$\times \log(\text{income})$		0.081	DCT	-1.322	-1.037
		(0.018)		(0.195)	(0.199)
Horsepower	-0.009	0.060	MT	-2.012	-2.012
	(0.018)	(0.025)		(0.124)	(0.138)
Fuel efficiency	0.077	0.237	Aisin	-0.486	0.149
	(0.080)	(0.096)		(0.149)	(0.155)
Size (length)	0.093	0.216	ZF	0.003	0.622
	(0.049)	(0.061)		(0.258)	(0.209)
Pickup	-0.168	-0.634	JATCO	0.463	0.455
	(0.286)	(0.286)		(0.305)	(0.254)
SUV	0.108	0.296	Other supplier	-0.146	0.224
	(0.118)	(0.124)		(0.269)	(0.187)
Van	-0.256	-0.347			
	(0.213)	(0.235)			
Foreign	-0.832	-0.652			
	(0.140)	(0.155)			
Observations			3,848		
Year fixed effects			Yes		
Company fixed effects			Yes		

Notes: This table reports logit and random-coefficients demand estimates. For unobserved heterogeneity and demographics, I use a product rule with level 12. Standard errors, clustered at the product level, are in parentheses. For the random-coefficients model, I use `pyblp` with optimal instruments; the feasibility and optimality tolerances are both 10^{-6}, as in Dubé et al. (2012).

Table 4: Price Elasticities, Marginal Costs, and Margins

Variable	Mean	Std. Dev.	10%	Median	90%
Price (10^3)	16.695	9.455	8.439	14.050	27.749
Own-price elasticity	-2.579	0.495	-3.234	-2.607	-1.894
Marginal cost (10^3)	10.118	5.873	3.969	8.662	19.103
Margin	0.413	0.088	0.315	0.392	0.541

Notes: This table reports summary statistics for own-price elasticities, marginal costs, and margins. Margin is defined as $1 -$ marginal cost/price.

Table 5: Supply Estimation Results

Vehicle		Transmission	
log(hp)	12.398	c_1	0.829
	(0.294)		(0.513)
log(mpg)	4.435	c_2	-13.125
	(0.443)		(2.363)
log(size)	7.339	c_3	16.873
	(0.896)		(5.650)
Foreign	1.288	$c_{1,DA}$	-3.604
	(0.104)		(0.673)
Labor cost	-0.011	$c_{1,DA}^{trend}$	-0.233
	(0.004)		(0.083)
Pickup	-2.226	$c_{1,HY}$	1.470
	(0.192)		(0.685)
SUV	0.565	$c_{1,VW}$	-1.214
	(0.099)		(0.327)
Van	-0.957	$c_{1,HO}$	2.292
	(0.186)		(0.513)
Observations		3,848	
R-squared		0.840	
Company and year fixed effects		Yes	

Notes: Here c_1, c_2, and c_3 are in-house transmission cost-function parameters. I fit a third-order polynomial and allow for heterogeneity among selected downstream firms: $c(N_t g_{jt}) = c_{1jt}(N_t g_{jt}) + c_2(N_t g_{jt})^2 + c_3(N_t g_{jt})^3$. DA stands for Daimler Group, HO for Honda, HY for Hyundai, and VW for Volkswagen.

Table 6: Decomposing the Value and Price of Sourcing Flexibility in an Economic Bust

Scenario	CS	DS π	Strategic π	Upstream π	Total welfare
Levels					
Baseline	8.814	86.550	14.894	0.195	95.560
Scenario 1	8.782	56.874	8.732	0.087	65.742
Scenario 2	8.807	57.437	9.830	0.133	66.378
Scenario 3	8.806	57.422	9.813	0.195	66.424
Decomposition					
Direct bust effect:	-0.032	-29.677	-6.163	-0.108	-29.817
Flexibility value:	0.026	0.564	1.098	0.046	0.635
Repricing effect	-0.001	-0.015	-0.016	0.062	0.046
Net bust effect:	-0.008	-29.128	-5.081	-0.000	-29.136

Notes: This table decomposes the effect of an economic bust in the 2018 market. The bust reduces market size by one third. Scenario 1 is "No strategic sourcing flexibility" forces the strategic firm–transmission pairs to use the full-in-house action while non-strategic products keep their observed sourcing assignments. Scenario 2 is "Sourcing flexibility, fixed τ" allows the Stage-2 sourcing equilibrium to adjust while holding upstream prices fixed at baseline values. Scenario 3 is "Reoptimized τ" allows upstream suppliers to reset prices using the Stage-1 first-order conditions and the recovered 2018 upstream marginal costs. Strategic-core profit is the downstream profit of products belonging to the firm–transmission pairs whose sourcing decisions are solved in the Stage-2 game. Direct bust effect measures no strategic flexibility vs. baseline comparison. Production-flexibility value is when we hold τ as fixed. Net bust effect is when we compare reoptimized τ vs. baseline. Downstream-side surplus is consumer surplus plus downstream profit. Monetary values are in billions of 2015 dollars.

Table 7: Upstream Price Response Relative to Baseline Markups

Supplier	Baseline markup ($)	$\Delta\tau$ ($)	$\Delta\tau$ / baseline markup
Aisin	40.7	16.4	40.2%
ZF	43.8	22.8	52.1%
JATCO	33.8	17.6	52.1%
Other	37.4	17.7	47.3%

Notes: Baseline markup is $\tau_{0s} - mc_s^U$. $\Delta\tau$ is the change from baseline upstream prices to reoptimized upstream prices in the economic-bust counterfactual. Prices are reported in 2015 dollars.

Table 8: Change in North American Transmission Sourcing

Make	NA transmission share			Avg. annual sales	
	Pre (%)	Post (%)	Δ (pp)	Pre units	Post units
Panel A. Firms moving toward North American sourcing					
Infiniti	28.0	75.9	47.9	83,800	52,200
Buick	36.3	83.9	47.6	203,900	100,400
Cadillac	56.9	100.0	43.1	149,200	79,600
Lexus	5.6	28.8	23.3	255,800	305,900
Honda	71.7	93.7	22.0	425,300	340,900
Volkswagen	0.0	18.0	18.0	111,400	99,200
GMC	85.6	100.0	14.4	238,300	153,600
Nissan	82.1	96.0	13.9	555,500	394,400
Acura	85.4	96.1	10.7	72,300	42,800
Subtotal, movers toward NA	**58.4**	**76.6**	**18.2**	**2,095,500**	**1,569,000**

Notes: The table reports sales-weighted average shares of U.S. calendar-year units whose transmission is sourced from the United States, Canada, or Mexico. The sample is restricted to ICE, non-electrified passenger cars and crossovers; pickups, EVs, and PHEVs are excluded. The pre-period is MY 2017–2019 and the post-period is MY 2022–2025. The table includes firms whose North American transmission share changed by at least 5 percentage points between the two windows. Hyundai, Kia, and Genesis are excluded because Korean OEMs face a separate preferential-trade margin under KORUS.

Table 9: Aisin Monopoly: Repricing under Normal Demand and an Economic Bust

Scenario	CS	DS π	Strategic DS π	Non-strategic DS π	Upstream π	Total welfare
Levels						
Normal demand, fixed τ	8.719	85.273	14.060	71.213	0.151	94.144
Normal demand, reoptimized τ	8.719	85.271	14.059	71.212	0.161	94.151
Bust, fixed τ	8.713	56.580	9.254	47.326	0.107	65.400
Bust, reoptimized τ	8.712	56.563	9.239	47.324	0.170	65.446
Key decomposition						
Repricing effect, normal demand	-0.000	-0.003	-0.001	-0.001	0.010	0.007
Repricing effect, bust	-0.001	-0.017	-0.015	-0.002	0.064	0.046
Additional repricing effect in bust	-0.001	-0.014	-0.014	-0.001	0.054	0.039

Notes: This table decomposes upstream repricing under the Aisin-monopoly counterfactual. In this counterfactual, downstream firms may still produce transmissions in-house, but all outsourced transmissions are assigned to Aisin. Normal demand uses the baseline 2018 market size. The bust reduces market size by one third. "Fixed τ" holds Aisin's upstream price at its baseline value. "Reoptimized τ" allows Aisin to reset its price using the Stage-1 first-order condition and the recovered 2018 upstream marginal cost. Strategic downstream profit is the profit of products belonging to firm–transmission pairs whose sourcing choices are solved in the Stage-2 game. Non-strategic downstream profit is the profit of all remaining products, whose sourcing assignments are held fixed but whose vehicle prices and quantities adjust in equilibrium. Monetary values are in billions of 2015 dollars.

Appendix A A Simple Model of Production Flexibility and Upstream Pricing

This appendix presents a simple model that clarifies the economic mechanism behind the full empirical model. The purpose is not to reproduce the full estimation environment, but to show why outsourcing can be valuable under nonlinear in-house production costs and why upstream suppliers with market power price this flexibility.

A.1 Environment

There are two downstream firms and two upstream suppliers. Each downstream firm produces one final product. Firm i's inverse demand is

$$p_i = \delta_i - \alpha q_i - \eta q_{-i}, \qquad \alpha > \eta > 0, \tag{A.1}$$

where $\delta_i = X_i + \xi_i$ contains observable product quality X_i and a demand shock ξ_i. Equivalently, demand can be written as $q_i(p, \xi)$.

Each downstream firm chooses whether to produce the input in-house or outsource it. Let $I_i \in \{0, 1\}$, where $I_i = 1$ denotes in-house production. If the input is outsourced, firm i pays upstream price $\tau_{s(i)}$. If it is produced in-house, the input cost is $c(q_i)$. Product profit is

$$\pi_i = q_i \left[p_i - mc_i - (1 - I_i)\tau_{s(i)} \right] - I_i c(q_i), \tag{A.2}$$

where mc_i is the marginal cost of all other vehicle components. To illustrate the role of nonlinear in-house costs, suppose

$$c(q_i) = c_1 q_i + c_2 q_i^2. \tag{A.3}$$

The timing mirrors the full model. Upstream suppliers set input prices. Downstream firms choose sourcing before the demand shock is realized, based on expected profits. After ξ is realized, downstream firms set final-good prices.

A.2 Outsourcing as Production Flexibility

For a given realization of demand, outsourcing is attractive when the upstream price is lower than the relevant in-house production cost. Abstracting from strategic price effects for exposition, the state-by-state value of outsourcing is

$$\Delta_i(\xi; \tau) = \pi_i^O(\xi; \tau) - \pi_i^I(\xi) = q_i(\xi) \left[\frac{c(q_i(\xi))}{q_i(\xi)} - \tau_{s(i)} \right]. \tag{A.4}$$

Thus, outsourcing is valuable in states where realized output places the downstream firm in a region of the in-house cost curve where production is relatively inefficient.

The expected value of having access to outsourcing is

$$E_{\xi}[\Delta_i(\xi;\tau)] = E_{\xi}\left[q_i(\xi)\left(c^{AC}(q_i(\xi)) - \tau_{s(i)}\right)\right], \quad \text{(A.5)}$$

where $c^{AC}(q) = c(q)/q$. Demand fluctuations matter because they change the distribution of realized output $q_i(\xi)$. They do not have a mechanical sign. If more probability mass falls in regions where in-house production is expensive relative to outsourcing, the outsourcing option becomes more valuable. If more probability mass falls in regions where in-house production is efficient, the outsourcing option becomes less valuable.

This distinction is important. The model does not rely on risk aversion. Firms are risk neutral and choose sourcing to maximize expected profits. The source of value is the interaction between demand realizations and nonlinear in-house production costs.

A.3 A Private-Information Sourcing Game

Now suppose each downstream firm has private action-specific payoff shocks $\epsilon_i(I_i)$. These shocks are i.i.d. Type-I extreme value. Let

$$E\pi_i(I_i, I_{-i}, \tau)$$

be firm i's expected equilibrium profit after integrating over demand shocks and downstream price competition. If firm i believes that its rival chooses in-house with probability σ_{-i}, then the deterministic value of choosing $I_i = k$ is

$$E\Pi_i(k;\sigma_{-i},\tau) = \sum_{I_{-i}\in\{0,1\}} E\pi_i(k, I_{-i}, \tau)Pr(I_{-i}|\sigma_{-i}). \quad \text{(A.6)}$$

The logit best response is

$$\sigma_i = Pr(I_i = 1) = \frac{\exp\left(E\Pi_i(1;\sigma_{-i},\tau)/\lambda_{\epsilon}\right)}{\exp\left(E\Pi_i(1;\sigma_{-i},\tau)/\lambda_{\epsilon}\right) + \exp\left(E\Pi_i(0;\sigma_{-i},\tau)/\lambda_{\epsilon}\right)}. \quad \text{(A.7)}$$

A Bayesian Nash equilibrium is a fixed point

$$\sigma^* = \Psi(\sigma^*, \tau). \quad \text{(A.8)}$$

The response of sourcing probabilities to upstream prices has two components. First, changing τ directly changes the expected payoff from outsourcing. Second, changing τ changes rivals' sourcing probabilities, which affects expected downstream competition. Differentiating the fixed point gives

$$\frac{d\sigma^*}{d\tau} = \left[I - \frac{\partial\Psi}{\partial\sigma}\right]^{-1} \frac{\partial\Psi}{\partial\tau}. \tag{A.9}$$

This is the same logic used in the full model to compute how equilibrium sourcing probabilities respond to upstream prices.

A.4 Upstream Pricing and Shock Propagation

Let $Q_s^O(\tau)$ be the expected outsourced demand faced by supplier s, taking into account downstream price competition and equilibrium sourcing probabilities. Supplier s's profit is

$$\Pi_s^U(\tau) = (\tau_s - mc_s^U)Q_s^O(\tau). \tag{A.10}$$

The first-order condition is

$$Q_s^O(\tau) + (\tau_s - mc_s^U)\frac{\partial Q_s^O(\tau)}{\partial \tau_s} = 0. \tag{A.11}$$

When downstream firms' incentive to outsource rises, $Q_s^O(\tau)$ shifts outward. A supplier with market power raises τ_s, extracting part of the production-flexibility value generated by outsourcing. This is the shock-propagation channel: demand or cost shocks change the value of downstream sourcing flexibility, and upstream market power changes the price of accessing that flexibility.

A.5 Implications

The simple model delivers three implications that guide the empirical analysis.

First, outsourcing is not valuable because firms directly value risk reduction. With risk-neutral firms, outsourcing is valuable because it can improve expected profits when demand realizations place firms in output regions where in-house production is inefficient.

Second, the effect of demand volatility is theoretically ambiguous. It depends on how the induced distribution of output overlaps with the in-house cost curve and the upstream price. This motivates the empirical analysis of the recovered demand shocks and the location of realized quantities on the estimated in-house cost curves.

Third, upstream market power affects not only the static level of input prices but also the transmission of shocks. When downstream firms' desire to outsource increases, concentrated upstream suppliers can raise prices and absorb part of the gains from production flexibility.

Appendix B Algorithm for Solving the Model

This appendix describes the computational procedure used to solve the model and to recover upstream marginal costs. The model is solved backward. For a given upstream price vector and sourcing assignment, I first solve the downstream pricing equilibrium. I then use these pricing outcomes to evaluate payoffs in the Stage-2 sourcing game. Finally, I aggregate expected outsourced demand across the Stage-2 equilibrium distribution and invert the upstream first-order conditions.

B.1 Stage 3: Downstream Pricing Equilibrium

For a given year t, upstream price vector τ_t, product-level sourcing assignment I_t, and shock realization $e_t = (\xi_t, \omega_t)$, downstream firms set vehicle prices in a differentiated-products Bertrand game. Let $g_j(p_t, e_t)$ denote the market share of product j. The Stage-3 first-order condition for product j is

$$g_j(p_t, e_t) + \sum_{k \in \mathcal{J}_t} \Gamma_{jk,t} \left[p_k - mc_k(g_k, \tau_t, I_t, e_t) \right] \frac{\partial g_k(p_t, e_t)}{\partial p_j} = 0. \qquad \text{(A.12)}$$

Here $\Gamma_{jk,t} = 1$ if products j and k are priced by the same downstream pricing unit and zero otherwise. The marginal cost $mc_k(\cdot)$ includes the vehicle cost component, the product-level marginal-cost shock, the upstream input price if the product is outsourced, and the nonlinear in-house transmission cost if the product is produced internally.

I solve the system in (A.12) by fixed-point iteration with damping and step-size controls. For each solved equilibrium, I store equilibrium prices, market shares, product profits, consumer surplus, and supplier-level outsourced demand. I drop draws only when the pricing equilibrium fails to converge or produces invalid price/profit objects. Derivatives used for the upstream first-order condition are computed after the final price equilibrium is solved and do not affect the validity of the pricing draw.

B.2 Derivatives with Respect to Upstream Prices

To compute the upstream first-order condition, I need the derivative of product shares and product profits with respect to upstream prices. Let D_p be the matrix of demand derivatives with element

$$(D_p)_{kj} = \frac{\partial g_k(p_t, e_t)}{\partial p_j}.$$

Let

$$A_t = \Gamma_t \circ D_p,$$

where $\circ$ denotes element-by-element multiplication. Let Z_t be the product-by-supplier assignment matrix with element $Z_{js} = 1$ if product j is outsourced to supplier s, and zero otherwise. The direct derivative of marginal cost with respect to upstream prices is Z_t.

Differentiating the downstream pricing first-order conditions gives

$$\frac{\partial p_t^*}{\partial \tau_t} = \left(\frac{\partial F_t}{\partial p_t}\right)^{-1} A_t Z_t, \tag{A.13}$$

where $F_t(p_t, \tau_t)$ denotes the vector of downstream pricing first-order conditions. The Jacobian $\partial F_t / \partial p_t$ includes the demand derivative terms, the demand Hessian terms, and the derivative of in-house marginal cost with respect to market shares. Product-share derivatives are then

$$g_{\tau,t} = D_p \frac{\partial p_t^*}{\partial \tau_t}. \tag{A.14}$$

For supplier s's outsourced demand under a given action profile a_t, the derivative with respect to upstream price $\tau_{\ell t}$ is

$$\frac{\partial Q_{st}(a_t, \tau_t)}{\partial \tau_{\ell t}} = M_t \sum_{j \in \mathcal{J}_t} Z_{js,t}(a_t) g_{\tau,j\ell,t}(a_t). \tag{A.15}$$

This expression uses the supplier assignment mask $Z_{js,t}$. I do not set $g_{\tau,j\ell,t}$ to zero for in-house products before computing derivatives: in-house products can have nonzero share responses to upstream prices through substitution with outsourced products. The supplier mask determines which products enter supplier s's outsourced demand.

B.3 Simulation over Assignments and Shocks

For each Stage-2 action profile a_t, I simulate product-level sourcing assignments and product-level shock realizations. Let $n = 1, \ldots, N_A$ index assignment draws and $m = 1, \ldots, N_E$ index demand/cost shock draws. For each pair (n, m), I solve the Stage-3 pricing equilibrium and compute product profits $\pi_{jt}^*(a_t, n, m)$, product shares $g_{jt}^*(a_t, n, m)$, and supplier-level outsourced demand.

The expected product profit under action profile a_t is

$$E\pi^*_{jt}(a_t, \tau_t) = \frac{1}{N_A N_E} \sum_{n=1}^{N_A} \sum_{m=1}^{N_E} \pi^*_{jt}(a_t, n, m; \tau_t). \tag{A.16}$$

The expected payoff for firm–transmission pair (f, h) is

$$E\pi_{fht}(a_t, \tau_t) = \sum_{j \in \mathcal{J}_{fht}} E\pi^*_{jt}(a_t, \tau_t). \tag{A.17}$$

Similarly, supplier s's expected outsourced demand under action profile a_t is

$$Q_{st}(a_t, \tau_t) = \frac{M_t}{N_A N_E} \sum_{n=1}^{N_A} \sum_{m=1}^{N_E} \sum_{j \in \mathcal{J}_t} Z^n_{js,t}(a_t) g^*_{jt}(a_t, n, m; \tau_t). \tag{A.18}$$

The demand shock draws are generated from the residual shock process described in Section 4.2. I use common random numbers across counterfactuals: the same pre-generated demand and cost shock panels are used when comparing equilibria. This reduces simulation noise in counterfactual differences.

B.4 Stage 2: Sourcing Equilibrium

For each year and upstream price vector τ_t, I evaluate all action profiles in the strategic core. The payoff matrix from (A.17) is used to solve the private-information sourcing game.

Let $\sigma_{fht}(a)$ denote the probability that firm–transmission pair (f, h) chooses action a. Given a strategy profile σ_t, the deterministic expected value of action a is

$$E\Pi_{fht}(a, \tau_t; \sigma_{-fht}) = \sum_{a_{-fht}} E\pi_{fht}(a, a_{-fht}, \tau_t) \prod_{(f',h') \neq (f,h)} \sigma_{f'h't}(a_{f'h'}). \tag{A.19}$$

The logit best response is

$$\sigma_{fht}(a) = \frac{\exp\left(E\Pi_{fht}(a, \tau_t; \sigma_{-fht})/\lambda_\epsilon\right)}{\sum_{a' \in \mathcal{A}_{fht}} \exp\left(E\Pi_{fht}(a', \tau_t; \sigma_{-fht})/\lambda_\epsilon\right)}. \tag{A.20}$$

I iterate on (A.20) with damping until the strategy profile converges. The equilibrium probability of an action profile is

$$Pr^*_t(a_t|\tau_t) = \prod_{(f,h)} \sigma^*_{fht}(a_{fht}|\tau_t). \tag{A.21}$$

The full sourcing game is high dimensional. In the empirical implementation, I solve the game for a strategic core of firm–transmission pairs that are active on the make-or-buy margin and account for large supplier-level demand. Products outside the strategic core keep their observed sourcing assignment, but they continue to enter supplier demand and respond through the downstream pricing equilibrium. I report sensitivity to the strategic-core size, the action grid, assignment draws, and shock draws below.

B.5 Derivative of Sourcing Probabilities

The upstream first-order condition requires the derivative of the Stage-2 equilibrium action probabilities with respect to upstream prices. I compute this derivative analytically by differentiating the Stage-2 logit fixed point.

Let $\mu_\tau(g, a, s)$ denote the derivative of player g's action probability with respect to upstream price τ_s:

$$\mu_\tau(g, a, s) = \frac{\partial \sigma_g(a)}{\partial \tau_s}.$$

The derivative has a direct component, through the effect of τ_s on expected payoffs, and an equilibrium feedback component, through the effect of τ_s on rivals' strategy probabilities. In matrix form, for each supplier s, I solve a linear system of the form

$$\mu_{\tau,s} = J_\sigma \left[b_s + B\mu_{\tau,s} \right], \tag{A.22}$$

or equivalently,

$$(I - J_\sigma B)\mu_{\tau,s} = J_\sigma b_s. \tag{A.23}$$

Here J_σ is the block-diagonal Jacobian of logit choice probabilities with respect to expected values, b_s is the direct derivative of expected payoffs with respect to τ_s, and B collects how each player's expected payoff changes with rivals' choice probabilities. The resulting derivatives satisfy the adding-up restriction

$$\sum_{a \in \mathcal{A}_g} \mu_\tau(g, a, s) = 0$$

for every player g and supplier s.

B.6 Stage 1: Upstream Marginal-Cost Recovery

Given the Stage-2 equilibrium strategy and the expected supplier-level outsourced demand $Q_{st}(a_t, \tau_t)$, supplier s's expected demand is

$$ED_{st}(\tau_t) = \sum_{a_t} Pr_t^*(a_t|\tau_t) Q_{st}(a_t, \tau_t). \tag{A.24}$$

Its derivative is

$$\frac{\partial ED_{st}(\tau_t)}{\partial \tau_{\ell t}} = \sum_{a_t} \frac{\partial Pr_t^*(a_t|\tau_t)}{\partial \tau_{\ell t}} Q_{st}(a_t, \tau_t) + \sum_{a_t} Pr_t^*(a_t|\tau_t) \frac{\partial Q_{st}(a_t, \tau_t)}{\partial \tau_{\ell t}}. \tag{A.25}$$

The first term is the sourcing-probability response, and the second term is the downstream demand response holding the action profile fixed.

Under the supplier-year price-index restriction used in the empirical implementation, supplier s's expected profit is

$$\Pi_{st}^U(\tau_t) = (\tau_{st} - mc_{st}^U) ED_{st}(\tau_t). \tag{A.26}$$

The supplier first-order condition implies

$$mc_{st}^U = \tau_{st} + \frac{ED_{st}(\tau_t)}{\partial ED_{st}(\tau_t)/\partial \tau_{st}}. \tag{A.27}$$

I use (A.27) to recover a supplier-year effective upstream marginal cost. The recovered marginal cost is local to the observed equilibrium and should be interpreted under the maintained supplier-year price-index and linear-contract restrictions.

B.7 Counterfactual Upstream Repricing

In counterfactuals where upstream suppliers reoptimize prices, I hold the recovered supplier-year marginal costs fixed and solve the upstream first-order conditions. Given a candidate τ_t, I compute $ED_{st}(\tau_t)$ and $\partial ED_{st}(\tau_t)/\partial \tau_{st}$ using the steps above. The first-order condition implies the update

$$\tau_{st}^{BR} = mc_{st}^U - \frac{ED_{st}(\tau_t)}{\partial ED_{st}(\tau_t)/\partial \tau_{st}}. \tag{A.28}$$

I iterate on this update with damping,

$$\tau_t^{new} = (1 - \lambda)\tau_t + \lambda \tau_t^{BR}, \tag{A.29}$$

where $\lambda \in (0, 1]$. For the main counterfactuals, I use this procedure to compare equilibria with fixed upstream prices to equilibria in which upstream prices are reoptimized.

B.8 Simulation Specification and Sensitivity

The baseline simulation uses the strategic core described in the main text, a finite action grid for each strategic firm–transmission pair, balanced random assignment draws, and pre-generated demand and cost shock panels. I use common random numbers across counterfactuals. The main text reports the baseline specification. This appendix reports sensitivity to:

1. the number of strategic firm–transmission pairs;
2. the action grid;
3. the number of product-level assignment draws;
4. the number of demand and cost shock draws;

Active firm transmission pairs are those that changed the in-house proportions in my data sample. The active firm-transmission pair, which has the largest market share of each upstream firm, is defined as the strategic firm. I use the sensitivity test to see if I need to include the second largest firms. The set of strategic firm transmission pairs for each year are listed below. I additionally provide the sensitivity test for simulation specifications. The baseline simulation specification is five firm-transmission pairs. The action space is divided into five discrete choices adjusted to different action set. Then the discrete in-house proportions are $\{0, 0.25, 0.5, 0.75, 1\}$. Each firm-transmission pair would have a different choice set due to the data patterns.

From the sensitivity test table, one can see the importance of including the largest consumer(firm-transmission pair) for each upstream firm. However, the marginal gain is very small when moving to 6 firm-transmission pairs. In addition, adding more simulation draws for shocks and random assignment is also quantitatively less important. Using a coarser action grid has a small but non-negligible effect on the recovered marginal costs. This suggests that more computational effort should be devoted to refining the choice grid.

Table A.1: Sensitivity Test for Simulation Specifications

#Players	#Action	#Assignment (N)	#Shock (M)	mc_{Aisin}	mc_{ZF}	mc_{JATCO}	mc_{Other}
4	5	10	50	0.21%	-0.11%	-0.08%	0.36%
5	5	10	50				
6	5	10	50	-0.07%	0.11%	0.09%	-0.30%
5	4	10	50	-0.22%	0.01%	-0.03%	0.36%

Notes: This table reports the sensitivity test for the simulation specifications. The analysis is based on year 2018 and the reference group is the row in red. Column (5)-(8) shows the changes in marginal cost estimated compared to the reference group.